FTOS Configuration Guide

Version 7.5.1.0 August 2007

FORCE10™

Copyright 2007 Force10 Networks

Force10 Networks® reserves the right to change, modify, revise this publication without notice.

Trademarks

Force10 Networks® and E-Series® are registered trademarks of Force10 Networks, Inc. Force10, the Force10 logo, E1200, E600, E600i, E300, EtherScale, TeraScale, FTOS, C-Series, and S-Series are trademarks of Force10 Networks, Inc. All other brand and product names are registered trademarks or trademarks of their respective holders.

Statement of Conditions

In the interest of improving internal design, operational function, and/or reliability, Force10 Networks reserves the right to make changes to products described in this document without notice. Force10 Networks does not assume any liability that may occur due to the use or application of the product(s) described herein.

USA Federal Communications Commission (FCC) Statement

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designated to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy. If it is not installed and used in accordance to the instructions, it may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to take whatever measures necessary to correct the interference at their own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Force10 Networks is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications in the equipment. Unauthorized changes or modification could void the user’s authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Canadian Department of Communication Statement

The digital apparatus does not exceed the Class A limits for radio noise emissions from digital apparatus set out in the Radio Interference Regulations of the Canadian Department of Communications.

Attention: Le present appareil numerique n’ emet pas de perturbations radioelectriques depassant les normes applicables aux appareils numeriques de la Class A prescrites dans le Reglement sur les interferences radioelectriques etabli par le ministere des Communications du Canada.

European Union EMC Directive Conformance Statement

This product is in conformity with the protection requirements of EU Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. Force 10 Networks can not accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of this product, including the fitting of non-Force10 option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to CISPR 22/ European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Warning: This device is a Class A product. In a domestic environment, this device can cause radio interference, in which case, the user may be required to take appropriate measures.

VCCI Compliance for Class A Equipment (Japan)

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

This is Class A product based on the standard of the Voluntary Control Council For Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may arise. When such trouble occurs, the user may be required to take corrective actions.\

Danger: AC Power cords are for use with Force10 Networks equipment only, do not use Force10 Networks AC Power cords with any unauthorized hardware.

本製品に同梱いたしております電源コードセットは、本製品専用です。本電源コードセットは、本製品以外の製品ならびに他の用途でご使用いただくことは出来ません。製品本体には同梱された電源コードセットを使用し、他製品の電源コードセットを使用しないで下さい。

Feedback on Documentation?
Send email to techpubs@force10networks.com

New Features

Listed below are the new features for FTOS version 7.5.1.0:

Note: FTOS 7.5.1.0 is the first version that supports both the C-Series and E-Series product lines.

Layer 2

- CLI Option in **show ip bgp** Commands to Dump BGP Internal Database
- **clear frrp**
- Config/Show CLI to Field-upgrade SFM FPGA and Display Current Version
- Description Field for MSDP Neighbor
- Description Field for VRRP Neighbor
- Default MED in **show ip bgp**
- GVRP
- ISIS Throttling Enhancements to Control LSP Generation and SPF Calculations
- LACP Slow Mode
- Layer 2 Port-based QOS
- Mechanism to Capture Incoming and Outgoing PDUs per Peer
- Mtrace Originator
- Native VLAN
- Runtime Detection of MAC not Forwarding Traffic.
- **show console lp**
- Store Last Notification Sent/Received, Last Bad PDU Received

Layer 3

- CLI to change ECMP or LAG hash
- OSPF Ack Bundling
- RFC 2328-based LSA flooding

Management

- AAA Authorization Configuration Commands

- CLI:
    - **show console lp**
    - **show hardware poe-status**
- Force SSH to use v1 or v2
- OID for line card memory and CPU utilization
- VTY ACL Logging
- VTY MAC-SA Filter Support

System

- Banner Enhancement
- Config-save Timestamp
- HA between Primary and Secondary RPM
- netBSD on RPM
- SNMP-based Copy Config from Remote to Local
- SNMP trap on access

High Availability

- Move **show hardware phy** to **show interface** phy
- **show interface link-status**
- Auto-save of Command History and Hardware Log Trace Files upon Reload and Buffer Rollover

Contents

Chapter 2

Chapter 3

Chapter 4

Chapter 5

Chapter 6

Chapter 7

Chapter 8

Chapter 9

Chapter 14

Chapter 15

Chapter 16

Chapter 17

Chapter 18

Chapter 19

Chapter 20

Chapter 21

Chapter 22

Chapter 23

Chapter 24

Chapter 25

Chapter 26

Chapter 27

Chapter 28

Chapter 34

Chapter 35

Chapter 36
FTOS XML Feature 603

Chapter 37

Chapter 38

Chapter 39

Chapter 40
E-Series

Chapter 41

Appendix A

Appendix B

Appendix C

Appendix D

Appendix E

List of Figures

List of Tables

Preface About this Guide

Objectives

This guide describes the protocols and features supported by the Force10 Operating System (FTOS) and provides configuration instructions and examples for implementing them. It is a combined resource for the C-Series® and E-Series®.

Note: Due to a difference in hardware architecture and the rapid development of C-Series, features may occassionaly differ between C-Series and E-Series. These differences are identified by the information symbols on page 46.

Though this guide contains information on protocols, it is not intended to be a complete reference. This guide is a reference for configuring protocols on Force10 systems. For complete information on protocols, refer to other documentation including IETF Requests for Comment (RFCs). The instructions in this guide cite relevant RFCs, and Appendix D contains a complete list of the supported RFCs and Management Information Base files (MIBs).

Audience

This document is intended for system administrators who are responsible for configuring and maintaining networks and assumes you are knowledgeable in Layer 2 and Layer 3 networking technologies.

Conventions

This document uses the following conventions to describe command syntax:

Convention	Description
keyword	Keywords are in bold and should be entered in the CLI as listed.
parameter	Parameters are in italics and require a number or word to be entered in the CLI.
{X}	Keywords and parameters within braces must be entered in the CLI.
[X]	Keywords and parameters within brackets are optional.
x | y	Keywords and parameters separated by bar require you to choose one.

Information Symbols

Table 1 describes symbols contained in this guide.

Table 1 Information Symbols

Symbol	Warning	Description
⚡	Danger	This symbol warns you that improper handling and installation could result in bodily injury. Before you work on this equipment, be aware of electrical hazards and take appropriate safety precautions.
⛔	Caution	This symbol informs you that improper handling and installation could result in equipment damage or loss of data.
⚠	Warning	This symbol informs you that improper handling could reduce your component or system performance.
➡	Note	This symbol informs you of important operational information.
C-Series NO E-Series ✓	C-Series Only	This symbol informs you of a feature that is supported by the C-Series only, or of a difference between the C-Series and E-Series.
C-Series NO E-Series ✓	E-Series Only	This symbol informs you that a feature that is supported by the E-Series only.

Related Documents

For more information about the Force10 Networks E-Series and C-Series, refer to the following documents:

- *FTOS Command Line Reference*
- *E-Series Network Operations Guide*
- *Installing and Maintaining the <Force10 Chassis> System*
- *Release Notes*

Chapter 1 Configuration Fundamentals

Command Line Interface

The FTOS Command Line Interface (CLI) is a text-based interface through which you to configure interfaces and protocols. The CLI is largely the same for the E-Series and C-Series with the exception of some commands and command outputs. The CLI is structured in modes for security and management purposes. Different sets of commands are available in each mode, and you can limit user access to modes using privilege levels.

In FTOS, after a command is enabled, it is entered into the running configuration file. You can view the current configuration for the whole system or for a particular CLI mode. To save the current configuration copy the running configuration to another location.

➡ **Note:** Some features, commands, and command outputs may occassionaly differ between C-Series and E-Series due to a difference in hardware architecture and the rapid development of C-Series. Differences are identified by the information symbols on page 46.

Accessing the Command Line

Access the E-Series commands through a serial console port or a Telnet session (Figure 1). When the system successfully boots, you enter the command line in the EXEC mode.

➡ **Note:** You must have a password configured on a virtual terminal line before you can Telnet into the system. Therefore, you must use a console connection when connecting to the system for the first time.

Figure 1 Logging into the System using Telnet

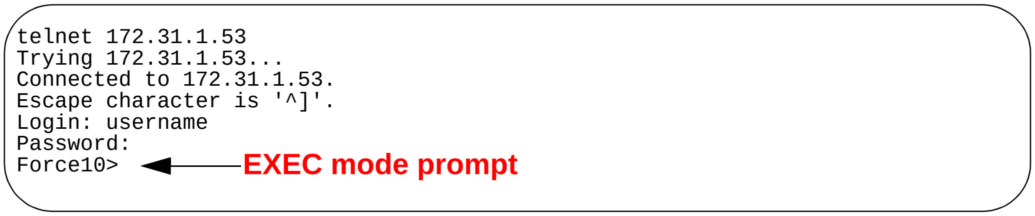

```
telnet 172.31.1.53
Trying 172.31.1.53...
Connected to 172.31.1.53.
Escape character is '^]'.
Login: username
Password:
Force10>
```

EXEC mode prompt

CLI Modes

Different sets of commands are available in each mode. A command found in one mode cannot be executed from another mode (with the exeption of EXEC mode commands preceded by the command **do**; see The do Command on page 52). You can set user access rights to commands and command modes using privilege levels; for more information on privilege levels and security options, refer to Chapter 6, Security, on page 127.

The FTOS CLI is divided into three major mode levels:

- **EXEC mode** is the default mode and has a privilege level of 1, which is the most restricted level. Only a limited selection of commands is available, notably **show** commands, which allow you to view system information.
- **EXEC privilege** has commands to view configurations, clear counters, manage configuration files, run diagnostics, and enable or disable debug operations. The privilege level is 15, which is unrestricted. You can configure a password for this mode; see Configuring the Enable Password on page 63.
- **CONFIGURATION mode** enables you to configure security features, time settings, set logging and SNMP functions, configure static ARP and MAC addresses, and set line cards on the system.

Beneath CONFIGURATION mode are sub-modes for interfaces, protocols, and features. Figure 2 illustrates the command mode structure. Two sub-CONFIGURATION modes are important when configuring the chassis for the first time:

- **INTERFACE mode** is the mode in which you configure Layer 2 and Layer 3 protocols and IP services specific to an interface. An interface can be physical (Management interface, 1-Gigabit Ethernet, 10-Gigabit Ethernet, or SONET) or logical (Loopback, Null, port channel, or VLAN).
- **LINE mode** is the mode in which you to configure the console and virtual terminal lines.

Figure 2 CLI Modes in FTOS

- AS-PATH ACL
- FVRP
- INTERFACE
    - GIGABIT ETHERNET
    - 10 GIGABIT ETHERNET
    - INTERFACE RANGE
    - LOOPBACK
    - MANAGEMENT ETHERNET
    - NULL
    - PORT-CHANNEL
    - SONET
    - VLAN
    - VRRP
- IP COMMUNITY-LIST
- IP ACCESS-LIST
    - STANDARD ACCESS-LIST
    - EXTENDED ACCESS-LIST
- LINE
    - AUXILIARY
    - CONSOLE
    - VIRTUAL TERMINAL
- MAC ACCESS-LIST
- MULTIPLE SPANNING TREE
- Per-VLAN SPANNING TREE
- PREFIX-LIST
- RAPID SPANNING TREE
- REDIRECT
- ROUTE-MAP
- ROUTER BGP
- ROUTER ISIS
- ROUTER OSPF
- ROUTER RIP
- SPANNING TREE
- TRACE-LIST

fnEC001mp

Navigating CLI Modes

The FTOS prompt changes to indicate the CLI mode. Table 2 lists the CLI mode, its prompt, and information on how to access and exit this CLI mode. You must move linearly through the command modes, with the exception of the **end** command which takes you directly to EXEC privilege mode; the **exit** command moves you up one command mode level.

Note: Sub-CONFIGURATION modes all have the letter “conf” in the prompt with addtional modifiers to identify the mode and slot/port information.

Table 2 FTOS Command Modes

CLI Command Mode	Prompt	Access Command	Exit Command
EXEC	Force10>	Access the router through the console or Telnet.	

Table 2 FTOS Command Modes

CLI Command Mode		Prompt	Access Command	Exit Command
EXEC privilege		Force10#	- From EXEC mode, enter the command **enable**. - From any other mode, use the command **end**.	
CONFIGURATION		Force10(conf)#	- From EXEC privilege mode, enter the command **configure**. - From every mode except EXEC and EXEC privilege, enter the command **exit**.	
➡	**Note:** Access all of the following modes from CONFIGURATION mode.			
	AS-PATH ACL	Force10(config-as-path)#	**ip as-path access-list**	
	FVRP	Force10(config-fvrp)#	**protocol FVRP**	
INTERFACE modes	Gigabit Ethernet Interface	Force10(conf-if-gi-0/0)#	**interface**	
	10 Gigabit Ethernet Interface	Force10(conf-if-te-0/0)#		
	Interface Range	Force10(conf-if-range)#		
	Loopback Interface	Force10(conf-if-lo-0)#		
	Management Ethernet Interface	Force10(conf-if-ma-0/0)#		
	Null Interface	Force10(conf-if-nu-0)#		
	Port-channel Interface	Force10(conf-if-po-0)#		
	SONET Interface	Force10(conf-if-so-0/0)#		
	VLAN Interface	Force10(conf-if-vl-0)#		
IP ACCESS-LIST	STANDARD ACCESS- LIST	Force10(config-ext-nacl)#	**ip access-list standard**	
	EXTENDED ACCESS- LIST	Force10(config-std-nacl)#	**ip access-list extended**	
	IP COMMUNITY-LIST	Force10(config-community-list)#	**ip community-list**	
LINE	AUXILIARY	Force10(config-line-aux)#	**line**	
	CONSOLE	Force10(config-line-console)#		
	VIRTUAL TERMINAL	Force10(config-line-vty)#		

Table 2 FTOS Command Modes

CLI Command Mode		Prompt	Access Command	Exit Command
MAC ACCESS-LIST	STANDARD ACCESS- LIST	Force10(config-std-macl)#	**mac access-list standard**	
	EXTENDED ACCESS- LIST	Force10(config-ext-macl)#	**mac access-list extended**	
	MULTIPLE SPANNING TREE	Force10(config-mstp)#	**protocol spanning-tree mstp**	
	Per-VLAN SPANNING TREE Plus	Force10(config-pvst)#	**protocol spanning-tree pvst**	
	PREFIX-LIST	Force10(conf-nprefixl)#	**ip prefix-list**	
	RAPID SPANNING TREE	Force10(config-rstp)#	**protocol spanning-tree rstp**	
	REDIRECT	Force10(conf-redirect-list)#	**ip redirect-list**	
	ROUTE-MAP	Force10(config-route-map)#	**route-map**	
	ROUTER BGP	Force10(conf-router_bgp)#	**router bgp**	
	ROUTER ISIS	Force10(conf-router_isis)#	**router isis**	
	ROUTER OSPF	Force10(conf-router_ospf)#	**router ospf**	
	ROUTER RIP	Force10(conf-router_rip)#	**router rip**	
	SPANNING TREE	Force10(config-span)#	**protocol spanning-tree 0**	
	TRACE-LIST	Force10(conf-trace-acl)#	**ip trace-list**	

Figure 3 illustrates how to change the command mode from CONFIGURATION mode to PROTOCOL SPANNING TREE.

Figure 3 Changing CLI Modes

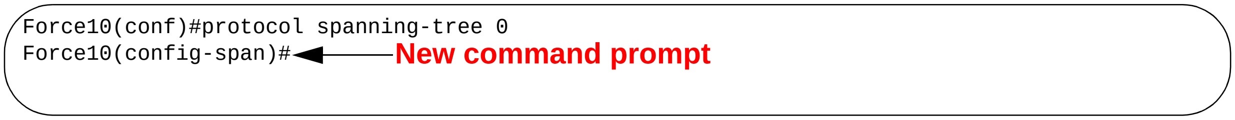

```
Force10(conf)#protocol spanning-tree 0
Force10(config-span)#
```

The do Command

Enter an EXEC mode command from any CONFIGURATION mode (CONFIGURATION, INTERFACE, SPANNING TREE, etc.) without returning to EXEC mode by preceding the EXEC mode command with the command **do**. Figure 4 illustrates the **do** command.

Note: The following commands cannot be modified by the **do** command: **enable, disable, exit**, and **configure**.

Figure 4 Using the **do** Command

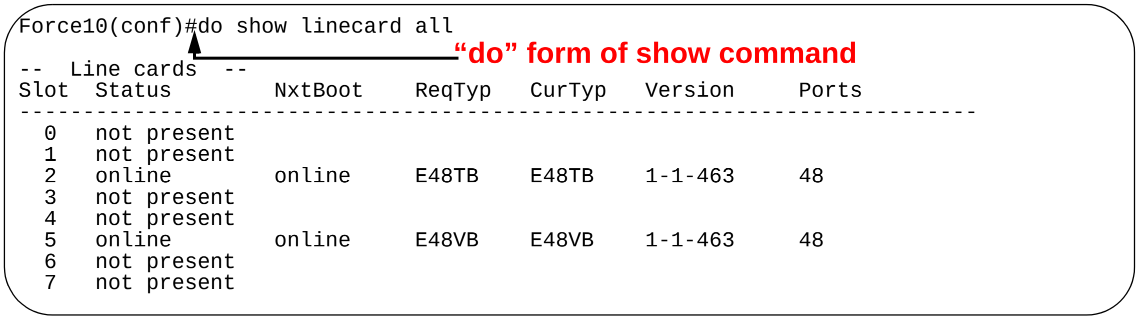

```
Force10(conf)#do show linecard all

--  Line cards  --
Slot  Status        NxtBoot    ReqTyp   CurTyp   Version     Ports
---------------------------------------------------------------------------
  0   not present
  1   not present
  2   online        online     E48TB    E48TB    1-1-463     48
  3   not present
  4   not present
  5   online        online     E48VB    E48VB    1-1-463     48
  6   not present
  7   not present
```

Undoing Commands

When you enter a command, the command line is added to the running configuration file. Disable a command and remove it from the running-config by entering the original command preceded by the command **no**. For example, to delete an ip address configured on an interface, use the **no ip address** *ip-address* command, as shown in Figure 5.

Note: Use the **help** command to help you construct the “no” form of a command.

Figure 5 Undoing a command with the **no** Command

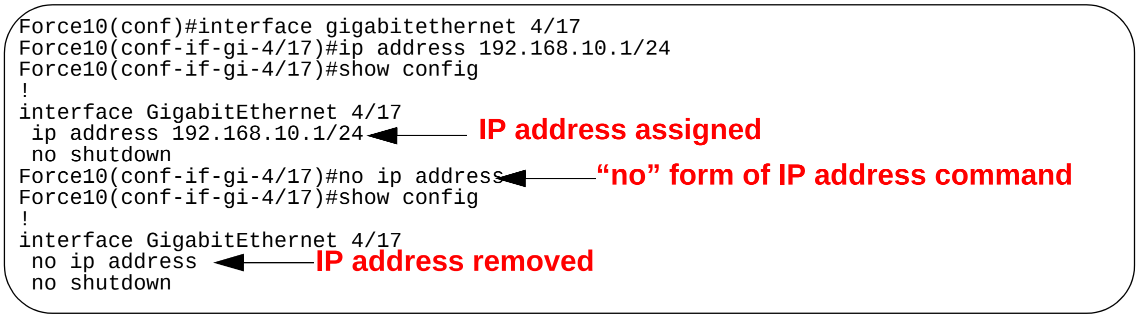

```
Force10(conf)#interface gigabitethernet 4/17
Force10(conf-if-gi-4/17)#ip address 192.168.10.1/24
Force10(conf-if-gi-4/17)#show config
!
interface GigabitEthernet 4/17
 ip address 192.168.10.1/24
 no shutdown
Force10(conf-if-gi-4/17)#no ip address
Force10(conf-if-gi-4/17)#show config
!
interface GigabitEthernet 4/17
 no ip address
 no shutdown
```

IP address assigned

“no” form of IP address command

IP address removed

Layer 2 protocols are disabled by default. Enable them using the **no disable** command. For example, in PROTOCOL SPANNING TREE mode, enter **no disable** to enable Spanning Tree.

Obtaining Help

Obtain a list of keywords and a brief functional description of those keywords at any CLI mode using the **?** or **help** command.

- Enter **?** at the prompt or after a keyword to list the keywords available in the current mode.
    - **?** after a prompt lists all of the available keywords. The output of this command is the same for the **help** command.

Figure 6 ? Command Example

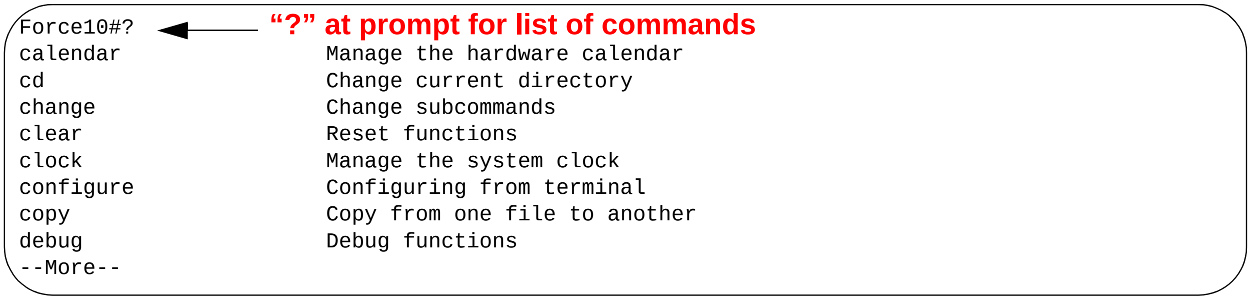

```
Force10#?
calendar                Manage the hardware calendar
cd                      Change current directory
change                  Change subcommands
clear                   Reset functions
clock                   Manage the system clock
configure               Configuring from terminal
copy                    Copy from one file to another
debug                   Debug functions
--More--
```

“?” at prompt for list of commands

- **?** after a partitial keyword lists all of the keywords that begin with the specified letters.

Figure 7 Keyword? Command Example

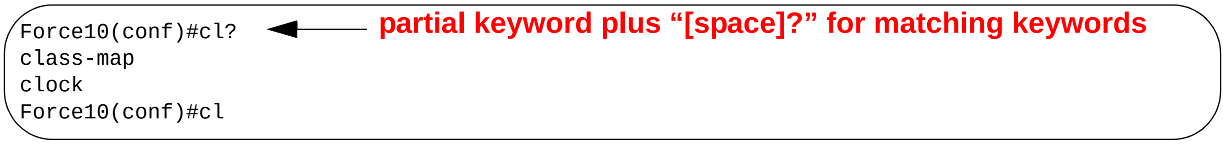

```
Force10(conf)#cl?
class-map
clock
Force10(conf)#cl
```

partial keyword plus “[space]?” for matching keywords

- A keyword followed by [space]**?** lists all of the keywords that can follow the specified keyword.

Figure 8 Keyword ? Command Example

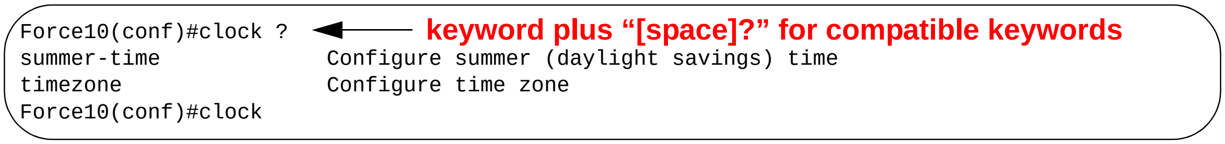

```
Force10(conf)#clock ?
summer-time             Configure summer (daylight savings) time
timezone                Configure time zone
Force10(conf)#clock
```

keyword plus “[space]?” for compatible keywords

Entering and Editing Commands

When entering commands:

- The CLI is not case sensitive.
- You can enter partial CLI keywords.
    - You must enter the minimum number of letters to uniquely identify a command. For example, **cl** cannot be entered as a partial keyword because both the **clock** and **class-map** commands begin with the letters “cl.” **clo**, however, can be entered as a partial keyword because only one command begins with those three letters.
- The TAB key auto-completes keywords in commands. You must enter the minimum number of letters to uniquely identify a command.
- The UP and DOWN arrow keys display previously entered commands (see Command History).
- The BACKSPACE and DELETE keys erase the previous letter.
- Key combinations are available to move quickly across the command line, as described in Table 3.

Table 3 Short-Cut Keys and their Actions

Key Combination	Action
CNTL-A	Moves the cursor to the beginning of the command line.
CNTL-B	Moves the cursor back one character.
CNTL-D	Deletes character at cursor.
CNTL-E	Moves the cursor to the end of the line.
CNTL-F	Moves the cursor forward one character.
CNTL-I	Completes a keyword.
CNTL-K	Deletes all characters from the cursor to the end of the command line.
CNTL-L	Re-enters the previous command.
CNTL-N	Return to more recent commands in the history buffer after recalling commands with CTRL-P or the UP arrow key.
CNTL-P	Recalls commands, beginning with the last command
CNTL-R	Re-enters the previous command.
CNTL-U	Deletes the line.
CNTL-W	Deletes the previous word.
CNTL-X	Deletes the line.
CNTL-Z	Ends continuous scrolling of command outputs.
Esc B	Moves the cursor back one word.

Table 3 Short-Cut Keys and their Actions (continued)

Key Combination	Action
Esc F	Moves the cursor forward one word.
Esc D	Deletes all characters from the cursor to the end of the word.

Command History

FTOS maintains a history of previously-entered commands for each mode. For example:

- When you are in EXEC mode, the UP and DOWN arrow keys display only the previously-entered EXEC mode commands.
- When you are in CONFIGURATION mode, the UP or DOWN arrows keys recall only the previously-entered CONFIGURATION mode commands.

Filtering show Command Outputs

Filter the output of a **show** command to display specific information by adding **|** [**except** | **find** | **grep |** **no-more**] *specified_text* after the command. *specified_text* is the text for which you are filtering, and it is case sensitive.

Note: FTOS accepts a space or no space before and after the pipe. To filter on a phrase with spaces, underscores, or ranges, enclose the phrase with double quotation marks.

- **except** displays text that does not match the specified text. Figure 9 shows this command used in combination with the command **show linecard all**.

Figure 9 Filtering Command Outputs with the **except** Command

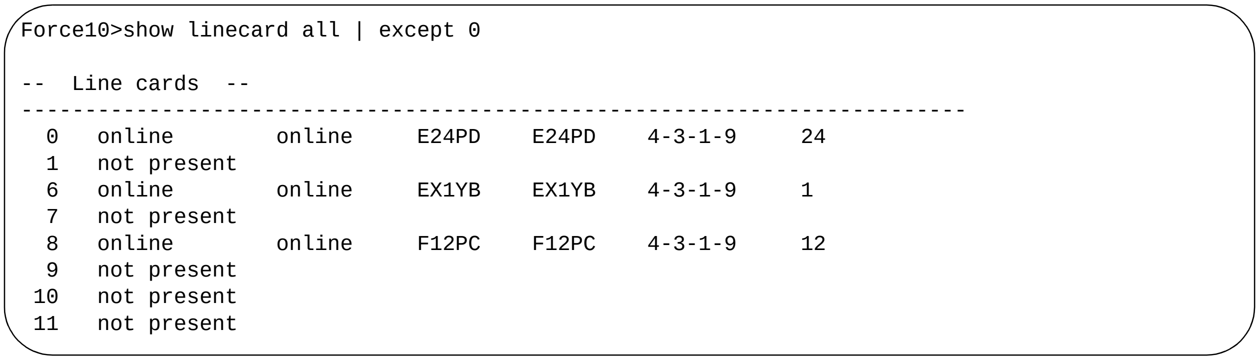

```
Force10>show linecard all | except 0

--  Line cards  --
--------------------------------------------------------------------------
  0   online        online     E24PD     E24PD     4-3-1-9     24
  1   not present
  6   online        online     EX1YB     EX1YB     4-3-1-9     1
  7   not present
  8   online        online     F12PC     F12PC     4-3-1-9     12
  9   not present
 10   not present
 11   not present
```

- **find** displays the output of the show command beginning from the first occurrence of specified text Figure 10 shows this command used in combination with the command **show linecard all**.

Figure 10 Filtering Command Outputs with the **find** Command

```
Force10(conf)#do show linecard all | find 0
  0   not present
  1   not present
  2   online        online     E48TB    E48TB    1-1-463     48
  3   not present
  4   not present
  5   online        online     E48VB    E48VB    1-1-463     48
  6   not present
  7   not present
```

- **grep** displays only the lines containing specified text. Figure 11 shows this command used in combination with the command **show linecard all**.

Figure 11 Filtering Command Outputs with the **grep** Command

```
Force10(conf)#do show linecard all | grep 0
  0   not present
```

- **display** displays additional configuration information.
- **no-more** displays the output all at once rather than one screen at a time. This is similar to the command **terminal length** except that the **no-more** option affects the output of the specified command only.

Note: You can filter a single command output multiple times. For example:

Force10# *command* | **grep** *regular-expression* | **except** *regular-expression* | **grep** *other-regular-expression* | **find** *regular-expression* | **no-more**

Multiple Users in Configuration mode

FTOS notifies all users in the event that there are multiple users logged into CONFIGURATION mode. A warning message indicates the username, type of connection (console or vty), and in the case of a vty connection, the IP address of the terminal on which the connection was established. For example:

- On the system that Telnets into the switch, Message 1 appears:

Message 1 Multiple Users in Configuration mode Telnet Message

```
% Warning: The following users are currently configuring the system:
User "<username>" on line console0
```

- On the system that is connected over the console, Message 2 appears:

Message 2 Multiple Users in Configuration mode Telnet Message

```
% Warning: User "<username>" on line vty0 "10.11.130.2" is in configuration mode
```

If either of these messages appears, Force10 recommends that you coordinate with the users listed in the message so that you do not unintentionally overwrite each other’s configuration changes.

Determining the Chassis Mode

The chassis mode in FTOS determines which hardware is being supported in a chassis. The chassis mode is programmed into an EEPROM on the backplane of the chassis, and the change takes place only after the chassis is rebooted. Configuring the appropriate chassis mode enables the system to use all the ports on the card and recognize all software features.

For more information on the command, see the *FTOS Configuration Guide*.

Chapter 2 Getting Started

When you power up the chassis, the system performs a power-on self-test (POST) during which Route Processor Modules (RPMs), switch fabric modules (SFMs), and line cards status LEDs blink green. The system then load FTOS; boot messages scroll up the terminal during this process. No user interaction is required if the boot process proceeds without interruption.

When the boot process is complete, the RPM and line card status LEDs remain online (green), and the console monitor displays the Force10 banner and EXEC mode prompt, as shown in Figure 12.

Figure 12 Completed Boot Process

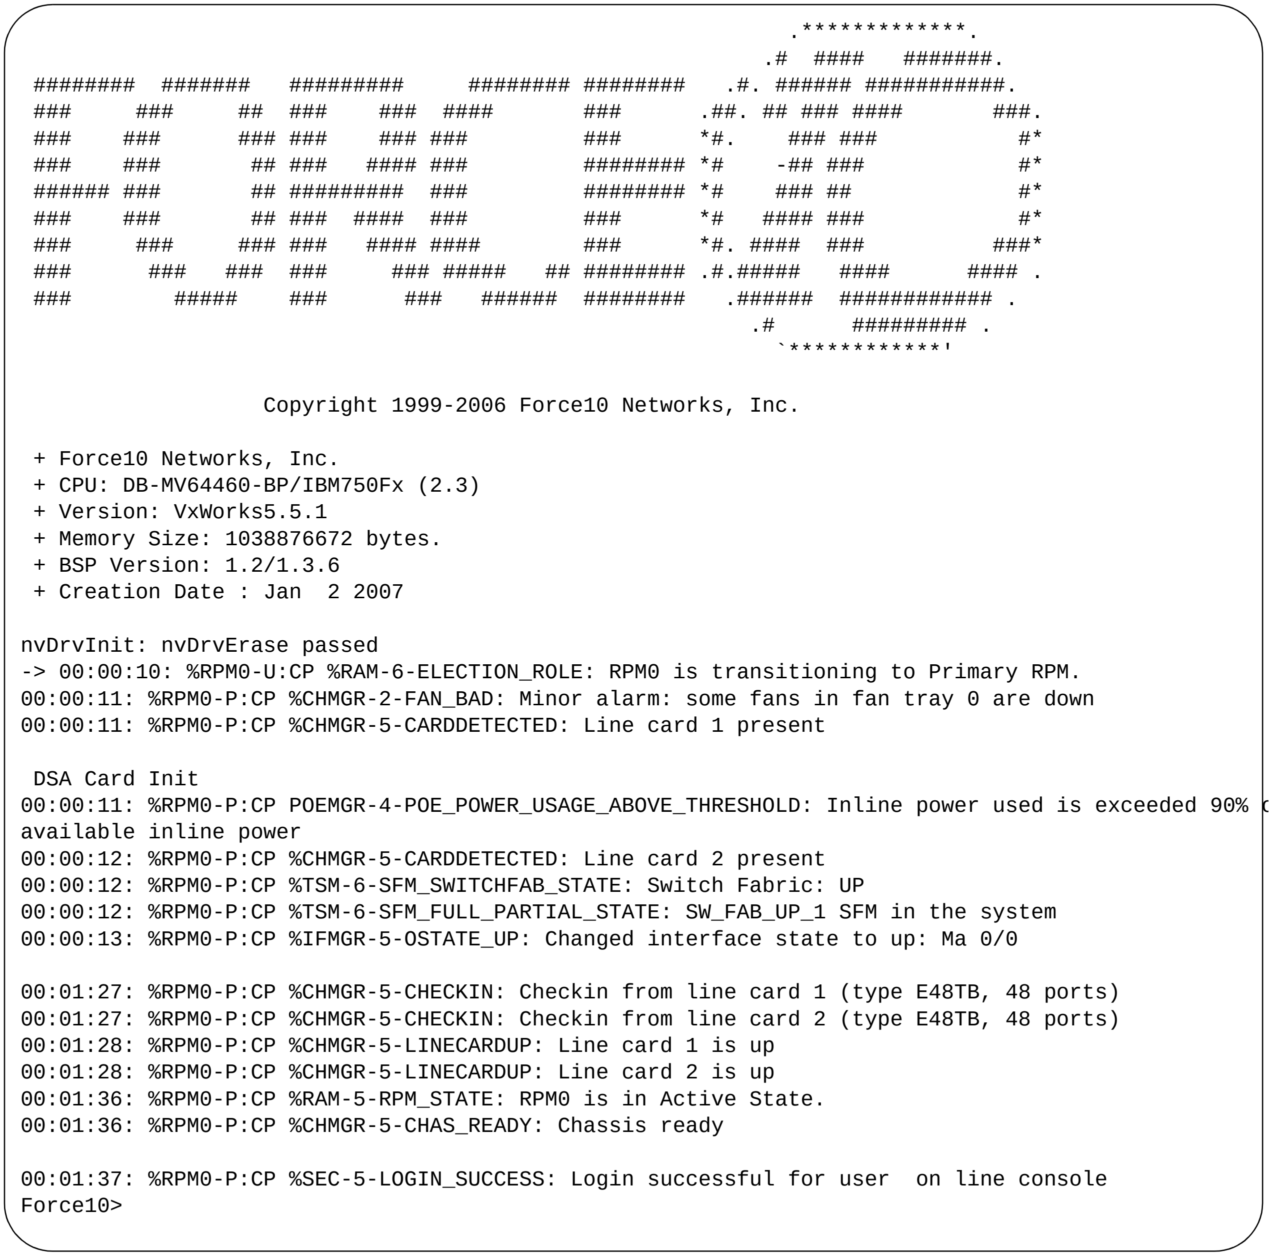

```
                                                            .*************.
                                                          .#  ####   #######.
 ########  #######   #########     ######## ########   .#. ###### ###########.
 ###     ###     ##  ###    ###  ####       ###      .##. ## ### ####       ###.
 ###    ###      ### ###    ### ###         ###      *#.    ### ###           #*
 ###    ###       ## ###   #### ###         ######## *#    -## ###            #*
 ###### ###       ## #########  ###         ######## *#    ### ##             #*
 ###    ###       ## ###  ####  ###         ###      *#   #### ###            #*
 ###     ###     ### ###   #### ####        ###      *#. ####  ###          ###*
 ###      ###   ###  ###     ### #####   ## ######## .#.#####   ####      #### .
 ###        #####    ###      ###   ######  ########   .######  ############ .
                                                         .#      ######### .
                                                           `************'


                  Copyright 1999-2006 Force10 Networks, Inc.

 + Force10 Networks, Inc.
 + CPU: DB-MV64460-BP/IBM750Fx (2.3)
 + Version: VxWorks5.5.1
 + Memory Size: 1038876672 bytes.
 + BSP Version: 1.2/1.3.6
 + Creation Date : Jan  2 2007

nvDrvInit: nvDrvErase passed
-> 00:00:10: %RPM0-U:CP %RAM-6-ELECTION_ROLE: RPM0 is transitioning to Primary RPM.
00:00:11: %RPM0-P:CP %CHMGR-2-FAN_BAD: Minor alarm: some fans in fan tray 0 are down
00:00:11: %RPM0-P:CP %CHMGR-5-CARDDETECTED: Line card 1 present

 DSA Card Init
00:00:11: %RPM0-P:CP POEMGR-4-POE_POWER_USAGE_ABOVE_THRESHOLD: Inline power used is exceeded 90% o
available inline power
00:00:12: %RPM0-P:CP %CHMGR-5-CARDDETECTED: Line card 2 present
00:00:12: %RPM0-P:CP %TSM-6-SFM_SWITCHFAB_STATE: Switch Fabric: UP
00:00:12: %RPM0-P:CP %TSM-6-SFM_FULL_PARTIAL_STATE: SW_FAB_UP_1 SFM in the system
00:00:13: %RPM0-P:CP %IFMGR-5-OSTATE_UP: Changed interface state to up: Ma 0/0

00:01:27: %RPM0-P:CP %CHMGR-5-CHECKIN: Checkin from line card 1 (type E48TB, 48 ports)
00:01:27: %RPM0-P:CP %CHMGR-5-CHECKIN: Checkin from line card 2 (type E48TB, 48 ports)
00:01:28: %RPM0-P:CP %CHMGR-5-LINECARDUP: Line card 1 is up
00:01:28: %RPM0-P:CP %CHMGR-5-LINECARDUP: Line card 2 is up
00:01:36: %RPM0-P:CP %RAM-5-RPM_STATE: RPM0 is in Active State.
00:01:36: %RPM0-P:CP %CHMGR-5-CHAS_READY: Chassis ready

00:01:37: %RPM0-P:CP %SEC-5-LOGIN_SUCCESS: Login successful for user  on line console
Force10>
```

Default Configuration

A version of FTOS is pre-loaded onto the chassis, however the system is not configured when you power up for the first time (except for the default hostname, which is Force10). You must configure the system using commands.

Configuring a Host Name

The host name appears in the prompt. The default host name is **force10**.

- Host names must start with a letter and end with a letter or digit.
- Characters within the string can be letters, digits, and hyphens.

To configure a host name:

Step	Task	Command Syntax	Command Mode
1	Create a new host name.	**hostname** *name*	CONFIGURATION

Figure 13 illustrates the **hostname** command.

Figure 13 Configuring a Hostname

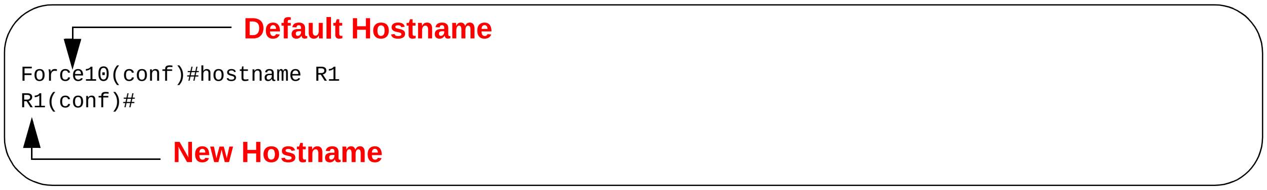

```
Force10(conf)#hostname R1
R1(conf)#
```

Accessing the System Remotely

You can Telnet to the management port to access the system remotely. Configuring the system for Telnet is a three-step process:

1. Configure an IP address for the management port. See page 61.
2. Configure a managment route with a default gateway. See page 62.
3. Configure a username and password. See page 62.

Configuring the Management Port IP Address

You must assign IP addresses to the management ports in order to access the system remotely.

Note: Assign different IP addresses to the management port of each RPM.

To configure the management port IP address:

Step	Task	Command Syntax	Command Mode
1	Enter INTERFACE mode for the Management port.	**interface ManagementEthernet** *slot/port* • *slot* range: 0 to 1 • *port* range: 0	CONFIGURATION
2	Assign an IP address to the interface.	**ip address** *ip-address/mask* • *ip-address:* an address in dotted-decimal format (A.B.C.D). • *mask:* a subnet mask in /prefix-length format (/xx).	INTERFACE
3	Enable the interface.	**no shutdown**	INTERFACE

Configuring a Management Route

Define a path from the system to the network from which you are accessing the system remotely. Management routes are separate from IP routes and are only used to manage the system via the management port.

To configure a management route:

Step	Task	Command Syntax	Command Mode
1	Configure a management route to the network from which you are accessing the system.	**management route** *ip-address*/*mask gateway* • *ip-address:* the network address in dotted-decimal format (A.B.C.D). • *mask:* a subnet mask in /prefix-length format (/xx). • *gateway*: the next hop for network traffic originating from the management port.	CONFIGURATION

Configuring a Username and Password

You must configure a system username and password to access the system remotely.

To configure a username and password:

Step	Task	Command Syntax	Command Mode
1	Configure a username and password to access the system remotely.	**username** *username* **password** [*encryption-type*] *password* *encryption-type* specifes how you are inputting the password, is 0 by default, and is not required. • 0 is for inputting the password in clear text. • 7 is for inputting a password that is already encrypted using a Type 7 hash. Obtaining the encrypted password from the configuration of another Force10 system.	CONFIGURATION

Configuring the Enable Password

Access EXEC Privilege mode using the command **enable**. This mode is unrestricted by default; configure it with a password as a basic security measure. There are two types of enable passwords:

- **enable password** stores the password in the running/startup configuration using a Type 7 encryption method.
- **enable secret** is stored in the running/startup configuration in using a stronger, MD5 encryption method.

Force10 recommends using the enable secret password.

To configure an enable password:

Step	Task	Command Syntax	Command Mode
1	Create a password to access EXEC Privilege mode.	**enable** [**password** | **secret**] [**level** *level*] [*encryption-type*] *password* *level* is the privilege level, is 15 by default, and is not required. *encryption-type* specifes how you are inputting the password, is 0 by default, and is not required. • 0 is for inputting the password in clear text. • 7 is for inputting a password that is already encrypted using a Type 7 hash. Obtaining the encrypted password from the configuration file of another Force10 system. • 5 is for inputting a password that is already encrypted using a Type 7 hash. Obtaining the encrypted password from the configuration file of another Force10 system.	CONFIGURATION

Configuration File Management

Files can be stored on and accessed from various storage media. Rename, delete, and copy files on the system from EXEC Privilege mode.

The EtherScale platform architecture uses MMC cards for both the internal and external Flash memory. MMC cards support a maximum of 100 files. The TeraScale platform architecture uses Compact Flash for the internal and external Flash memory. It has a space limitation but does not limit the number of files it can contain.

Copying Files to and from the System

The command syntax for copying files is similar to UNIX. The **copy** command uses the format **copy** *file-origin file-destination*.

Note: See the FTOS Command Line Reference for a detailed description of the command **copy**.

- To copy a local file to a remote system, use Table 4 combine the *file-origin* syntax for a local file location with the *file-destination* syntax for a remote file location.
- To copy a remote file to Force10 system, cuse Table 4 combine the *file-origin* syntax for a remote file location with the *file-destination* syntax for a local file location.

Table 4 Forming a copy Command

	file-origin **Syntax**	*file-destination* **Syntax**
Local File Location		
Internal flash:		
primary RPM	**copy flash://***filename*	**flash://***filename*
standby RPM	**copy rpm{0|1}flash://***filename*	**rpm{0|1}flash://***filename*
External flash:		
primary RPM	**copy rpm{0|1}slot0://***filename*	**rpm{0|1}slot0://***filename*
standby RPM	**copy rpm{0|1}slot0://***filename*	**rpm{0|1}slot0://***filename*
Remote File Location		
FTP server	**copy ftp://***username*:*password*@{*hostip* | *hostname*}/*filepath*/*filename*	**ftp://***username*:*password*@{*hostip* | *hostname*}/*filepath*/*filename*
TFTP server	**copy tftp://**{*hostip* | *hostname*}/*filepath*/*filename*	**tftp://**{*hostip* | *hostname*}/*filepath*/*filename*
SCP server	**copy scp://**{*hostip* | *hostname*}/*filepath*/*filename*	**scp://**{*hostip* | *hostname*}/*filepath*/*filename*

Important Points to Remember

- You may not copy a file from one remote system to another.
- You may not copy a file from one location to the same location.
- The internal flash memories on the RPMs are synchronized whenever there is a change, but only if the RPMs are running the same version of FTOS.
- When copying to a server, a hostname can only be used if a DNS server is configured.

Figure 14 shows an example of using the command **copy** to save a file to an FTP server.

Figure 14 Saving a file to a Remote System

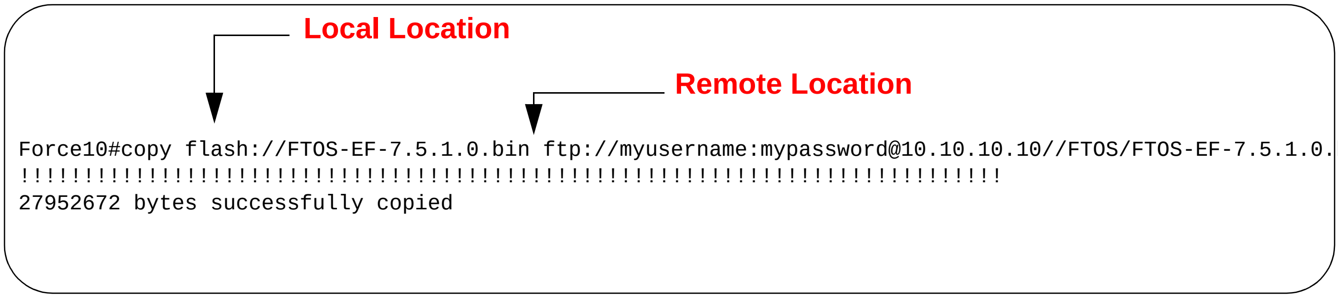

```
Force10#copy flash://FTOS-EF-7.5.1.0.bin ftp://myusername:mypassword@10.10.10.10//FTOS/FTOS-EF-7.5.1.0.
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
27952672 bytes successfully copied
```

Figure 15 shows an example of using the command **copy** to import a file to the Force10 system from an an FTP server.

Figure 15 Saving a file to a Remote System

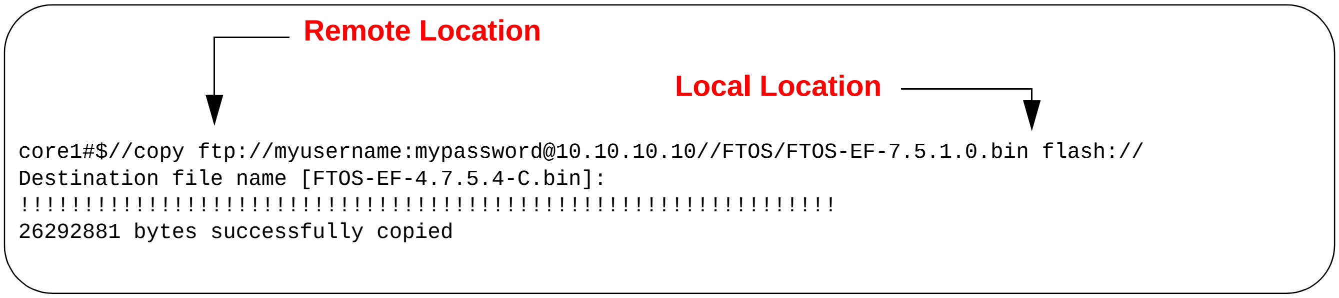

```
core1#$//copy ftp://myusername:mypassword@10.10.10.10//FTOS/FTOS-EF-7.5.1.0.bin flash://
Destination file name [FTOS-EF-4.7.5.4-C.bin]:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
26292881 bytes successfully copied
```

Saving the Running-configuration

The running-configuration contains the current system configuration. Force10 recommends that you copy your running-configuration to the startup-configuration. The system uses the startup-configuration during boot-up to configure the system. The startup-configuration is stored in the internal flash on the primary RPM by default, but it can be saved onto an external flash (on an RPM) or a remote server.

To save the running-configuration:

➡ **Note:** The commands in this section follow the same format as those in Copying Files to and from the System on page 64 but use the filenames *startup-configuration* and *running-configuration*. These commands assume that current directory is the internal flash, which is the system default.

Task	Command Syntax	Command Mode
Save the running-configuration to:		EXEC Privilege
the startup-configuration on the internal flash of the primary RPM	**copy running-config startup-config**	
the internal flash on an RPM	**copy running-config rpm**{0|1}**flash://startup-config**	
Note: The internal flash memories on the RPMs are synchronized whenever there is a change, but only if the RPMs are running the same version of FTOS.		
the startup-configuration on the external flash of an RPM	**copy running-config rpm**{0|1}**slot0://startup-config**	
an FTP server	**copy running-config ftp://** *username*:*password*@{*hostip* | *hostname*}/*filepath*/**running-config**	
a TFTP server	**copy running-config tftp://**{*hostip* | *hostname*}/*filepath*/**running-config**	
an SCP server	**copy running-config scp://**{*hostip* | *hostname*}/*filepath*/**running-config**	
Note: When copying to a server, a hostname can only be used if a DNS server is configured.		
Save the running-configuration to the startup-configuration on the internal flash of the primary RPM. Then copy the new startup-config file to the external flash of the primary RPM.	**copy running-config startup-config duplicate**	EXEC Privilege

Viewing Files

File information and content can only be viewed on local file systems.

To view a list of files on the internal or external Flash:

Step	Task	Command Syntax	Command Mode
1	View a list of files on:		
	the internal flash of an RPM	**dir flash:**	EXEC Privilege
	the external flash of an RPM	**dir slot:**	

The output of the command **dir** also shows the read/write privileges, size (in bytes), and date of modification for each file, as shown in Figure 16.

Figure 16 Viewing a List of Files in the Internal Flash

```
Force10#dir
Directory of flash:

  1  drw-      32768   Jan 01 1980 00:00:00  .
  2  drwx        512   Jul 23 2007 00:38:44  ..
  3  drw-       8192   Mar 30 1919 10:31:04  TRACE_LOG_DIR
  4  drw-       8192   Mar 30 1919 10:31:04  CRASH_LOG_DIR
  5  drw-       8192   Mar 30 1919 10:31:04  NVTRACE_LOG_DIR
  6  drw-       8192   Mar 30 1919 10:31:04  CORE_DUMP_DIR
  7  d---       8192   Mar 30 1919 10:31:04  ADMIN_DIR
  8  -rw-   33059550   Jul 11 2007 17:49:46  FTOS-EF-7.4.2.0.bin
  9  -rw-   27674906   Jul 06 2007 00:20:24  FTOS-EF-4.7.4.302.bin
 10  -rw-   27674906   Jul 06 2007 19:54:52  boot-image-FILE
 11  drw-       8192   Jan 01 1980 00:18:28  diag
 12  -rw-       7276   Jul 20 2007 01:52:40  startup-config.bak
 13  -rw-       7341   Jul 20 2007 15:34:46  startup-config
 14  -rw-   27674906   Jul 06 2007 19:52:22  boot-image
 15  -rw-   27674906   Jul 06 2007 02:23:22  boot-flash
--More--
```

To view the contents of a file:

Step	Task	Command Syntax	Command Mode
1	View the:		
	contents of a file in the internal flash of an RPM	**show file rpm{0|1}flash://*filename***	EXEC Privilege
	contents of a file in the external flash of an RPM	**show file rpm{0|1}slot0://*filename***	
	running-configuration	**show running-config**	
	startup-configuration	**show startup-config**	

File System Management

The Force10 system can use the internal Flash, external Flash, or remote devices to store files. It stores files on the internal Flash by default but can be configured to store files elsewhere.

To view file system information:

Step	Task	Command Syntax	Command Mode
1	View information about each file system.	**show file-systems**	EXEC Privilege

The output of the command **show file-systems** (Figure 17) shows the total capacity, amount of free memory, file structure, media type, read/write privileges for each storage device in use.

Figure 17 show file-system Command Example

```
Force10#show file-systems
Size(b)     Free(b)      Feature      Type   Flags  Prefixes
    520962048   213778432     dosFs2.0 USERFLASH      rw  flash:
    127772672    21936128     dosFs2.0 USERFLASH      rw  slot0:
            -           -            -   network      rw  ftp:
            -           -            -   network      rw  tftp:
            -           -            -   network      rw  scp:
```

You can change the default file system so that file management commands apply to a particular device or memory.

To change the default storage location:

Step	Task	Command Syntax	Command Mode
1	Change the default directory.	**cd** *directory*	EXEC Privilege

In Figure 18, the default storage location is changed to the external Flash of the primary RPM. File management commands then apply to the external Flash rather than the internal Flash.

Figure 18 show file-system Command Example

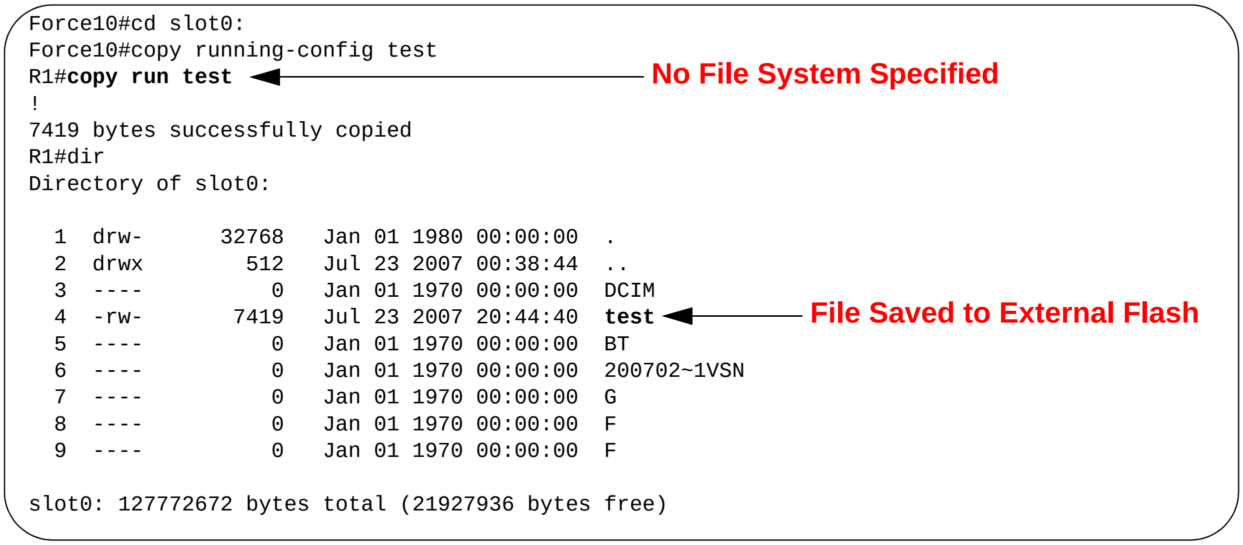

```
Force10#cd slot0:
Force10#copy running-config test
R1#copy run test
!
7419 bytes successfully copied
R1#dir
Directory of slot0:

  1  drw-      32768   Jan 01 1980 00:00:00  .
  2  drwx        512   Jul 23 2007 00:38:44  ..
  3  ----          0   Jan 01 1970 00:00:00  DCIM
  4  -rw-       7419   Jul 23 2007 20:44:40  test
  5  ----          0   Jan 01 1970 00:00:00  BT
  6  ----          0   Jan 01 1970 00:00:00  200702~1VSN
  7  ----          0   Jan 01 1970 00:00:00  G
  8  ----          0   Jan 01 1970 00:00:00  F
  9  ----          0   Jan 01 1970 00:00:00  F

slot0: 127772672 bytes total (21927936 bytes free)
```

No File System Specified

File Saved to External Flash

Loading a Different FTOS Version

To load a different FTOS version, see the upgrade procedure in the latest release notes.

Chapter 3 CAM Profiling

C-Series	NO
E-Series	✓

Platform Specific Feature: CAM Profiling is supported on the E-Series only.

This chapter includes the following major sections:

- CAM Profiling
- CAM IPv4flow Commands on page 83
- Configuring the CAM for IPv4Flow

CAM Profiling

The CAM Profiling feature enables you to partition the CAM entries to optimize your application. For example:

- When the E-Series is deployed as a switch, more Layer 2 FIB entries can be configured.
- When the E-Series is deployed as a router, more Layer 3 FIB entries can be configured.
- When IPv6 is not employed, additional CAM space can be allocated for ACL.

Warning: The CAM Profiling features are for users not wanting to use the default CAM profile. If you are using these features for the first time, contact Force10 TAC (Technical Assistance Center) for assistance.

Important Points to Remember

- The CAM Profiling feature has two levels: The top level—basic partitioning of the CAM—is a TeraScale-only feature. The basic partitioning of CAMs in EtherScale systems is fixed. However, both EtherScale and TeraScale systems can sub-partition the IPv4Flow CAM partition. EtherScale systems can also sub-partition the IPv4ACL partition. See the section CAM IPv4flow Commands on page 83 for details on sub-partitioning.
- CAM Profiling commands are similar to Boot commands; you must save the new profile (use the **copy running-config startup-config duplicate** command) and then reboot your chassis to apply the new profile.
- When budgeting your CAM allocations for ACLs and QoS configurations, remember to take into account that TCP and UDP rules with port range options may require more than one CAM entr.y

- Your chassis, RPMs, and all online line cards must have the same CAM profile.
- CAM Profiling includes Configuration and EXEC mode commands. Contact Force10 TAC before you use the EXEC-level commands.

The commands for the first level of CAM Profiling are (for syntax details, see the CAM Profiling chapter in the *FTOS Command Line Interface Reference*):

- **cam-profile default microcode** (see configuring the chassis with the default CAM profile and microcode on page 74)
- **cam-profile default microcode chassis** (see Debugging CAM Profiling on page 82)
- **cam-profile default microcode linecard** (see changing a card to the default profile on page 75)
- **cam-profile ipv4-egacl-16k microcode** (see CAM profile for the VLAN ACL group feature on page 77)
- **cam-profile ipv4-egacl-16k microcode chassis** (see Debugging CAM Profiling on page 82)
- **cam-profile ipv4-egacl-16k microcode linecard** (see Debugging CAM Profiling on page 82)
- **show cam-profile** (see Viewing CAM Profiles on page 78)
- **show cam-profile default microcode** (see Viewing CAM Profiles on page 78)
- **show cam-profile ipv4-egacl-16k microcode** (see Viewing CAM Profiles on page 78)
- **cam-profile ipv4-320k microcode lag-hash-mpls** (see LAG hashing on page 76)
- **cam-profile unified-default microcode ipv6-extacl** (see Unified-default CAM Profile on page 73)

On TeraScale systems, you can configure the CAM (content-addressable memory) so that its partitions better suit the needs of your network. Table 5 shows the default partitions (EtherScale partitions differ).

Table 5 FTOS Default CAM partitions

CAM Partition	Default
Layer 2 FIB (Forwarding Information Base)	32K entries
Layer 2 ACL (Access Control Lists)	1K entries
Layer 3 FIB	256K entries
Layer 3 ACL	12K entries
Layer 3 Flow	24K entries
EgL2ACL (egress Layer 2 ACL)	1K entries
EgIPv4ACL	1K entries
Reserved	8K entries
IPv6FIB	0 entries
IPv6ACL	0 entries
IPv6Flow	0 entries
EgIPv6ACL	0 entries

In addition to the partition for Layer 3 Flow shown above, FTOS supports an internal partitioning of that partition on TeraScale systems and of the Layer 3 ACL partition on EtherScale systems, as detailed in CAM IPv4flow Commands on page 83.

Unified-default CAM Profile

The unified CAM profile uses the IPv6-ExtACL microcode. It maintains the CAM allocations for IPv6 and IPv4 FIB, while allocating more CAM space for L2ACL, EgL2ACL, and IPv4ACL regions.

Matching Microcode to the CAM Profile

On TeraScale systems, the CAM Profiling commands enable you to adjust the CAM profile and to select the version of microcode that further optimizes the selected partitioning. On EtherScale systems, the CAM Profiling commands only adjust the CAM partitions, not select the microcode.

Note: A reference in this chapter to a CAM profile includes the associated microcode.

When to Use CAM Profiling

The CAM Profiling feature is available on TeraScale systems running at least FTOS 6.3.1.1. Using the feature is almost always optional. The feature is provided as a way to improve network performance in certain situations, as discussed below. Those situations include:

- When you want LAG hashing for MPLS packets to be done based on IP source and destination: See LAG hashing on page 76
- When you want to do LAG hashing based on bidirectional flow: See LAG hashing based on bidirectional flow on page 77
- When you want to optimize for the VLAN ACL Group feature, which permits group VLANs for the IP egress ACL: See CAM profile for the VLAN ACL group feature on page 77

A CAM profile is stored on every card, including each RPM. Because the same profile must be on every card in the chassis, the following three situations are when using the CAM Profiling feature is required. All situations occur after you have used the feature to select a non-default profile, and then you:

- Move a card from the chassis to a chassis that uses a default CAM profile. This, of course, would include any chassis running a version of FTOS lower than version 6.3.1.1.
- Replace the version of FTOS on the chassis with a version lower than version 6.3.1.1.
- Add a secondary RPM to a chassis that is using the non-default CAM profile.

In cases where you move or install a line card in a chassis that is already using a non-default profile, you can simply install the line card without considering its current profile. The system will automatically reconfigure the card’s profile to match the CAM profile in the chassis. The system will save the profile on the card, and will then reboot the card.

CAM Profiling Preparation

As suggested above, before you use the CAM Profiling commands, you should first consider your plans for changing chassis cards and system software.

Use CAM Profiling in this sequence:

1. If you plan to move a line card between two chassis that are both using the default CAM profile, but you want to modify one chassis to a non-default profile, move the line card before changing from the default.
2. If you plan to downgrade to an FTOS version lower than version 6.3.1.1, ensure that the chassis is using the default profile. See configuring the chassis with the default CAM profile and microcode.
3. If you plan to move a card from a chassis that uses a non-default CAM profile to a chassis using a default CAM profile, change the card to the default profile before moving it. See changing a card to the default profile, below.
4. If you add a new secondary RPM to a chassis that is using a non-default profile, add the profile to the RPM after installing it. See add the CAM profile to a new RPM, below.
5. After installing or moving cards, use the CAM Profiling commands that tune your system to your network needs. See Configuration Task List for Non-default CAM Profiling on page 76.

configuring the chassis with the default CAM profile and microcode

Before downgrading to a version of FTOS below version 6.3.1.1, make sure that the chassis is using the default CAM profile and microcode.

Step	Command Syntax	Command Mode	Purpose
1	**show cam-profile [summary]**	EXEC	Verify that the CAM profile and/or microcode is non-default, needing to be reset to default. The **summary** keyword yields a shorter, sufficient result.
2	**cam-profile default microcode [default | lag-hash-mpls | lag-hash-align]**	CONFIGURATION	Enter **cam-profile default microcode default** to revert the chassis to the CAM and microcode defaults.
3	**copy running-config** *<file-URL>*	EXEC privilege	Save the changed configuration to NVRAM. Any of the keywords and variables of the command that you use will accomplish the task. For example, you can enter **copy running-config startup-config**.
4	**show running-config cam-profile**	EXEC	Verify that the **cam-profile** command was executed.
5	**reload**	EXEC privilege	Even though the running configuration shows that the default profile and microcode are selected, and you have copied the profile to memory, you must reboot the chassis for the change to be propagated to the line card memory.

changing a card to the default profile

Before you move a line card from a chassis with a non-default CAM profile to a chassis with a default CAM profile, use the CAM Profiling command **cam-profile default microcode linecard** in the EXEC mode to load the default profile.

Step	Command Syntax	Command Mode	Purpose
1	**show cam-profile [summary]**	EXEC	Verify that the CAM profile is non-default, needing to be reset on the line card.
2	**cam-profile default microcode [default | lag-hash-mpls | lag-hash-align] linecard** *number*	EXEC	Enter **cam-profile default microcode default linecard** *number*, where *number* is the slot number, to revert the line card to the CAM and microcode defaults.

Move the line card to the new chassis. Since the CAM profiles and microcode levels match, you do not need to execute any more commands.

add the CAM profile to a new RPM

If you add a secondary RPM to a chassis, add the existing CAM profile to the RPM after installing it.

Step	Command Syntax	Command Mode	Purpose
1			Install the secondary RPM.
2	**copy running-config** *<file-URL>*	EXEC privilege	Save the configuration on the primary. This writes the CAM profile to the NVRAM of the secondary RPM as well. There is no need to reboot the chassis.

Configuration Task List for Non-default CAM Profiling

This section describes the situations and associated commands in which you would use a non-default CAM profile, including associated microcode. The three non-default CAM Profiling options in Version 6.3.1 are:

- LAG hashing (optional)
- LAG hashing based on bidirectional flow (optional)
- CAM profile for the VLAN ACL group feature on page 77 (optional)

All cards in a chassis must use the same CAM profile. Therefore, you must install cards and run CAM Profiling commands in the proper sequence in order for that to occur. Those procedures are described above. After you have completed those procedures, you can use the CAM Profiling command sequences that are described below.

LAG hashing

The default system software and configuration in TeraScale systems include a CAM profile and microcode that treats MPLS packets as non-IP packets. Switching and LAG hashing is based on SA and DA MAC addresses. Alternatively, you can have LAG hashing for MPLS packets based on IP SA and DA and IP protocol. This is applicable for MPLS packets with 5 labels or less than 5 labels. To accomplish this, use the default CAM profile alone with the **cam-profile default microcode lag-hash-mpls** command (see Step 1 below) to load the LAG-Hash-MPLS microcode.

FTOS can look up to 5 labels deep into MPLS packets for an IP header. MPLS packets are treated as follows for the default load-balance:

- When MPLS IP packets are received, FTOS looks up to 5 labels deep for the IP header.
- When an IP header is present, hashing is based on IP 3 tuple (IP SA, IP DA and IP protocol).
- If an IP header is not found after the 5th label, hashing is based on the MPLS labels.
- If the packet has more than 5 MPLS labels, hashing is based on MAC SA and DA.

Note: MPLS packets are always hashed based on MAC SA, DA for load-balanced MAC configuration.

Step	Command Syntax	Command Mode	Purpose
1	**cam-profile [default | ipv4-320k] microcode lag-hash-mpls**	CONFIGURATION	Entering the keyword **lag-hash-mpls** option designates the lag-hash-mpls microcode. View the ipv4-320k CAM profile using the **show cam-profile** command.

2	**copy running-config** <*file-URL*>	EXEC privilege	Save the changed configuration to NVRAM. Any of the keywords and variables of the command that you use will accomplish the task. For example, you can enter **copy running-config startup-config**.
3	**show running-config cam-profile**	EXEC	Verify that the **cam-profile** command was executed.
4	**reload**	EXEC privilege	Reboot the chassis for the change to be propagated to all cards.

LAG hashing based on bidirectional flow

If you want to do LAG hashing such that both directions of a bidirectional flow (for example, VoIP or P2P file sharing) are mapped to the same output link in the LAG bundle, then you would load the LAG-hash-align microcode using the **cam-profile default** command in the CONFIGURATION mode.

Step	Command Syntax	Command Mode	Purpose
1	**cam-profile default microcode lag-hash-align**	CONFIGURATION	Enter the keyword **lag-hash-align** to designate the lag-hash-align microcode.
2	**copy running-config** <*file-URL*>	EXEC privilege	Save the changed configuration to NVRAM. Any of the keywords and variables of the command that you use will accomplish the task. For example, you can enter **copy running-config startup-config**.
3	**show running-config cam-profile**	EXEC	Verify that the **cam-profile** command was executed.
4	**reload**	EXEC privilege	Reboot the chassis for the change to be propagated to all cards.

CAM profile for the VLAN ACL group feature

Note: Do not use this CAM profile for Layer 2 egress ACLs.

If you want to optimize for the VLAN ACL Group feature, which permits group VLANs for the IP egress ACL, then you would use the **cam-profile ipv4-egacl-16k** command in the CONFIGURATION mode.

Step	Command Syntax	Command Mode	Purpose
1	**cam-profile ipv4-egacl-16k microcode acl-group**	CONFIGURATION	Enter the keyword **acl-group** to designate the acl-group microcode.
2	**copy running-config** *<file-URL>*	EXEC privilege	Save the changed configuration to NVRAM. Any of the keywords and variables of the command that you use will accomplish the task. For example, you can enter **copy running-config startup-config**.
3	**show running-config cam-profile**	EXEC	Verify that the **cam-profile** command was executed.
4	**reload**	EXEC privilege	Reboot the chassis for the change to take effect.

The IPv4-egacl-16 CAM profile contains the following CAM partitions (The partitions containing value changes from the default CAM profile are shown in bold type):

Profile Name	: IPv4-egACL-16k
L2FIB	: 32K entries
L2ACL	**: 7K entries**
IPv4FIB	**: 192K entries**
IPv4ACL	**: 8K entries**
IPv4Flow	: 24K entries
EgL2ACL	**: 0 entries**
EgIPv4ACL	**: 16K entries**
Reserved	: 8K entries
IPv6FIB	: 0 entries
IPv6ACL	: 0 entries
IPv6Flow	: 0 entries
EgIPv6ACL	: 0 entries

Viewing CAM Profiles

The CAM Profiling feature provides the three show commands described in this section. In addition, you can use the **show running-config** command, as shown below, to display the CAM profile.

If you use the **show running-config** command after you have used a CAM Profiling command and before you have saved the new profile or rebooted, be aware that the CAM profile displayed by **show running-config** is not necessarily the active CAM profile. If you set the CAM profile from the CONFIGURATION mode, then this command shows the selected profile, not the active one. Otherwise, it shows the pre-existing CAM profile. Remember that you must save the profile to NVRAM before you reboot, or the CAM profile will not take effect.

Command Syntax	Command Mode	Usage
show running-config cam-profile	EXEC	Appended with the **cam-profile** delimiter, this command displays just the CAM Profiling section of the running configuration report.

```
Force10#show running-config cam-profile
!
cam-profile default microcode default

Force10#
```

Figure 19 show running-config Command Example

The **show cam-profile** command produces very different displays depending on how you modify it:

- Entering only **show cam-profile** (without any parameters) results in the display shown in Figure 22 on page 81. The command **show cam-profile summary** generates a condensed form of the same report.
- Use the **show cam-profile** command with <*cam-profile*> and **microcode** <*microcode*> parameters as a help command, in other words, to preview the CAM partitions that are allocated by the specified CAM profile. Using the command **show cam-profile** <*cam-profile*> without the microcode parameters will return an error. To display the desired partition list, enter the name of the CAM profile and the name of the microcode. For example, use **show cam-profile ipv4-egacl-16k microcode acl-group** to display the partitions that would be allocated to the **ipv4-egacl-16k** profile. You can enter **show cam-profile ?** to see the profiles that are available in this version.

Command Syntax	Command Mode	Usage
show cam-profile <*cam-profile*> **microcode** <*microcode*>	EXEC	Use the command as a help command, in other words, to preview the CAM partitions that are allocated by the **cam-profile** commands available in this version.

```
Force10#show cam-profile default microcode default

---  Sample Configuration of CAM Profile  ---

CamSize          : 18-Meg
                 : Settings
Profile Name     : DEFAULT
L2FIB            : 32K entries
L2ACL            : 1K entries
IPv4FIB          : 256K entries
IPv4ACL          : 12K entries
IPv4Flow         : 24K entries
EgL2ACL          : 1K entries
EgIPv4ACL        : 1K entries
Reserved         : 8K entries
IPv6FIB          : 0  entries
IPv6ACL          : 0  entries
IPv6Flow         : 0  entries
EgIPv6ACL        : 0  entries
MicroCode Name   : Default
Force10#
```

Figure 20 show cam-profile default Command Example (partitions allocated to default profile)

Using the **show cam-profile** command without parameters always results in a three-column display, as shown below:

- Left column: A list of the parameter headings (Profile Name, partition names, and Microcode Name)
- Middle column: The currently configured CAM profile (and microcode name) for the chassis and all the online line cards
- Right column: The CAM profile and the microcode name *most recently saved*, that will be configured for the chassis and all online line cards *after the next boot*.

Command Syntax	Command Mode	Usage
show cam-profile [summary]	EXEC	(OPTIONAL) Enter the keyword **summary** to view a summary listing (profile name and microcode name) by card.

The following example shows the condense report produced by the **show cam-profile summary** command. For an example of the detailed report, see Figure 22 on page 81.

```
Force10#show cam-profile summary

-- Chassis CAM Profile --
                 : Current Settings : Next Boot
Profile Name     : Default          : Default
MicroCode Name   : Default          : Default

                 : Current Settings : Next Boot
-- Line card 1 --
Profile Name     : Default          : Default
MicroCode Name   : Default          : Default

                 : Current Settings : Next Boot
-- Line card 6 --
Profile Name     : Default          : Default
MicroCode Name   : Default          : Default
Force10#
```

Figure 21 show cam-profile summary Command Example

The following example of using the **show cam-profile** command to display details is a truncated display of the full report.

```
Force10#show cam-profile

-- Chassis Cam Profile --

CamSize          : 18-Meg
                 : Current Settings : Next Boot
Profile Name     : DEFAULT          : DEFAULT
L2FIB            : 32K entries      : 32K entries
L2ACL            : 1K entries       : 1K entries
IPv4FIB          : 256K entries     : 256K entries
IPv4ACL          : 12K entries      : 12K entries
IPv4Flow         : 24K entries      : 24K entries
EgL2ACL          : 1K entries       : 1K entries
EgIPv4ACL        : 1K entries       : 1K entries
Reserved         : 8K entries       : 8K entries
IPv6FIB          : 0  entries       : 0  entries
IPv6ACL          : 0  entries       : 0  entries
IPv6Flow         : 0  entries       : 0  entries
EgIPv6ACL        : 0  entries       : 0  entries
MicroCode Name   : Default          : Default

-- Line card 0 --
CamSize          : 18-Meg
                 : Current Settings : Next Boot
Profile Name     : DEFAULT          : DEFAULT
L2FIB            : 32K entries      : 32K entries
L2ACL            : 1K entries       : 1K entries
IPv4FIB          : 256K entries     : 256K entries
IPv4ACL          : 12K entries      : 12K entries
IPv4Flow         : 24K entries      : 24K entries
EgL2ACL          : 1K entries       : 1K entries
EgIPv4ACL        : 1K entries       : 1K entries
Reserved         : 8K entries       : 8K entries
IPv6FIB          : 0  entries       : 0  entries
IPv6ACL          : 0  entries       : 0  entries
IPv6Flow         : 0  entries       : 0  entries
EgIPv6ACL        : 0  entries       : 0  entries
MicroCode Name   : Default          : Default

-- Line card 13 --
CamSize          : 18-Meg
                : Current Settings : Next Boot
Profile Name     : Default          : Default
```

Figure 22 show cam-profile (details) Command Example

Debugging CAM Profiling

Issues that might be anticipated for the CAM Profiling feature primarily involve CAM profile mismatches on chassis components, because, while the whole system must use the same CAM profile, it is possible to configure different CAM profiles and microcodes on the various components.

The section CAM Profiling Preparation on page 74 discusses how to avoid those mismatches. However, the **cam-profile default microcode linecard** command discussed in that section for avoiding a mismatch, is one of a group of four similar commands that might be used by mistake to create a mismatch. Those commands are:

- **cam-profile default microcode chassis**
- **cam-profile default microcode linecard**
- **cam-profile ipv4-egacl-16k microcode acl-group chassis**
- **cam-profile ipv4-egacl-16k microcode acl-group linecard**

The commands are generally intended for reconciling mismatches in a debug mode, because, combined with the previously described procedures and configuration commands, TeraScale systems have a CAM profile conflict resolution algorithm that resolves mismatches. The software will detect a mismatch between line cards and RPM. RPM NVRAM settings always take precedence over the line card EEPROM (and the startup-config).

When a line card is inserted with a different CAM profile, the system will detect the mismatch, and the line card will be automatically set to the CAM profile and microcode of the RPM. Note that the line card reset requires a longer time for the line card to check in than usual. If the system cannot programmatically match the CAM profile on the line card to the chassis profile after three attempts to power cycle the card, the card will be put in a problem state.

When there is no startup-config after reload, the running-config will inherit (and display) the CAM profile that is currently stored in NVRAM.

When the running-config with a new CAM profile is saved in a chassis with dual RPMs, the CAM profile gets written to the NVRAM of the primary RPM, secondary RPM, and EEPROM of all line cards. If the primary RPM fails, the secondary RPM takes over, and CAM settings are unaffected.

To help you avoid problems, or at least isolate them, remember that:

- The CAM Profiling feature is available only for TeraScale systems running at least FTOS version 6.3.1.1.
- For normal CAM Profiling configuration, use the CONFIGURATION mode. Use the commands above in the EXEC mode primarily for troubleshooting.
- You must copy the new CAM profile to NVRAM and then reboot the chassis for a new CAM profile to take effect.
- When you insert a secondary RPM after changing and saving the CAM profile settings to the primary RPM, you must save the config again for the CAM profile to propagate to the secondary RPM.

CAM IPv4flow Commands

This section contains the following sections:

- Important Points to Remember on page 83
- Configuring the CAM for IPv4Flow on page 84
- configuring ip4flow for a line card on page 87
- Configuring CAM IPv4Flow for the System on page 84
- Viewing the CAM IPv4Flow on page 88

This section discusses IPv4Flow partitioning, which consists of sub-partitioning of the Layer 3 ACL and Layer 3 Flow partitions of the CAM on EtherScale systems and the Layer 3 Flow partition of the CAM on TerarScale systems.

Release 6.3.1 introduces the CONFIGURATION mode command—**cam ipv4flow**—one version for EtherScale IPv4Flow partitioning and one for TeraScale. As with the CAM Profiling of the parent partitions (see CAM Profiling on page 71), use the CONFIGURATION mode commands to configure globally. Use the EXEC mode commands primarily for adjusting the profile of the selected component to match the overall profile of the target system.

The commands for the IPv4Flow partitioning are as follows (for syntax details, see the CAM Profiling chapter in the FTOS Command Line Interface Reference):

- cam ipv4flow (see Configuring CAM IPv4Flow for the System on page 84)
- show cam-ipv4flow (see Viewing the CAM IPv4Flow on page 88)

  Maintenance commands (EXEC mode):

- cam-ipv4flow chassis all (see configuring the default ipv4flow for the chassis on page 86)
- cam ipv4flow linecard (see configuring ip4flow for a line card on page 87)

Important Points to Remember

- The chassis and all line cards within a single system must have the same CAM IPv4Flow profile.
- The CONFIGURATION mode IPv4Flow commands are similar to boot commands; you must save you new settings (**copy running-config**) and reboot your chassis for the new settings to take effect.
- Any change in the CAM IPv4Flow requires a reboot for the new profile to take effect.
- When rebooting the system, the software checks that all line cards in the chassis are configured with the same IPv4flow configuration. Starting with FTOS Release 6.3.1, line cards with IPv4flow configuration that do not match the system configuration are automatically reconfigured and rebooted. If three resets do not bring up the card, or if the cards are not brought up because the release is previous to 6.3.1—the system presents the error message “card problem – mismatch cam ipv4flow”—you must manually adjust the profile. To adjust the profile of the line card, see changing a card to the default profile on page 75 and configuring ip4flow for a line card on page 87.

Caution: Please contact Force10 Technical Assistance Center (TAC) before you use these commands.

Configuring the CAM for IPv4Flow

An FTOS ACL task controls and ensures that the CAM resources are properly allocated by changing the CAM entries according to the memory resources allowed in the system. With the CLI, the task of the ACL uses the user-defined IPv4Flow setting as its maximum size allowed.

IPv4Flow has the following space allocation restrictions:

- For IPv4flow, allocation is from a minimum of 1K to a maximum of 24K (TeraScale cards) or 32K (EtherScale cards).
- For System Flow, allocation is from a minimum of 4K to a maximum of 32K.

If you do not select the default IPv4Flow configuration, you must specify a size for each of the following IPv4Flow entries, and you must do so in this order:

- ACL (for EtherScale only)
- Multicast FIB/ACL
- PBR
- QoS
- System Flow
- Trace Lists

The IPv4Flow configuration is applied to the chassis and all line cards.

Configuring CAM IPv4Flow for the System

FTOS Release 6.3.1 introduces two CONFIGURATION mode commands (one for EtherScale, the other for TeraScale), in for modifying IPv4Flow settings. You can save the changes made by the new commands to the startup configuration file, which you can then copy to other systems to replicate IPv4Flow configuration changes.

The four pre-existing EXEC mode commands are still available, but you should only use those commands to correct mismatches. The configuration changes made by those commands are saved to the affected cards immediately (no need to save the running-config).

Use the CONFIGURATION mode command **cam-ipv4flow** to configure the chassis and all line cards with the same IPv4Flow profile.

Note: Note that command has a hyphen between **cam** and **ipv4flow**, in contrast to the syntax of the EXEC mode commands.

Step	Command Syntax	Command Mode	Purpose
1	**cam-ipv4flow** [{**acl** *value*} {**multicast-fib** *value*} {**pbr** *value*} {**qos** *value*} {**system-flow** *value*} {**trace-list** *value*}]	CONFIGURATION	To configure the chassis and all line cards to the same IPv4Flow profile. Note: The {**acl** *value*} parameter appears in EtherScale systems only.
2	**copy running-config startup-config**	EXEC privilege	Use the **copy running-config** command to save the configuration and write the profile to all cards. The change will take effect after the next reboot.
3	**reload**	EXEC privilege	Reboot the system.

The following example shows the process of using the **cam-ipv4flow** command (TeraScale version) with the default option, followed by **copy running-config** and **show cam-ipv4flow**:

```
Force10(conf)#cam-ipv4flow default
Force10#copy running-config startup-config
File with same name already exist.
Proceed to copy the file [confirm yes/no]: yes
!
3914 bytes successfully copied

Force10#sh cam-ipv4flow
-- Chassis Cam Ipv4Flow --
                     Current Settings   Next Boot
Multicast Fib/Acl :   8K                 9K
Pbr               :   2K                 1K
Qos               :   7K                 8K
System Flow       :   6K                 5K
Trace Lists       :   1K                 1K

-- Line card 0 --

                     Current Settings   Next Boot
Multicast Fib/Acl :   8K                 9K
Pbr               :   2K                 1K
Qos               :   7K                 8K
System Flow       :   6K                 5K
Trace Lists       :   1K                 1K

-- Line card 1 --
                     Current Settings   Next Boot

Multicast Fib/Acl :   8K                 9K
Pbr               :   2K                 1K
Qos               :   7K                 8K
System Flow       :   6K                 5K
Trace Lists       :   1K                 1K
```

Figure 23 show ipv4flow cam Command Example

Performing Maintenance Functions

This section discusses the use of the EXEC mode commands for setting IPv4Flow entries. The cases that involve the use of these commands are typically where you are resolving mismatches, such as moving a line card from a system running FTOS Release 6.3.1 or later to a system running an earlier version of FTOS. These are the same conditions as described for the high-level CAM partitions in When to Use CAM Profiling on page 73.

configuring the default ipv4flow for the chassis

Below is an example of using the EXEC mode command **cam ipv4flow chassis all default** to configure the IPv4Flow CAM profile of a TeraScale chassis to the default profile. You would typically do this before installing a previous version of FTOS that does not support the current profile.

Note the default sizes are shown in the example for each entry after the command is executed:

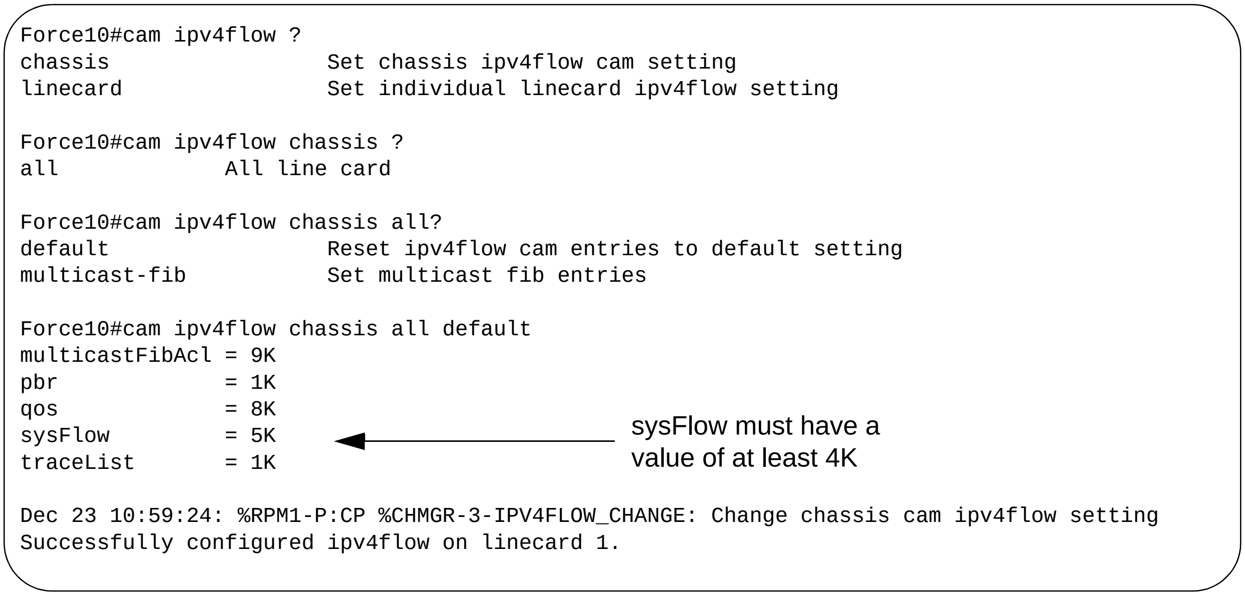

```
Force10#cam ipv4flow ?
chassis                 Set chassis ipv4flow cam setting
linecard                Set individual linecard ipv4flow setting

Force10#cam ipv4flow chassis ?
all             All line card

Force10#cam ipv4flow chassis all?
default                 Reset ipv4flow cam entries to default setting
multicast-fib           Set multicast fib entries

Force10#cam ipv4flow chassis all default
multicastFibAcl = 9K
pbr             = 1K
qos             = 8K
sysFlow         = 5K
traceList       = 1K

Dec 23 10:59:24: %RPM1-P:CP %CHMGR-3-IPV4FLOW_CHANGE: Change chassis cam ipv4flow setting
Successfully configured ipv4flow on linecard 1.
```

sysFlow must have a value of at least 4K

Figure 24 cam ipv4flow chassis Command Example

configuring ip4flow for a line card

Use the **cam ipv4flow linecard** command in EXEC mode to adjust the IPv4Flow profile of a single selected line card (designated by **linecard** *slot number*).

Starting with FTOS Release 6.3.1, executing the command does not incur a reset prompt, because you are generally using this command to prepare a card to be moved to another chassis. However, if you are moving a card into a chassis with a non-default profile, you do not need to prepare its profile. The target system should recognize that the card does not have the same profile, and then should automatically adjust the profile on the card to match the CAM profile on the target chassis.

```
Force10#cam ipv4flow linecard 0 ?
default                 Reset IPv4Flow CAM entries to default setting
multicast-fib           Set multicast fib entries
Force10#$ecard 0 multicast-fib 7 pbr 3 qos 7 system-flow 6 trace-list 1

New linecard cam ipv4flow setting:
multicastFibAcl = 7K
pbr             = 3K
qos             = 7K
sysFlow         = 6K
traceList       = 1K

00:52:27: %RPM1-P:CP %CHMGR-3-IPV4FLOW_CHANGE: Changed linecard 0 CAM IPv4Flow setting
Force10#sh cam-ipv4flow

-- Chassis Cam Ipv4Flow --
                     Current Settings   Next Boot
Multicast Fib/Acl :   9K                 8K
Pbr               :   1K                 2K
Qos               :   8K                 7K
System Flow       :   5K                 6K
Trace Lists       :   1K                 1K

-- Line card 0 --
                     Current Settings   Next Boot
Multicast Fib/Acl :   9K                 7K
Pbr               :   1K                 3K
Qos               :   8K                 7K
System Flow       :   5K                 6K
Trace Lists       :   1K                 1K

-- Line card 1 --
                     Current Settings   Next Boot
Multicast Fib/Acl :   9K                 8K
Pbr               :   1K                 2K
Qos               :   8K                 7K
System Flow       :   5K                 6K
Trace Lists       :   1K                 1K
```

Figure 25 cam ipv4flow linecard Command Example

Viewing the CAM IPv4Flow

To view IPv4Flow CAM entries on line cards (on TeraScale system):

```
Force10#show cam-ipv4flow ?
all                      status of all line cards
|                        Pipe through a command
<cr>

Force10#show cam-ipv4flow
-- Chassis Ipv4Flow Entries --

Multicast Fib/Acl :   9K
Pbr               :   1K
Qos               :   8K                  Maximum total of 24K
System Flow       :   5K
Trace Lists       :   1K

-- Line card 1 --
Multicast Fib/Acl :   9K
Pbr               :   1K
Qos               :   8K
System Flow       :   5K
Trace Lists       :   1K
```

Figure 26 show ipv4flow cam Command Example

Configuring the 32K Ingress ACLs Profile

The CAM profile *l2-ipv4-inacl* also called *32K Ingress ACLs* profile has the CAM entries shown in Figure 27, and uses the default microcode.

Figure 27 32K Ingress ACL CAM Profile

```
Force10#show cam-profile

-- Chassis CAM Profile --

CamSize          : 18-Meg
                 : Current Settings : Next Boot
Profile Name     : l2-ipv4-inacl    : l2-ipv4-inacl
L2FIB            : 32K entries      : 32K entries
L2ACL            : 32K entries      : 32K entries
IPv4FIB          : 64K entries      : 64K entries
IPv4ACL          : 27K entries      : 27K entries
IPv4Flow         : 8K entries       : 8K entries
EgL2ACL          : 2K entries       : 2K entries
EgIPv4ACL        : 2K entries       : 2K entries
Reserved         : 2K entries       : 2K entries
IPv6FIB          : 0  entries       : 0  entries
IPv6ACL          : 0  entries       : 0  entries
IPv6Flow         : 0  entries       : 0  entries
EgIPv6ACL        : 0  entries       : 0  entries
MicroCode Name   : Default          : Default
```

A CAM profile can be selected from either EXEC Privilege mode or CONFIGURATION mode.

- The CAM profile can be applied to the entire system or particular line cards if selected from EXEC Privilege mode. In this case, the system is immediately configured to write the new profile upon the next system boot.
- The CAM profile is applied to entire system if it is selected from CONFIGURATION mode, however the running-config must be saved for the change to take affect.

To change the CAM profile to *l2-ipv4-inacl* on the entire system:

Step	Task	Command Syntax	Command Mode
1	Change the CAM profile to *l2-ipv4-inacl.*	**cam-profile l2-ipv4-inacl microcode default**	EXEC Privilege
2	Verify that the new CAM profile will be written to the CAM on the next boot (Figure 28).	**show cam-profile summary**	EXEC Privilege
3	Reload the system.	**reload**	EXEC Privilege

To change the CAM profile to *l2-ipv4-inacl* on the entire system or a particular line card:

Step	Task	Command Syntax	Command Mode
1	Change the CAM profile to *l2-ipv4-inacl.*	**cam-profile l2-ipv4-inacl microcode default**	CONFIGURATION
2	Save the running-configuration	**copy running-config startup-config**	EXEC Privilege
3	Verify that the new CAM profile will be written to the CAM on the next boot (Figure 28).	**show cam-profile summary**	EXEC Privilege
4	Reload the system.	**reload**	EXEC Privilege

Figure 28 32K Ingress ACL CAM Profile Summary

```
Force10#show cam-profile summary

-- Chassis CAM Profile --
                 : Current Settings : Next Boot
Profile Name     : l2-ipv4-inacl    : l2-ipv4-inacl
MicroCode Name   : Default          : Default
--More--
```

Chapter 4 Management

This chapter explains the different protocols or services used to manage the Force10 system including:

- System Log Management on page 91
- SNMP on page 98
- Network Time Protocol on page 108
- File Transfer Services on page 112
- Terminal Lines on page 114

With FTOS you can choose among several different options for monitoring and troubleshooting the software and the system. By enabling debug commands, you can perform some troubleshooting. To get help with troubleshooting, you can view logs and different show command outputs.

System Log Management

Use the logging commands track changes in the system. With FTOS you can configure, save, and view system messages and error messages.

All error messages, except those beginning with %BOOTUP, are stored in the logging buffer. Below is an example of a message not stored in the logging buffer:

```
%BOOTUP:RPM0:CP %PORTPIPE-INIT-SUCCESS: Portpipe 0 enabled
```

Configuration Task List for System Log Management

By default, logging is enabled to the internal buffer, console and terminal lines, and any configured Syslog servers.

The following list includes the configuration tasks for system log management:

- enable logging on page 92 (default behavior)
- specify logging to a Syslog server on page 93 (optional)
- change logging settings on page 93 (optional)
- configure a Syslog server on page 94 (optional)
- FTOS support for software errors—core dumps on page 95
- configure a UNIX logging facility level on page 96 (optional)
- synchronize log messages on page 97 (optional)
- enable timestamp on Syslog messages on page 97 (optional)

For a complete listing of logging commands, refer to .

enable logging

By default, logging is enabled and log messages are sent to the internal buffer, all terminal lines, console, and Syslog servers. However, you must configure the IP address (with the **logging** command) of a Syslog server for a Syslog server to receive the log messages.

To disable logging except to the console, enter **no logging on** in the CONFIGURATION mode. To disable logging to the console, enter **no logging console** in the CONFIGURATION mode.

To re-enable full logging, enter **logging on** in the CONFIGURATION mode.

To view the current contents of the logging buffer and the logging settings for the system, use the **show logging** command (Figure 29) in the EXEC privilege mode.

```
Force10#show logging
Syslog logging: enabled
    Console logging: level Debugging
    Monitor logging: level Debugging
    Buffer logging: level Debugging, 40 Messages Logged, Size (40960 bytes)
    Trap logging: level Informational
%IRC-6-IRC_COMMUP: Link to peer RPM is up
%RAM-6-RAM_TASK: RPM1 is transitioning to Primary RPM.
%RPM-2-MSG:CP1 %POLLMGR-2-MMC_STATE: External flash disk missing in 'slot0:'
%CHMGR-5-CARDDETECTED: Line card 0 present
%CHMGR-5-CARDDETECTED: Line card 2 present
%CHMGR-5-CARDDETECTED: Line card 4 present
%CHMGR-5-CARDDETECTED: Line card 5 present
%CHMGR-5-CARDDETECTED: Line card 8 present
%CHMGR-5-CARDDETECTED: Line card 10 present
%CHMGR-5-CARDDETECTED: Line card 12 present
%TSM-6-SFM_DISCOVERY: Found SFM 0
%TSM-6-SFM_DISCOVERY: Found SFM 1
%TSM-6-SFM_DISCOVERY: Found SFM 2
%TSM-6-SFM_DISCOVERY: Found SFM 3
%TSM-6-SFM_DISCOVERY: Found SFM 4
%TSM-6-SFM_DISCOVERY: Found SFM 5
%TSM-6-SFM_DISCOVERY: Found SFM 6
%TSM-6-SFM_DISCOVERY: Found SFM 7
%TSM-6-SFM_SWITCHFAB_STATE: Switch Fabric: UP
%TSM-6-SFM_DISCOVERY: Found SFM 8
%TSM-6-SFM_DISCOVERY: Found 9 SFMs
%CHMGR-5-CHECKIN: Checkin from line card 5 (type EX1YB, 1 ports)
%TSM-6-PORT_CONFIG: Port link status for LC 5 => portpipe 0: OK portpipe 1: N/A
%CHMGR-5-LINECARDUP: Line card 5 is up
%CHMGR-5-CHECKIN: Checkin from line card 12 (type S12YC12, 12 ports)
%TSM-6-PORT_CONFIG: Port link status for LC 12 => portpipe 0: OK portpipe 1: N/A
%CHMGR-5-LINECARDUP: Line card 12 is up
%IFMGR-5-CSTATE_UP: changed interface Physical state to up: So 12/8
%IFMGR-5-CSTATE_DN: changed interface Physical state to down: So 12/8
```

Figure 29 show logging Command Example (Partial)

specify logging to a Syslog server

By default, all system messages are stored in the logging internal buffer and Syslog servers. You can add external devices and change the settings for storing messages in an internal buffer.

To specify different Syslog servers on the system, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
logging {*ip-address* | *hostname*}	CONFIGURATION	Configure a Syslog server to receive log messages from the system. Enter the IP address or host name of the server. You can configure up to eight Syslog servers to store system messages.

To view any changes made, use the **show running-config logging** command (Figure 30) in the EXEC privilege mode.

change logging settings

You can change the default settings of the system logging by changing the severity level and the storage location. The default is to log all messages up to debug level, that is, all system messages. By changing the severity level in the logging commands, you control the number of system messages logged.

To change one of the settings for logging system messages, use any or all of the following commands in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
logging buffered [*level*] [*size*]	CONFIGURATION	Specify the minimum severity level and number of the system messages logged to an internal buffer. Configure the following optional parameters: • *level* range: 0 to 7 or one of the following message levels (emergencies, alerts, critical, errors, warning, notifications, informational, or debugging) • *size* range: 40960 to 524288 bytes. The default setting is size 40960 and level 7. To return to the default setting, enter **default logging buffered**. When you decrease the buffer size, all messages stored in the buffer are lost. Increasing the buffer size does not affect messages stored in the buffer.

Command Syntax	Command Mode	Purpose
logging console [*level*]	CONFIGURATION	Specify the severity level for messages logged to the console. Configure the following optional parameter: • *level* range: 0 to 7 or one of the following message levels (emergencies, alerts, critical, errors, warning, notifications, informational, or debugging) The default setting is level 7. To return to the default setting, enter **default logging console**.
logging history *level*	CONFIGURATION	Specify the severity level for messages saved to the system history table and sent to the SNMP server: • *level* range: 0 to 7 or one of the following message levels (emergencies, alerts, critical, errors, warning, notifications, informational, or debugging). The default setting is level 4.
logging history size *size*	CONFIGURATION	Specify the number of messages saved to the system history table: • *size* range: 0 to 500 messages. The default setting is 1 message.
logging monitor [*level*]	CONFIGURATION	Specify the severity level for messages sent to terminal lines: • *level* range: 0 to 7 or one of the following message levels (emergencies, alerts, critical, errors, warning, notifications, informational, or debugging). The default setting is level 7. To return to the default setting, enter **default logging monitor**.

To view the logging buffer and configuration, use the **show logging** command (Figure 29) in the EXEC privilege mode.

configure a Syslog server

You can configure a BSD or SunOS UNIX system as a Syslog server. For system messages to be stored on a Syslog server, you must configure the `syslog.conf` file in the Syslog server and assign write permission to the file.

The following examples configure a Syslog daemon for messages up to the debugging level in two different operating systems:

- for a 4.1 BSD UNIX system, include this line in the `/etc/syslog.conf` file

```
local7.debugging /var/log/force10.log
```

- for a 5.7 SunOS UNIX system, include this line in the `/etc/syslog.conf` file

```
local7.debugging /var/adm/force10.log
```

In the lines above, local7 is the logging facility and debugging is the Syslog level. Therefore the Syslog daemon sends all messages since debugging is the lowest Syslog level. Refer to **logging facility** and **logging console** command descriptions for more information on those keywords.

To change the severity level of messages logged to a Syslog server, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
logging trap [*level*]	CONFIGURATION	Specify the severity level for messages sent to a Syslog server: • *level* range: 0 to 7 or one of the following message levels (emergencies, alerts, critical, errors, warning, notifications, informational, or debugging). The default setting is level 6. To return to the default setting, enter **default logging trap**.

To view the logging configuration, use the **show running-config logging** command (Figure 30) in the EXEC privilege mode.

FTOS support for software errors—core dumps

Two types of core dumps—application and Kernel—are enabled by default. In addition, the user may turn off the core dump for Kernel crashes by using the CLI. The High Availability module is aware of the core dump upload and it does not reboot the crashed RPM until the core dump has completed or is aborted.

The Flash should have enough memory to hold core dumps, however users are encouraged to configure an FTP server as the core dump destination.

Kernel core dump—By default the Kernel core dump would be sent to the Flash device in the CORE_DUMP_DIR directory, however if Flash is out of memory, the core-dump is aborted. Using the CLI, the user may configure a server as the FTP target location for the core dump. The kernel core dumps are overwritten every time there is a new core dump. The user should upload kernel core dump manually if an FTP server is not configured and should subsequently delete it from flash. The kernel core dump is named `f10rp1.kcore.gz`.

Note: The Kernel core dump can be large and may take up to 10 to 15 minutes to upload.

Application core dump—By default, the application core dump can only be sent to the Flash device, however if Flash is out of memory, the core dump is aborted. Application core dumps have a timestamp embedded in them that prevents them from being overwritten by default. It is up to the user to delete the core dump files. Application core dumps are named as `f10rp1<yymmddhhmmss>.acore.gz`

You can configure a system to enable Kernel core dumps:

Step	Command Syntax	Command Mode	Purpose
1	**logging kernel-coredump disable**	CONFIGURATION	To disable the kernel core dump function. The default setting is core dump enable.
2	**no logging kernel-coredump disable**	CONFIGURATION	To enable kernel core dump function.
3	**logging kernel-coredump server [server IP address/hostname] [login name] [password]**	CONFIGURATION	To specify the server.

configure a UNIX logging facility level

You can save system log messages with a UNIX system logging facility.

To configure a UNIX logging facility level, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
logging facility [*facility-type*]	CONFIGURATION	Specify one of the following parameters. • auth (for authorization messages) • cron (for system scheduler messages) • daemon (for system daemons) • kern (for kernel messages) • local0 (for local use) • local1 (for local use) • local2 (for local use) • local3 (for local use) • local4 (for local use) • local5 (for local use) • local6 (for local use) • local7 (for local use). This is the default. • lpr (for line printer system messages) • mail (for mail system messages) • news (for USENET news messages) • sys9 (system use) • sys10 (system use) • sys11 (system use) • sys12 (system use) • sys13 (system use) • sys14 (system use) • syslog (for Syslog messages) • user (for user programs) • uucp (UNIX to UNIX copy protocol) The default is local7.

To view nondefault settings, use the **show running-config logging** command (Figure 30) in the EXEC mode.

```
Force10#show running-config logging
!
logging buffered 524288 debugging
service timestamps log datetime msec
service timestamps debug datetime msec
!
logging trap debugging
logging facility user
logging source-interface Loopback 0
logging 10.10.10.4
Force10#
```

Figure 30 show running-config logging Command Example

synchronize log messages

You can configure FTOS to filter and consolidate the system messages for a specific line by synchronizing the message output. Only the messages with a severity at or below the set level appear. This feature works on the terminal and console connections available on the system.

To synchronize log messages, use these commands in the following sequence starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**line** {**console 0** | **vty** *number* [*end-number*] | **aux 0**}	CONFIGURATION	Enter the LINE mode. Configure the following parameters for the virtual terminal lines: • *number* range: zero (0) to 8. • *end-number* range: 1 to 8. You can configure multiple virtual terminals at one time by entering a *number* and an *end-number.*
2	**logging synchronous** [**level** *severity-level* | **all**] [*limit*]	LINE	Configure a level and set the maximum number of messages to be printed. Configure the following optional parameters: • **level** *severity-level* range: 0 to 7. Default is 2. Use the **all** keyword to include all messages. • *limit* range: 20 to 300. Default is 20.

To view the logging synchronous configuration, use the **show config** command in the LINE mode.

enable timestamp on Syslog messages

Syslog messages, by default, do not include a time/date stamp stating when the error or message was created.

To have FTOS include a timestamp with the Syslog message, use the following command syntax in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
service timestamps [**log** | **debug**] [**datetime** [**localtime**] [**msec**] [**show-timezone**] | **uptime**]	CONFIGURATION	Add timestamp to Syslog messages. Specify the following optional parameters: • **datetime**: You can add the keyword localtime to include the **localtime, msec,** and **show-timezone**. If you do not add the keyword **localtime**, the time is UTC. • **uptime**. To view time since last boot. If neither parameter is specified, FTOS configures **uptime**.

To view the configuration, use the **show running-config logging** command in the EXEC privilege mode.

To disable time stamping on Syslog messages, enter **no service timestamps** [**log** | **debug**].

SNMP

The Simple Network Management Protocol (SNMP) is used to communicate management information between the network management stations and the agents in the network elements. The FTOS supports SNMP versions 1, 2c, and 3, supporting both read-only and read-write modes. FTOS sends SNMP traps, which are messages informing the SNMP manager about the network. The software supports up to 16 SNMP trap receivers.

FTOS SNMP implementation conforms to RFC 1157 and RFC 2274 and the enterprise-specific MIB Force10-COPY-CONFIG-MIB, which supports SNMP SET requests.

Important Points to Remember

- Typically, 5-second timeout and 3-second retry values on an SNMP server are sufficient for both LAN and WAN applications. If you experience a timeout with these values, Force10 recommended best practice is to increase the timeout and retry values on your SNMP server to the following to accommodate the high port density provided by the Force10 system:

  SNMP Timeout—greater than 3 seconds

  SNMP Retry count—greater than 2 seconds
- If you want to query the chassis using SNMP v1/v2/v3 with an IPv6 address, configure the IPv6 address on an non-management port on the chassis.
- If you want to send SNMP v1/v2/v3 traps from the chassis using an IPv6 address, use a non-management port.
- SNMP v3 informs are not currently supported with IPv6 addresses.

- If you are using ACLs in SNMP v3 configuration, group ACL overrides user ACL if the user is part of that group.

Configuration Task List for SNMP

To enable SNMP on the system, enter the **snmp-server community** command. A system message appears after you enable SNMP. The following list contains configuration tasks for SNMP:

- configure access to an SNMP community on page 99 (mandatory)
- configure the system to send SNMP notifications on page 100 (mandatory)
- set SNMP information on page 102 (optional)

configure access to an SNMP community

You enable SNMP when you configure the community string to be used by the SNMP manager and agent. Without the community string set, you cannot query SNMP data.

To configure the SNMP community string, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
snmp-server community *community-string* {**ro** | **rw**} [**security-name** *name* [*access name*]]	CONFIGURATION	Enter a community string that both the SNMP manager and agent understand. Configure the following parameters: • *community-string*: some community strings are: community; public • **ro:** read-only • **rw:** read-write • **security-name**: (Optional) Enter the keyword followed by the security name as defined by the community MIB. • *access- name*: (Optional) Enter the standard access list name (a string up to 16 characters long).

To view the SNMP configuration, use the **show running-config snmp** command (Figure 31) in the EXEC privilege mode.

```
Force10#show running-config snmp
!
snmp-server enable traps bgp
snmp-server enable traps snmp
snmp-server enable traps envmon
snmp-server host 12.31.1.3 traps version 2c force10networks udp-port 162
snmp-server location labsun3
snmp-server trap-source Loopback 0
Force10#
```

Figure 31 show running-config snmp Command Example

configure the system to send SNMP notifications

SNMP traps can be collected and sent to an SNMP host (manager). Traps are not saved on the system, so to analyze the information collected in the traps, you should have the traps sent to a device or the SNMP manager. You can configure up to 16 SNMP hosts.

To configure an SNMP host to store traps, use these commands in the following sequence in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**snmp-server host** *ip-address* [**traps** | **informs**] [**version 1** | **2c** | **3**] [**auth** | **no auth** | **priv**] [*community-string*] [**udp-port** *port-number*] [*notification-type*]}	CONFIGURATION	In *ip-address*, enter an IP address of the device to store the SNMP traps. Configure at least one of the following parameters: • **traps**: (OPTIONAL) Enable all traps. • **informs**: (OPTIONAL) Enter this keyword to send inform notifications to the specified host. • **version**: (OPTIONAL) Enter the keyword followed by either **1** or **2c**. If neither is entered, the default is **1**. • **auth**: (OPTIONAL) Enter the keyword to specify authentication of a packet without encryption. • **noauth**: (OPTIONAL) Enter the keyword to specify no authentication of a packet. • **priv**: (OPTIONAL) Enter the keyword to specify both authentication and then scrambling of the packet. • *community-string*: (OPTIONAL) Enter a text string. You can also enter one of the optional notification types (**bgp**, **envmon**, **snmp**). **Note**: For version 1 and version 2c security models, this string represents the name of the SNMP community. While the string can be set using this command, it is recommended that you set the community string using the **snmp-server community** command before executing this command. For the version 3 security model, this string is the USM user security name. • **udp-port**: (OPTIONAL) Enter the keyword followed by the port number of the remote host to use. Range: 0 to 65535. Default: 162 • *notification-type*: (OPTIONAL) Enter one of the following keywords as the type of trap to be sent to the host: •**bgp** - allow BGP state change traps •**envmon** - allow environment monitor traps •**snmp** - allow SNMP-type notification (RFC 1157) traps Default: All trap types are sent to host

Step	Command Syntax	Command Mode	Purpose
2	**snmp-server enable traps** [*notification-type*] [*notification-option*]	CONFIGURATION	Enable the generation of SNMP traps. Configure up to 16 traps. Configure the optional parameters to specify which types of traps are sent: • *notification-type*: Enter one of the optional notification types (**bgp**, **envmon**, **snmp**). • *notification-option*: For the **envmon** notification-type, you can specify an additional option (**fan**, **supply**, **temperature**). For the **snmp** notification type, you can specify an additional option (**authentication**, **coldstart**, **linkdown**, **linkup**). The notification options for the snmp notification-type comply with the “generic traps” defined in RFC 1157. If you enter **snmp-server enable traps**, all traps are sent.

To view the SNMP configuration, use the **show running-config snmp** command (Figure 31) in the EXEC mode.

To delete an SNMP host configuration, use the **no snmp-server host** *ip-address* **traps** command.

To disable traps, use the **no snmp-server enable traps** [*notification-type*] [*notification-option*] command syntax.

To specify an interface to transmit the SNMP traps, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
snmp-server trap-source *interface*	CONFIGURATION	Enter the following keywords and slot/port or number information: • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a loopback interface, enter the keyword **loopback** followed by a number between 0 and 16383. • For a SONET interface, enter the keyword **sonet** followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information.

To view the configuration, use the **show running-config snmp** command syntax (Figure 31) in the EXEC privilege mode.

set SNMP information

To set the contact and location information, use either or both of the following commands in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
snmp-server contact *text*	CONFIGURATION	Specify a name or phone number up to 55 characters. Do not use spaces.
snmp-server location *text*	CONFIGURATION	Specify the location of the system with up to 255 alphanumeric characters. Do not use spaces.

To view the SNMP configuration, use the **show running-config snmp** command (Figure 31) in the EXEC mode. Other SNMP commands (for details, see Chapter 7, SNMP and Syslog Commands, in the *FTOS Command Line Interface Reference*) include:

- **show snmp:** View the status of SNMP network elements.
- **show snmp engineID:** View the identification of the local SNMP engine and all remote engines that are configured on the router.
- **show snmp group**: Display the group name, security model, view status, and storage type of each group.
- **show snmp user**: View the information configured on each SNMP user name.
- **snmp-server engineID**: Configure name for both the local and remote SNMP engines on the router.
- **snmp-server group**: Configure a new SNMP group or a table that maps SNMP users to SNMP views.
- **snmp-server user**: Configure a new user to an SNMP group.
- **snmp-server view**:snmp-server view Configure an SNMPv3 view.
- **snmp trap link-status**: Enable the interface to send SNMP link traps, which indicate whether the interface is up or down.

Copying Configuration Files Using SNMP

Use SNMP from a remote client to:

- copy the running-config file to the startup-config file, or
- copy configuration files from the Force10 system to a server
- copy configuration files from a server to the Force10 system

The relevant MIBs for these functions are:

Table 6 MIB Objects for Copying Configuration Files via SNMP

MIB Object	OID	Object Values	Description
copySrcFileType	.1.3.6.1.4.1.6027.3.5.1.1.1.2	1 = FTOS file 2 = running-config 3 = startup-config	Specifies the type of file to copy from. Valid values are: • If the copySrcFileType is running-config or startup-config, the default copySrcFileLocation is flash. • If the copySrcFileType is a binary file, the copySrcFileLocation and copySrcFileName must also be specified.
copySrcFileLocation	.1.3.6.1.4.1.6027.3.5.1.1.1.3	1 = flash 2 = slot0 3 = tftp 4 = ftp 5 = scp	Specifies the location of source file. • If the copySrcFileLocation is FTP or SCP, copyServerAddress, copyUserName, and copyUserPassword must be specified.
copySrcFileName	.1.3.6.1.4.1.6027.3.5.1.1.1.4	Path (if file is not in current directory) and filename.	Specifies name of the file. • If copySourceFileType is set to running-config or startup-config, copySrcFileName is not required.
copyDestFileType	.1.3.6.1.4.1.6027.3.5.1.1.1.5	1 = FTOS file 2 = running-config 3 = startup-config	Specifies the type of file to copy to. • If the copySourceFileType is running-config or startup-config, the default copyDestFileLocation is flash. • If the copyDestFileType is a binary the copyDestFileLocation and copyDestFileName must be specified.
copyDestFileLocation	.1.3.6.1.4.1.6027.3.5.1.1.1.6	1 = flash 2 = slot0 3 = tftp 4 = ftp 5 = scp	Specifies the location of destination file. • If the copyDestFileLocation is FTP or SCP, copyServerAddress, copyUserName, and copyUserPassword must be specified.
copyDestFileName	.1.3.6.1.4.1.6027.3.5.1.1.1.7	Path (if file is not in default directory) and filename.	Specifies the name of destination file.
copyServerAddress	.1.3.6.1.4.1.6027.3.5.1.1.1.8	IP Address of the server	The IP address of the server. • If the copyServerAddress is specified so must copyUserName, and copyUserPassword.

Table 6 MIB Objects for Copying Configuration Files via SNMP

MIB Object	OID	Object Values	Description
copyUserName	.1.3.6.1.4.1.6027.3.5.1.1.1.9	Username for the server.	Username for for the FTP, TFTP, or SCP server. • If the copyUserName is specified so must copyUserPassword.
copyUserPassword	.1.3.6.1.4.1.6027.3.5.1.1.1.10	Password for the server.	Password for for the FTP, TFTP, or SCP server.

To copy a configuration file:

Step	Task	Command Syntax	Command Mode
1	Create an SNMP community string with read/write privileges.	**snmp-server community** *community-name* **rw**	CONFIGURATION
2	Copy the *f10-copy-config.mib* MIB from the Force10 iSupport webpage to the server to which you are copying the configuration file.		
3	On the server, use the command **snmpset** as shown: **snmpset -v** *snmp-version* **-c** *community-name* **-m** *mib_path*/**f10-copy-config.mib** *force10system-ip-address mib-object.index* {**i** | **a** | **s**} *object-value...* • Every specified object must have an object value, which must be preceded by the keyword **i**. See Table 6 for valid values. • *index* must be unique to all previously executed **snmpset** commands. If an index value has been used previously, a message like the one in Message 1 appears. In this case, increment the index value and enter the command again. • Use as many MIB Objects in the command as required by the MIB Object descriptions in Table 6 to complete the command. See Table 7 or examples. **Note:** You can use the entire OID rather than the object name. Use the form: *OID.index* **i** *object-value*, as shown in Figure 33.		

Message 1 snmpset Index Value Error

```
Error in packet.
Reason: notWritable (that object does not support modification)
Failed object: FORCE10-COPY-CONFIG-MIB::copySrcFileType.101
```

Table 7 shows examples of using the command **snmpset** to copy a configuration. These examples assume that:

- the server OS is Unix
- you are using SNMP version 2c
- the community name is public, and
- the file *f10-copy-config.mib* is in the current directory or in the snmpset tool path.

Note: In Unix, enter the command **snmpset** for help using this command. Place the file *f10-copy-config.mib* the directory from which you are executing the **snmpset** command or in the snmpset tool path.

Table 7 Copying Configuration Files via SNMP

Task

Copy the running-config to the startup-config using the following command from the Unix server:

snmpset -v 2c -c public -m ./f10-copy-config.mib *force10system-ip-address* **copySrcFileType.***index* **i 2 copyDestFileType.***index* **i 3**

Figure 32 show the command syntax using MIB object names, and Figure 33 shows the same command using the object OIDs. In both cases, the object is followed by a unique index number.

Figure 32 Copying Configuration Files via SNMP using Object-Name Syntax

```
> snmpset -v 2c -r 0 -t 60 -c private -m ./f10-copy-config.mib 10.10.10.10 copySrcFileType.101
i 2 copyDestFileType.101 i 3
FORCE10-COPY-CONFIG-MIB::copySrcFileType.101 = INTEGER: runningConfig(2)
FORCE10-COPY-CONFIG-MIB::copyDestFileType.101 = INTEGER: startupConfig(3)
```

Figure 33 Copying Configuration Files via SNMP using OID Syntax

```
> snmpset -v 2c -c private -m ./f10-copy-config.mib 10.10.10.10
.1.3.6.1.4.1.6027.3.5.1.1.1.1.2.100 i 2 .1.3.6.1.4.1.6027.3.5.1.1.1.1.5.100 i 3
FORCE10-COPY-CONFIG-MIB::copySrcFileType.100 = INTEGER: runningConfig(2)
FORCE10-COPY-CONFIG-MIB::copyDestFileType.100 = INTEGER: startupConfig(3)
```

Copy the startup-config to the server via FTP using the following command from the Unix server:

snmpset -v 2c -c public -m ./f10-copy-config.mib *force10system-ip-address* **copySrcFileType.***index* **i 2 copyDestFileName.***index* **s** *filepath/filename* **copyDestFileLocation.***index* **i 4 copyServerAddress.***index* **a** *server-ip-address* **copyUserName.***index* **s** *server-login-id* **copyUserPassword.***index* **s** *server-login-password*

- *server-ip-address* must be preceded by the keyword **a**.
- values for copyUsername and copyUserPassword must be preceded by the keyword **s**.

Figure 34 Copying Configuration Files via SNMP and FTP to a Remote Server

```
> snmpset -v 2c -c private -m ./f10-copy-config.mib 10.10.10.10 copySrcFileType.110 i 2
copyDestFileName.110 s /home/startup-config copyDestFileLocation.110 i 4 copyServerAddress.110
a 11.11.11.11 copyUserName.110 s mylogin copyUserPassword.110 s mypass
FORCE10-COPY-CONFIG-MIB::copySrcFileType.110 = INTEGER: runningConfig(2)
FORCE10-COPY-CONFIG-MIB::copyDestFileName.110 = STRING: /home/startup-config
FORCE10-COPY-CONFIG-MIB::copyDestFileLocation.110 = INTEGER: ftp(4)
FORCE10-COPY-CONFIG-MIB::copyServerAddress.110 = IpAddress: 11.11.11.11
FORCE10-COPY-CONFIG-MIB::copyUserName.110 = STRING: mylogin
FORCE10-COPY-CONFIG-MIB::copyUserPassword.110 = STRING: mypass
```

Copy the startup-config to the server via TFTP using the following command from the Unix server:

Note: Verify that the file exists and its permissions are set to 777, and specify the relative path to the TFTP root directory.

Table 7 Copying Configuration Files via SNMP

Task

snmpset -v 2c -c public -m ./f10-copy-config.mib *force10system-ip-address* **copySrcFileType.***index* **i 3 copyDestFileType.***index* **i 1 copyDestFileName.***index* **s** *filepath/filename* **copyDestFileLocation.***index* **i 3 copyServerAddress.***index* **a** *server-ip-address*

Figure 35 Copying Configuration Files via SNMP and TFTP to a Remote Server

```
.snmpset -v 2c -c private -m ./f10-copy-config.mib 10.10.10.10
copySrcFileType.4 i 3
copyDestFileType.4 i 1
copyDestFileLocation.4 i 3
copyDestFileName.4 s /home/myfilename
copyServerAddress.4 a 11.11.11.11
```

Copy a binary file from the server to the startup-configuration on the Force10 system via FTP using the following command from the Unix server:

snmpset -v 2c -c public -m ./f10-copy-config.mib *force10system-ip-address* **copySrcFileType.***index* **i 1 copySrcFileLocation.***index* **i 4 copySrcFileName.***index* **s** *filepath/filename* **copyDestFileType.***index* **i 3 copyServerAddress.***index* **a** *server-ip-address* **copyUserName.***index* **s** *server-login-id* **copyUserPassword.***index* **s** *server-login-password*

Figure 36 Copying Configuration Files via SNMP and FTP from a Remote Server

```
> snmpset -v 2c -c private -m ./f10-copy-config.mib 10.10.10.10 copySrcFileType.10 i 1
copySrcFileLocation.10 i 4 copyDestFileType.10 i 3 copySrcFileName.10 s /home/myfilename
copyServerAddress.10 a 172.16.1.56 copyUserName.10 s mylogin copyUserPassword.10 s mypass
```

Force10 provides additional MIB Objects to view copy statistics. These are provided in Table 8.

Table 8 MIB Objects for Copying Configuration Files via SNMP

MIB Object	OID	Values	Description
copyState	.1.3.6.1.4.1.6027.3.5.1.1.1.11	1= running 2 = successful 3 = failed	Specifies the state of the copy operation.
copyTimeStarted	.1.3.6.1.4.1.6027.3.5.1.1.1.12	Time value	Specifies the point in the up-time clock that the copy operation started.
copyTimeCompleted	.1.3.6.1.4.1.6027.3.5.1.1.1.13	Time value	Specifies the point in the up-time clock that the copy operation completed.

Table 8 MIB Objects for Copying Configuration Files via SNMP

MIB Object	OID	Values	Description
copyFailCause	.1.3.6.1.4.1.6027.3.5.1.1.1.14	1 = bad file name 2 = copy in progress 3 = disk full 4 = file exists 5 = file not found 6 = timeout 7 = unknown	Specifies the reason the copy request failed.
copyEntryRowStatus	.1.3.6.1.4.1.6027.3.5.1.1.1.15	Row status	Specifies the state of the copy operation. Uses CreateAndGo when you are performing the copy. The state is set to active when the copy is completed.

To obtain a value for any of the MIB Objects in Table 8:

Step	Task
1	Get a copy-config MIB object value. **snmpset -v 2c -c public -m ./f10-copy-config.mib** *force10system-ip-address* [*OID.index* | *mib-object.index* • *index* is the index value used in the **snmpset** command used to complete the copy operation.
	Note: You can use the entire OID rather than the object name. Use the form: *OID.index*, as shown in Figure 38.

Figure 37 and Figure 38 are examples of using the command **snmpget** to obtain a MIB object value. These examples assume that:

- the server OS is Unix
- you are using SNMP version 2c
- the community name is public, and
- the file *f10-copy-config.mib* is in the current directory.

Note: In Unix, enter the command **snmpset** for help using this command.

Figure 37 shows the command syntax using MIB object names, and Figure 38 shows the same command using the object OIDs. In both cases, the object is followed by same index number used in the **snmpset** command.

Figure 37 Obtaining MIB Object Values for a Copy Operation using Object-name Syntax

```
>snmpget-v2c-cprivate-m./f10-copy-config.mib10.11.131.140copyTimeCompleted.110
FORCE10-COPY-CONFIG-MIB::copyTimeCompleted.110 = Timeticks: (1179831) 3:16:38.31
```

Figure 38 Obtaining MIB Object Values for a Copy Operation using OID Syntax

```
> snmpget -v 2c -c private 10.11.131.140 .1.3.6.1.4.1.6027.3.5.1.1.1.1.1.13.110
SNMPv2-SMI::enterprises.6027.3.5.1.1.1.1.1.13.110 = Timeticks: (1179831) 3:16:38.31
```

Network Time Protocol

Network Time Protocol (NTP) is defined in RFC 1305 and synchronizes timekeeping among a set of distributed time servers and clients. The protocol also coordinates time distribution in a large, diverse network with a variety of interfaces. In NTP, servers maintain the time and NTP clients synchronize with a time-serving host. NTP clients choose from among several NTP servers to determine which offers the best available source of time and the most reliable transmission of information.

In a LAN, you can configure NTP to broadcast its messages.

For more information on NTP, refer to RFC 1305, *Network Time Protocol (Version 3) Specification, Implementation and Analysis*.

Configuration Task List for NTP

Force10 Networks recommends configuring NTP for the most accurate time. In FTOS, other time sources can be configured (the hardware clock and the software clock) for a single device, but NTP clients within a network redistribute reference time via local routing algorithms and time daemons to ensure that all network devices have the correct time.

By default, NTP is not enabled on the system. Configure the **ntp server** command to enable NTP globally.

The following list includes the configuration tasks for NTP:

- specify an NTP server on page 109 (mandatory)
- configure NTP broadcasts on page 110 (optional)
- configure NTP authentication on page 110 (optional)
- set the hardware clock with NTP on page 111 (optional)
- disable NTP on an interface on page 111 (optional)
- configure a source IP address for NTP packets on page 112 (optional)

For more detailed information on the commands related to NTP, refer to .

specify an NTP server

FTOS synchronizes with a time-serving host to get the correct time. You can set FTOS to poll specific NTP time-serving hosts for the current time. From those time-serving hosts, the system chooses one NTP host with which to synchronize and serve as a client to the NTP host. As soon as a host-client relationship is established, the networking device propagates the time information throughout its local network.

Since a large number of polls to NTP time serving hosts can impact network performance, Force10 Networks recommends that you limit the number of polls in your network. Instead, configure FTOS to send NTP broadcasts to distribute the NTP information throughout the network.

To specify a time-serving host for the system, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ntp server *ip-address* [**key** *keyid*] [**prefer**] [**version** *number*]	CONFIGURATION	Configure an NTP server. Configure the IP address of a server and the following optional parameters: • **key** *keyid:* Configure a text string as the key exchanged between the NTP server and client. • **prefer:** Enter the keyword to set this NTP server as the preferred server. • **version** *number:* Enter a number 1 to 3 as the NTP version.

You can use this command to configure multiple time serving hosts, one at a time.

To view the NTP status, use the **show ntp status** command (Figure 39) in the EXEC privilege mode.

```
Force10#sh ntp sta
Clock is synchronized, stratum 2, reference is 100.10.10.10
frequency is -32.000 ppm, stability is 15.156 ppm, precision is 4294967290
reference time is BC242FD5.C7C5C000 (10:15:49.780 UTC Mon Jan 10 2000)
clock offset is clock offset msec, root delay is 0.01656 sec
root dispersion is 0.39694 sec, peer dispersion is peer dispersion msec
peer mode is client
Force10#
```

Figure 39 show ntp status Command Example

To view the configured NTP time servers and their status, use the **show ntp associations** command (Figure 40) in the EXEC privilege mode.

```
Force10#show ntp associations
   remote        ref clock     st when poll reach   delay   offset    disp
==========================================================================
 100.10.10.10    .LOCL.           1 710d   16    0    13.41    5.100 16000.0
* master (synced), # master (unsynced), + selected, - candidate
Force10#
```

Figure 40 show ntp associations Command Example

configure NTP broadcasts

With FTOS, you can receive broadcasts of time information. You can set interfaces within the system to receive NTP information through broadcast.

To configure an interface to receive NTP broadcasts, use the following commands in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ntp broadcast client	INTERFACE	Set the interface to broadcast NTP packets.

To view the NTP configuration on the interface, use the **show config** command in the INTERFACE mode.

configure NTP authentication

NTP authentication and the corresponding trusted key provide a reliable means of exchanging NTP packets with trusted time sources. NTP authentication begins when the first NTP packet is created following the configuration of keys. NTP authentication in FTOS uses the MD5 algorithm and the key is embedded in the synchronization packet that is sent to an NTP time source.

To configure NTP authentication, use these commands in the following sequence in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ntp authenticate**	CONFIGURATION	Enable NTP authentication.
2	**ntp authentication-key** *number* **md5** *key*	CONFIGURATION	Set an authentication key. Configure the following parameters: *number:* Range 1 to 4294967295. This *number* must be the same as the *number* in the **ntp trusted-key** command. *key:* Enter a text string. This text string is encrypted.

Step	Command Syntax	Command Mode	Purpose
3	**ntp trusted-key** *number*	CONFIGURATION	Define a trusted key. Configure a number from 1 to 4294967295. The *number* must be the same as the *number* used in the **ntp authentication-key** command.

To view the NTP configuration, use the **show running-config ntp** command (Figure 41) in the EXEC privilege mode. Figure 41 shows an encrypted authentication key. All keys are encrypted.

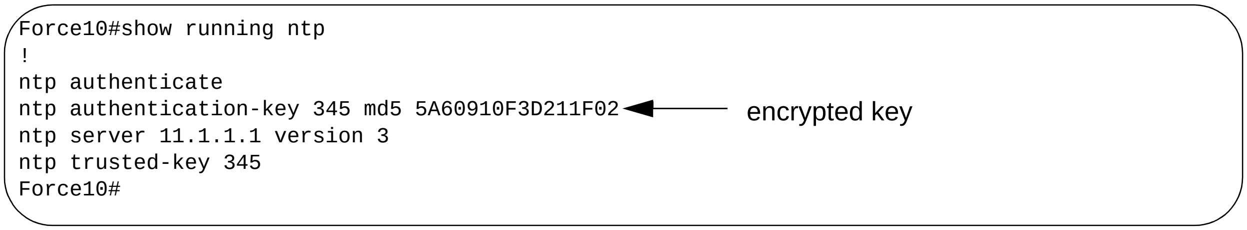

```
Force10#show running ntp
!
ntp authenticate
ntp authentication-key 345 md5 5A60910F3D211F02
ntp server 11.1.1.1 version 3
ntp trusted-key 345
Force10#
```

Figure 41 show running-config ntp Command Example

set the hardware clock with NTP

You can configure FTOS to periodically set the system hardware clock from NTP.

To set the system hardware clock from NTP, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ntp update-calendar	CONFIGURATION	Set FTOS to periodically update the hardware clock from NTP.

To view the NTP configuration, use the **show running-config ntp** command in the EXEC privilege mode.

disable NTP on an interface

By default, NTP is enabled on all active interfaces. If you disable NTP on an interface, FTOS drops any NTP packets sent to that interface.

To disable NTP on an interface, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ntp disable	INTERFACE	Disable NTP on the interface.

To re-enable NTP on an interface, enter **no ntp disable**.

To view whether NTP is configured on the interface, use the **show config** command in the INTERFACE mode. If **ntp disable** is not listed in the **show config** command output, then NTP is enabled. (The **show config** command displays only nondefault configuration information.)

configure a source IP address for NTP packets

By default, the source address of NTP packets is the IP address of the interface used to reach the network. You can configure one interface's IP address to be included in all NTP packets.

To configure an IP address as the source address of NTP packets, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ntp source *interface*	CONFIGURATION	Enter the following keywords and slot/port or number information: • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a loopback interface, enter the keyword **loopback** followed by a number between 0 and 16383. • For a port channel interface, enter the keyword **lag** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN interface, enter the keyword **vlan** followed by a number from 1 to 4094.

To view the configuration, use the **show running-config ntp** command (Figure 39) in the EXEC privilege mode.

File Transfer Services

With FTOS, you can configure the system to transfer files over the network using File Transfer Protocol (FTP). One FTP application is copying the system image files over an interface on to the system; however, FTP is not supported on VLAN interfaces.

For more information on FTP, refer to RFC 959, *File Transfer Protocol*.

Configuration Task List for File Transfer Services

The following list includes the configuration tasks for file transfer services.

- enable FTP server on page 113 (mandatory)
- configure FTP server parameters on page 113 (optional)
- configure FTP client parameters on page 114 (optional)

For a complete listing of FTP related commands, refer to .

enable FTP server

To enable the system as an FTP server, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ftp-server enable	CONFIGURATION	Enable FTP on the system.

To view FTP configuration, use the **show running-config ftp** command (Figure 42) in the EXEC privilege mode.

```
Force10#show running ftp
!
ftp-server enable
ftp-server username nairobi password 0 zanzibar
Force10#
```

Figure 42 show running-config ftp Command Output

configure FTP server parameters

After the FTP server is enabled on the system, you can configure different parameters.

To configure FTP server parameters, use any or all of the following commands in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ftp-server topdir *dir*	CONFIGURATION	Specify the directory for users using FTP to reach the system. The default is the internal flash directory.
ftp-server username *username* **password** [*encryption-type*] *password*	CONFIGURATION	Specify a user name for all FTP users and configure either a plain text or encrypted password. Configure the following optional and required parameters: • username: • *encryption-type:* Enter 0 for plain text or 7 for encrypted text. • *password:* Enter a text string.

Note: You cannot use the change directory (**cd**) command until **ftp-server topdir** has been configured.

To view the FTP configuration, use the **show running-config ftp** command in EXEC privilege mode.

configure FTP client parameters

To configure FTP client parameters, use the following commands in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ip ftp source-interface *interface*	CONFIGURATION	Enter the following keywords and slot/port or number information: • For a Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a loopback interface, enter the keyword **loopback** followed by a number between 0 and 16383. • For a port channel interface, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN interface, enter the keyword **vlan** followed by a number from 1 to 4094.
ip ftp password *password*	CONFIGURATION	Configure a password.
ip ftp username *name*	CONFIGURATION	Enter username to use on FTP client.

To view FTP configuration, use the **show running-config ftp** command (Figure 42) in the EXEC privilege mode.

Terminal Lines

By using the terminal lines in the system, you can access the system remotely and restrict access to the system by creating user profiles. The terminal lines on the system provide different means of accessing the system. The console line (console) connects you through the Console port in the RPMs. The virtual terminal lines (VTY) connect you through Telnet to the system. The auxiliary line (aux) connects secondary devices such as modems.

Configuration Task List for Terminal Lines

The following list includes the configuration tasks for terminal lines:

- enter LINE mode on page 115 (optional)
- filter traffic on a line on page 115
- configure privilege on page 116 (mandatory)
- configure password and login authentication on page 116 (mandatory)

- limit IP traffic on a terminal connection on page 118 (optional)
- set timeout on page 118 (optional)
- telnet to another network device on page 119 (optional)

For more information on commands available on the terminal lines, refer to .

enter LINE mode

By default, the terminal lines on the system are not configured and you must configure the privilege and user access. You configure the terminal lines on the system by entering the LINE mode for each type of terminal connection.

To enter the LINE mode to configure a terminal connection, use one of the following commands in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
line aux 0	CONFIGURATION	Enter the LINE mode to configure an auxiliary terminal line.
line console 0	CONFIGURATION	Enter the LINE mode for the console port.
line vty *number* [*end-number*]	CONFIGURATION	Enter the LINE mode to configure virtual terminals. FTOS supports up to 10 virtual terminals for Telnet sessions. Specify a number from 0 to 9 for the virtual terminal. To configure multiple virtual terminals, enter an end number. For example, to enter and configure virtual terminals 0 through 3, enter **line vty 0 3**.

To view the current configuration for the terminal connection, enter **show config** in the LINE mode. Figure 43 shows the configuration for three virtual terminal lines.

```
Force10(config-line-vty)#show config
line vty 0
line vty 1
line vty 2
Force10(config-line-vty)#
```

Figure 43 show config Command Example for Multiple VTY Terminal Lines

You cannot delete a terminal connection.

filter traffic on a line

Use only Standard ACLs in the **access-class** command to filter traffic on Telnet sessions.

To configure and assign an IP ACL to a line, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**line vty** *number* [*end-number*]	CONFIGURATION	Enter the LINE mode. To configure multiple virtual terminals, enter an end number. For example, to enter and configure vitriol terminals 0 through 3, enter **line vty 0 3**.
2	**ip access-group** *access-list name*	CONFIGURATION	Assign a configured Standard ACL to the line.

To view the configuration, enter the **show config** command in the LINE mode. To view the status of the ACL, enter the **show ip accounting access-list** *access-list-name* command.

configure privilege

There is no default privilege level for the terminal lines.

To set a privilege level for terminal lines, use the following command in the LINE mode:

Command Syntax	Command Mode	Purpose
privilege level *level*	LINE	Configure a level for the terminal line. Range 0 to 15. The highest level is 15.

To view the configuration, use the **show config** command in the LINE mode.

To return to the default setting (that is, no privilege level assigned to the terminal lines), enter **no privilege level** in the LINE mode.

configure password and login authentication

Use passwords and login authentication to configure access according to different user needs while protecting the system. Users access certain commands by using passwords and login authentication and the privilege command.

To configure a password and assign login authentication to a terminal connection, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**aaa authentication login** {*method-list-name* | **default**} **line**	CONFIGURATION	State that all terminal lines use the LINE mode password for authentication. Set a login authentication scheme for terminal lines by specifying a *method-list-name*. The name is configured in this command must be the same name used in Step 3. To configure the default login authentication scheme, use the **default** keyword.
2	**line** {**aux 0** | **console 0** | **vty** *number* [*end-number*]}	CONFIGURATION	Enter one or more terminal lines.
3	**login authentication** {*method-list-name* | **default**}	LINE	Use the same *method-list-name* that you entered in Step 1 or enter default. For example, if you entered **test** as the name of the authentication scheme in Step 1, enter **test** as the name in this step. This command does not appear in the LINE mode unless you configured the **aaa authentication login** command.
4	**password** *password*	LINE	Enter a text string to be used as a password. Users on that terminal line are prompted for this password.
5	**show config**	LINE	View the configuration.

Figure 44 shows the steps used to configure a password and login authentication scheme for three virtual terminals.

```
Force10(conf)#aaa authentication login suzanne line
Force10(conf)#line vty 0 2
Force10(config-line-vty)#login authent suzanne
Force10(config-line-vty)#password dilling
Force10(config-line-vty)#show confi
line vty 0
 password dilling
login authentication suzanne
line vty 1
 password dilling
login authentication suzanne
line vty 2
 password dilling
login authentication suzanne
Force10(config-line-vty)#
```

Figure 44 Commands to Configure Login Authentication and Password

limit IP traffic on a terminal connection

You can apply a standard IP ACL to a terminal line to limit IP traffic over that terminal connection.

To assign a standard IP ACL, use the following command in the LINE mode:

Command Syntax	Command Mode	Purpose
access-class *access-list-name*	LINE	Apply a standard IP ACL to a terminal connection.

To view a terminal line configuration, use the **show config** command in the LINE mode.

set timeout

As a security feature, FTOS returns to the EXEC mode after a period of inactivity on the terminal lines. You can change the amount of time before FTOS times out.

To change the time interval, use the following command in the LINE mode:

Command Syntax	Command Mode	Purpose
exec-timeout *minutes* [*seconds*]	LINE	Set the number of minutes and seconds. *minutes* range: 0 to 35791. Default 10 minutes for console line and 30 minutes for virtual terminal lines. *seconds* range: 0 to 2147483. Default is 0.

To view the configuration, use the **show config** command in the LINE mode. To return to the default values, enter **no exec-timeout**.

telnet to another network device

To telnet to a peer RPM, use the following command in EXEC or EXEC privilege mode:

Command Syntax	Command Mode	Purpose
telnet-peer-rpm	EXEC or EXEC privilege	Open a Telnet connection to the peer RPM.

To telnet to a network device with an IPv4 address, use the following command in EXEC or EXEC privilege mode:

Command Syntax	Command Mode	Purpose
telnet *ip-address*	EXEC or EXEC privilege	Open a Telnet connection to a device with an IPv4 address, where the address is in dotted decimal format (A.B.C.D)..

To telnet to a network device with an IPv6 address, use the following command in EXEC or EXEC privilege mode:

Command Syntax	Command Mode	Purpose
telnet *ip-address*	EXEC or EXEC privilege	Open a Telnet connection to a device with an IPv6 address. Enter the address is in the format: 0000:0000:0000:0000:0000:0000:0000:0000. Elision of zeros is supported.

Chapter 5 RMON

This describes the Remote Monitoring (RMON):

- Implementation on page 121
- Fault Recovery on page 122

Remote Monitoring (RMON) is an industry-standard implementation that monitors network traffic by sharing network monitoring information. RMON provides both 32-bit and 64-bit monitoring facility and long-term statistics collection on Force10 Ethernet Interfaces.

RMON operates with SNMP and monitors all nodes on a LAN segment. RMON monitors traffic passing through the router and segment traffic not destined for the router. The monitored interfaces may be chosen by using alarms and events with standard MIBs.

Implementation

Configuring RMON requires using the RMON CLI and includes the following tasks:

- Adding RMON data collection
- Removing RMON data collection
- Event settings
- Alarm settings

RMON implements the following standard RFCs (for details see Appendix D):

- RFC-2819
- RFC-3273
- RFC-3434

Fault Recovery

RMON provides the following fault recovery functions:

Interface Down—When an RMON-enabled interface goes down, monitoring continues. However, all data values are registered as 0xFFFFFFFF (32 bits) or ixFFFFFFFFFFFFFFFF (64 bits). When the interface comes back up, RMON monitoring processes resumes.

Note: A Network Management System (NMS) should be ready to interpret a down interface and plot the interface performance graph accordingly.

Line Card Down—The same as Interface Down (see above).

RPM Down, RPM Failover—Master and standby RPMs run the RMON sampling process in the background. Therefore, when an RPM goes down, the other RPM maintains the sampled data—the new master RPM provides the same sampled data as did the old master—as long as the master RPM had been running long enough to sample all the data.

NMS backs up all the long-term data collection, and displays the failover downtime from the performance graph.

Chassis Down—When a chassis goes down, all sampled data is lost. But the RMON configurations are saved in the configuration file, and the sampling process continues after the chassis returns to operation.

Platform Adaptation—RMON supports all Force10 chassis and all Force10 Ethernet Interfaces.

setting rmon alarm

To set an alarm on any MIB object, use the **rmon alarm** or **rmon hc-alarm** command in global configuration mode. To disable the alarm, use the **no** form of this command:

Command Syntax	Command Mode	Purpose
[**no**] **rmon alarm** *number variable interval* {**delta** | **absolute**} **rising-threshold** [*value event-number]* **falling-threshold** *value event-number* [**owner** *string*] or [**no**] **rmon hc-alarm** *number variable interval* {**delta** | **absolute**} **rising-threshold** *value event-number* **falling-threshold** *value event-number* [**owner** *string*]	CONFIGURATION	To set an alarm on any MIB object. Use the **no** form of this command to disable the alarm. Configure the alarm using the following optional parameters: • *number*: Alarm number, should be an integer from 1 to 65,535, the value must be unique in the RMON Alarm Table • *variable*: The MIB object to monitor—the variable must be in the SNMP OID format. For example, 1.3.6.1.2.1.1.3. The object type must be a 32-bit integer for the **rmon alarm** command and 64 bits for the **rmon hc-alarm** command. • *interval*: Time in seconds the alarm monitors the MIB variable, the value must be between 1 to 3,600. • **delta**: Tests the change between MIB variables, this is the *alarmSampleType* in the RMON Alarm table. • **absolute**: Tests each MIB variable directly, this is the *alarmSampleType* in the RMON Alarm table. • **rising-threshold** *value*: Value at which the rising-threshold alarm is triggered or reset. For the **rmon alarm** command this is a 32-bits value, for **rmon hc-alarm** command this is a 64-bits value. • *event-number*: Event number to trigger when the rising threshold exceeds its limit. This value is identical to the *alarmRisingEventIndex* in the alarmTable of the RMON MIB. If there is no corresponding rising-threshold event, the value should be zero. • **falling-threshold** *value*: Value at which the falling-threshold alarm is triggered or reset. For the **rmon alarm** command, this is a 32-bits value, for **rmon hc-alarm** command this is a 64bits value. • *event-number*: Event number to trigger when the falling threshold exceeds its limit. This value is identical to the *alarmFallingEventIndex* in the alarmTable of the RMON MIB. If there is no corresponding falling-threshold event, the value should be zero. • **owner** *string*: (Optional) Specifies an owner for the alarm, this is the alarmOwner object in the alarmTable of the RMON MIB. Default is a null-terminated string.

The following example configures an RMON alarm using the **rmon alarm** command:

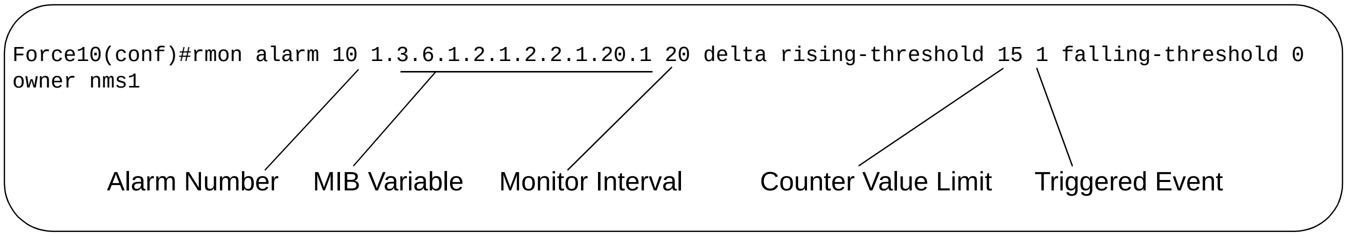

```
Force10(conf)#rmon alarm 10 1.3.6.1.2.1.2.2.1.20.1 20 delta rising-threshold 15 1 falling-threshold 0
owner nms1
```

Figure 45 rmon alarm Command Example

The above example configures RMON alarm number 10. The alarm monitors the MIB variable 1.3.6.1.2.1.2.2.1.20.1 (ifEntry.ifOutErrors) once every 20 seconds until the alarm is disabled, and checks the rise or fall of the variable. The alarm is triggered when the 1.3.6.1.2.1.2.2.1.20.1 value shows a MIB counter increase of 15 or more (such as from 100000 to 100015). The alarm then triggers event number 1, which is configured with the RMON event command. Possible events include a log entry or a SNMP trap. If the 1.3.6.1.2.1.2.2.1.20.1 value changes to 0 (falling-threshold 0), the alarm is reset and can be triggered again.

configuring an RMON event

To add an event in the RMON event table, use the **rmon event** command in global configuration mode. To disable RMON on the interface, use the **no** form of this command:

Command Syntax	Command Mode	Purpose
[**no**] **rmon event** *numb*er [*log*] [**trap** *community*] [**description** *string*] [**owner** *string*]	CONFIGURATION	*number*: Assigned event number, which is identical to the eventIndex in the eventTable in the RMON MIB. The value must be an integer from 1 to 65,535, the value must be unique in the RMON Event Table. *log*: (Optional) Generates an RMON log entry when the event is triggered and sets the eventType in the RMON MIB to log or log-and-trap. Default is no log. **trap** *community*: (Optional) SNMP community string used for this trap. Configures the setting of the eventType in the RMON MIB for this row as either snmp-trap or log-and-trap. This value is identical to the eventCommunityValue in the eventTable in the RMON MIB. Default is "public". **description** *string*: (Optional) Specifies a description of the event, which is identical to the event description in the eventTable of the RMON MIB. Default is a null-terminated string. **owner** *string*: (Optional) Owner of this event, which is identical to the eventOwner in the eventTable of the RMON MIB. Default is a null-terminated string.

The following example shows the **rmon event** command:

```
Force10(conf)#rmon event 1 log trap eventtrap description “High ifOutErrors” owner nms1
```

Figure 46 rmon event Command Example

The above configuration example creates RMON event number 1, with the description “High ifOutErrors”, and generates a log entry when the event is triggered by an alarm. The user *nms1* owns the row that is created in the event table by this command. This configuration also generates an SNMP trap when the event is triggered using the SNMP community string “eventtrap”.

configuring RMON collection statistics

To enable RMON MIB statistics collection on an interface, use the RMON collection statistics command in interface configuration mode. To remove a specified RMON statistics collection, use the **no** form of this command.

Command Syntax	Command Mode	Purpose
[**no**] **rmon collection statistics** {**controlEntry** *integer*} [**owner** *ownername*]	CONFIGURATION	**controlEntry**: Specifies the RMON group of statistics using a value. *integer*: A value from 1 to 65,535 that identifies the RMON Statistics Table. The value must be unique in the RMON Statistic Table. **owner**: (Optional) Specifies the name of the owner of the RMON group of statistics. *ownername*: (Optional) Records the name of the owner of the RMON group of statistics. Default is a null-terminated string

The following command enables the RMON statistics collection on the interface, with an ID value of 20 and an owner of “john”:

```
Force10(conf-if-mgmt)#rmon collection statistics controlEntry 20 owner john
```

Figure 47 rmon collection statistics Command Example

configuring RMON collection history

To enable the RMON MIB history group of statistics collection on an interface, use the **rmon collection history** command in interface configuration mode. To remove a specified RMON history group of statistics collection, use the **no** form of this command.

Command Syntax	Command Mode	Purpose
[**no**] **rmon collection history** {**controlEntry** *integer*} [**owner** *ownername*] [**buckets** *bucket-number*] [**interval** *seconds*]	CONFIGURATION	**controlEntry**: Specifies the RMON group of statistics using a value. *integer*: A value from 1 to 65,535 that identifies the RMON group of statistics. The value must be a unique index in the RMON History Table. **owner**: (Optional) Specifies the name of the owner of the RMON group of statistics.Default is a null-terminated string. *ownername*: (Optional) Records the name of the owner of the RMON group of statistics. **buckets**: (Optional) Specifies the maximum number of buckets desired for the RMON collection history group of statistics. *bucket-number*: (Optional) A value associated with the number of buckets specified for the RMON collection history group of statistics. The value is limited to from 1 to 1000. Default is 50 (as defined in RFC-2819). **interval**: (Optional) Specifies the number of seconds in each polling cycle. *seconds*: (Optional) The number of seconds in each polling cycle. The value is ranged from 5 to 3,600 (Seconds). Default is 1,800 as defined in RFC-2819.

enabling an RMON MIB collection history group

The following command enables an RMON MIB collection history group of statistics with an ID number of 20 and an owner of “john”, both the sampling interval and the number of buckets use their respective defaults.

```
Force10(conf-if-mgmt)#rmon collection history controlEntry 20 owner john
```

Figure 48 rmon collection history Command Example

Chapter 6 Security

Security in FTOS is based on the AAA Security model, which includes services for authentication, authorization, and accounting: FTOS also supports SSHv2 secure connections.

These features and their related configuration tasks are described in the following sections:

- AAA Accounting on page 127
- AAA Authentication on page 130
- AAA Authorization on page 133
- RADIUS on page 139
- TACACS+ on page 144
- VTY Line and Access-Class Configuration on page 148
- SCP and SSH on page 150
- Enabling and Disabling the Telnet Daemon on page 154
- Enabling and Disabling the SSH Daemon on page 154
- Trace List on page 154
- Protection Against TCP Tiny and Overlapping Fragment Attack on page 160

AAA Accounting

AAA Accounting enables tracking of services that users are accessing and the amount of network resources being consumed by those services. When AAA Accounting is enabled, the network server reports user activity to the security server in the form of accounting records. Each accounting record is comprised of accounting AV pairs and is stored on the access control server.

As with authentication and authorization, you must configure AAA Accounting by defining a named list of accounting methods, and then apply that list to various interfaces.

Configuration Task List for AAA Accounting

The following list includes the configuration tasks:

- enable AAA Accounting on page 128 (mandatory)
- suppress generation of accounting records for null username sessions on page 128 (optional)
- monitor accounting on page 129 (optional)
- accounting attribute-value pairs on page 129 (optional)
- configuring accounting for terminal lines on page 129 (optional)

For details on commands related to AAA Accounting, refer to the Security chapter in the .

enable AAA Accounting

The **aaa accounting** command enables you to create a record for any or all of the accounting functions monitored. To enable AAA accounting, perform the following task in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
aaa accounting {**system** | **exec** | **command** *level*} {*default* | *name*} {**start-stop** | **wait-start** | **stop-only**} {**tacacs+** | **radius**}	CONFIGURATION	Enable AAA Accounting and create a record for monitoring the accounting function. The variables are: • **system**—enables system accounting • **exec**—enables exec accounting • **command** *level*—enables exec accounting • *default* | *name*—Enter the name of a list of accounting methods. • **start-stop**—Use for more accounting information, to send a start-accounting notice at the beginning of the requested event and a stop-accounting notice at the end. • **wait-start**—ensures that the RADIUS or TACACS+ security server acknowledges the start notice before granting the user's process request • **stop-only**—Use for minimal accounting; instructs the specified authentication system (RADIUS or TACACS+) to send a stop record accounting notice at the end of the requested user process. • **tacacs+** | **radius**—designates the security service

suppress generation of accounting records for null username sessions

When AAA Accounting is activated, the FTOS software issues accounting records for all users on the system, including users whose username string, because of protocol translation, is NULL. An example of this is a user who comes in on a line where the AAA Authentication **login method-list none** command is applied. To prevent accounting records from being generated for sessions that do not have usernames associated with them, perform the following task in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
aaa accounting suppress null-username	CONFIGURATION	Prevent accounting records from being generated for users whose username string is NULL

monitor accounting

FTOS does not support periodic interim accounting, because the **periodic** command can cause heavy congestion when many users are logged in to the network.

No specific **show** command exists for either RADIUS or TACACS+ accounting. To obtain accounting records displaying information about users currently logged in, perform the following task in Privileged EXEC mode:

Command Syntax	Command Mode	Purpose
show accounting	CONFIGURATION	Step through all active sessions and print all the accounting records for the actively accounted functions.

accounting attribute-value pairs

The network access server monitors the accounting functions defined in either TACACS+ attribute/value (AV) pairs or RADIUS attributes, depending on which security method you have implemented.

In the following sample configuration, TACACS+-style accounting is used to track all usage of EXEC commands and commands on privilege level 15.

```
aaa accounting default exec start-stop tacacs+
```

```
aaa accounting command 15 start-stop tacacs+
```

System accounting can use only the default method list:

```
aaa accounting default system start-stop tacacs+
```

configuring accounting for terminal lines

Use the following commands to enable accounting with a named method list for a specific terminal line (where com15 and execAcct are the method list names):

```
Force10(config-line-vty)# accounting commands 15 com15
```

```
Force10(config-line-vty)# accounting exec execAcct
```

AAA Authentication

FTOS supports a distributed client/server system implemented through Authentication, Authorization, and Accounting (AAA) to help secure networks against unauthorized access. In the Force10 implementation, the E-Series acts as a RADIUS or TACACS+ client and sends authentication requests to a central RADIUS or TACACS+ server that contains all user authentication and network service access information.

Force10 Networks uses the AAA for login authentication. With AAA, you can specify the security protocol or mechanism for different login methods and different users. In FTOS, AAA uses a list of authentication methods, called method lists, to define the types of authentication and the sequence in which they are applied. You can define a method list or use the default method list. User-defined method lists take precedence over the default method list.

Configuration Task List for AAA Authentication

The following list includes the configuration tasks:

- Configure Login Authentication for Terminal Lines on page 130
- AAA Authentication Login on page 131
- AAA Authentication—RADIUS on page 132

For a complete listing of all commands related to login authentication, refer to .

Configure Login Authentication for Terminal Lines

You can assign up to five authentication methods to a method list. FTOS evaluates the methods in the order in which you enter them in each list. If the first method list does not respond or returns an error, FTOS applies the next method list until the user either passes or fails the authentication. If the user fails a method list, FTOS does not apply the next method list.

AAA Authentication Login

To configure an authentication method and method list, use these commands in the following sequence in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**aaa authentication login** {*method-list-nam*e | **default**} *method1* [... *method4*]	CONFIGURATION	Define an authentication method-list (*method-list-name*) or specify the **default**. The **default** method-list is applied to all terminal lines. Possible methods are: • **enable**—use the password defined by the **enable secret** or **enable password** command in the CONFIGURATION mode. • **line**—use the password defined by the password command in the LINE mode. • **local**—use the username/password database defined in the local configuration. • **none**—no authentication. • **radius**—use the RADIUS server(s) configured with the radius-server host command. • **tacacs+**—use the TACACS+ server(s) configured with the tacacs-server host command
2	**line** {**aux 0** | **console 0** | **vty** *number* [... *end-number*]}	CONFIGURATION	Enter the LINE mode.
3	**login authentication** {*method-list-name* | **default**}	LINE	Assign a *method-list-name* or the **default** list to the terminal line.

To view the configuration, use the **show config** command in the LINE mode or the **show running-config** in the EXEC privilege mode.

Note: Force10 Networks recommends that you use the **none** method only as a backup. This method does not authenticate users. The **none** and **enable** methods do not work with SSH.

You can create multiple method lists and assign them to different terminal lines.

AAA Authentication—Enable

To enable AAA authentication, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
aaa authentication enable {*method-list-nam*e **| default**} *method1* [... *method4*]	CONFIGURATION	• **default**—Uses the listed authentication methods that follow this argument as the default list of methods when a user logs in. • *method-list-name*—Character string used to name the list of enable authentication methods activated when a user logs in. • *method1* [... *method4*]—Any of the following: RADIUS, TACACS, enable, line, none.

If the default list is not set, only the local enable is checked. This has the same effect as issuing: **aaa authentication enable default enable**

AAA Authentication—RADIUS

To enable authentication from the RADIUS server, and use TACACS as a backup, use the following commands:

Step	Command Syntax	Command Mode	Purpose
1	**aaa authentication enable default radius tacacs**	CONFIGURATION	To enable RADIUS and to set up TACACS as backup.
2	**radius-server host x.x.x.x key some-password**	CONFIGURATION	To establish host address and password.
3	**tacacs-server host x.x.x.x key some-password**	CONFIGURATION	To establish host address and password.

To get enable authentication from the RADIUS server, and use TACACS as a backup, issue the following commands:

```
Force10(config)# aaa authentication enable default radius tacacs
Radius and TACACS server has to be properly setup for this.
Force10(config)# radius-server host x.x.x.x key <some-password>
Force10(config)# tacacs-server host x.x.x.x key <some-password>
```

To use local authentication for enable secret on console, while using remote authentication on VTY lines, perform the following steps:

```
Force10(config)# aaa authentication enable mymethodlist radius tacacs
Force10(config)# line vty 0 9
Force10(config-line-vty)# enable authentication mymethodlist
```

server-side configuration

In case of RADIUS, FTOS sends an authentication packet with the following:

```
Username: $enab15$
Password: <password-entered-by-user>
```

Therefore, the RADIUS server must have an entry for this user name.

When using TACACS+, Force10 sends an initial-packet with service type SVC_ENABLE, and then, a second packet with just the password. The tacacs-server must have an entry for username $enable$.

AAA Authorization

FTOS enables AAA new-model by default.You can set authorization to be either local or remote. Different combinations of authentication and authorization yield different results. By default, FTOS sets both to local.

Privilege Levels Overview

Limiting access to the E-Series is one method of protecting the E-Series and your network. However, at times, you might need to allow others access to the router and you can limit that access to a subset of commands. In FTOS, you can configure a privilege level for users who need limited access to the E-Series.

Every command in FTOS is assigned a privilege level of 0, 1 or 15. You can configure up to 16 privilege levels in FTOS. FTOS is pre-configured with 3 privilege levels and you can configure 13 more. The three pre-configured levels are:

- **Privilege level 1**—is the default level for the EXEC mode. At this level, you can interact with the router, for example, view some show commands and Telnet and ping to test connectivity, but you cannot configure the router. This level is often called the “user” level. One of the commands available in Privilege level 1 is the **enable** command, which you can use to enter a specific privilege level.
- **Privilege level 0**—contains only the **end**, **enable** and **disable** commands.
- **Privilege level 15**—the default level for the **enable** command, is the highest level. In this level you can access any command in FTOS.

Privilege levels 2 through 14 are not configured and you can customize them for different users and access.

After you configure other privilege levels, enter those levels by adding the level parameter after the **enable** command or by configuring a user name or password that corresponds to the privilege level. Refer to for more information on configuring user names.

By default, commands in FTOS are assigned to different privilege levels. You can access those commands only if you have access to that privilege level. For example, to reach the **protocol spanning-tree** command, you must log in to the router, enter the **enable** command for privilege level 15 (this is the default level for the command) and then enter the CONFIGURATION mode.

You can configure passwords to control access to the box and assign different privilege levels to users. FTOS supports the use of passwords when you log in to the E-Series and when you enter the **enable** command. If you move between privilege levels, you are prompted for a password if you move to a higher privilege level.

Configuration Task List for Privilege Levels

The following list includes the configuration tasks for privilege levels and passwords.

- configure a user name and password on page 134 (mandatory)
- configure enable password command on page 135 (mandatory)
- configure custom privilege levels on page 136 (mandatory)
- specify LINE mode password and privilege on page 138 (optional)
- enable and disable privilege levels on page 139 (optional)

For a complete listing of all commands related to privilege and passwords, refer to .

configure a user name and password

In FTOS, you can assign a specific user name to limit user access to the E-Series.

To configure a user name and password, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
username *name* [**access-class** *access-list-name*] [**nopassword** | **password** [*encryption-type*] *password*] [**privilege** *level*]	CONFIGURATION	Assign a user name and password. Configure the optional and required parameters: • *name:* Enter a text string up to 25 characters long. • **access-class** *access-list-name:* Enter the name of a configured IP ACL. • **nopassword:** Do not require the user to enter a password. • *encryption-type:* Enter 0 for plain text or 7 for encrypted text. • *password:* Enter a string. • **privilege** *level* range: 0 to 15.

To view user names, use the **show users** command in the EXEC privilege mode.

configure enable password command

To configure FTOS, you must use the **enable** command to enter the EXEC privilege level 15. After entering the command, FTOS requests that you enter a password. Privilege levels are not assigned to passwords, rather passwords are assigned to a privilege level. A password for any privilege level can always be changed. To change to a different privilege level, enter the **enable** command, followed by the privilege level. If you do not enter a privilege level, the default level 15 is assumed.

To configure a password for a specific privilege level, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
enable password [**level** *level*] [*encryption-mode*] *password*	CONFIGURATION	Configure a password for a privilege level. Configure the optional and required parameters: • **level** *level:* Specify a level 0 to 15. Level 15 includes all levels. • *encryption-type:* Enter 0 for plain text or 7 for encrypted text. • *password:* Enter a string. To change only the password for the enable command, configure only the *password* parameter.

To view the configuration for the **enable secret** command, use the **show running-config** command in the EXEC privilege mode.

In custom-configured privilege levels, the **enable** command is always available. No matter what privilege level you entered FTOS, you can enter the **enable 15** command to access and configure all CLI.

configure custom privilege levels

In addition to assigning privilege levels to the user, you can configure the privilege levels of commands so that they are visible in different privilege levels. Within FTOS, commands have certain privilege levels. With the privilege command, the default level can be changed or you can reset their privilege level back to the default.

- Assign the launch keyword (for example, **configure**) for the keyword's command mode.
- If you assign only the first keyword to the privilege level, all commands beginning with that keyword are also assigned to the privilege level. If you enter the entire command, the software assigns the privilege level to that command only.

To assign commands and passwords to a custom privilege level, you must be in privilege level 15 and use these commands in the following sequence in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**username** *name* [**access-class** *access-list-name*] [**privilege** *level*] [**nopassword** | **password** [*encryption-type*] *password*]	CONFIGURATION	Assign a user name and password. Configure the optional and required parameters: • *name:* Enter a text string. • **access-class** *access-list-name:* Enter the name of a configured IP ACL. • **privilege** *level* range: 0 to 15. • **nopassword:** Do not require the user to enter a password. • *encryption-type:* Enter 0 for plain text or 7 for encrypted text. • *password:* Enter a string.
2	**enable password** [**level** *level*] [*encryption-mode*] *password*	CONFIGURATION	Configure a password for privilege level. Configure the optional and required parameters: • **level** *level:* Specify a level 0 to 15. Level 15 includes all levels. • *encryption-type:* Enter 0 for plain text or 7 for encrypted text. • *password:* Enter a string up to 25 characters long. To change only the password for the enable command, configure only the *password* parameter.

Step	Command Syntax	Command Mode	Purpose
3	**privilege** *mode* {**level** *level command* | **reset** *command*}	CONFIGURATION	Configure level and commands for a mode or reset a command's level. Configure the following required and optional parameters: • *mode:* Enter a keyword for the modes (exec, configure, interface, line, route-map, router) • **level** *level* range: 0 to 15. Levels 0, 1 and 15 are pre-configured. Levels 2 to 14 are available for custom configuration. • *command:* A FTOS CLI keyword (up to 5 keywords allowed). • **reset:** Return the command to its default privilege mode.

To view the configuration, use the **show running-config** command in the EXEC privilege mode.

Figure 49 is an example of a configuration to allow a user "john" to view only the EXEC mode commands and all **snmp-serve**r commands. Since the **snmp-server** commands are "enable" level commands and, by default, found in the CONFIGURATION mode, you must also assign the launch command for the CONFIGURATION mode, **configure**, to the same privilege level as the **snmp-server** commands.

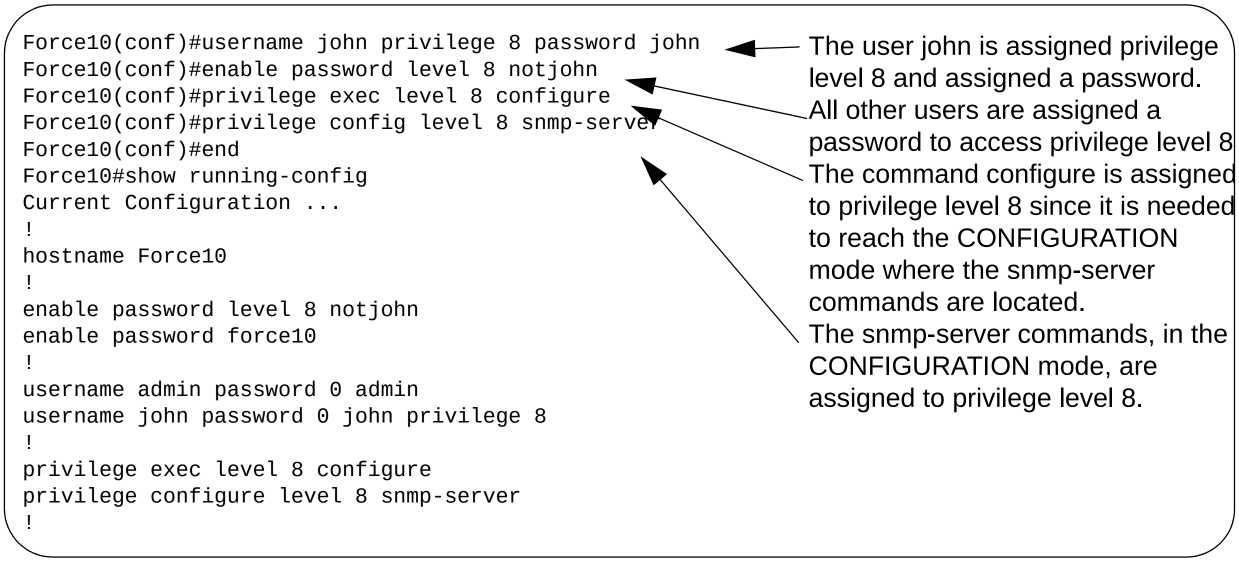

```
Force10(conf)#username john privilege 8 password john
Force10(conf)#enable password level 8 notjohn
Force10(conf)#privilege exec level 8 configure
Force10(conf)#privilege config level 8 snmp-server
Force10(conf)#end
Force10#show running-config
Current Configuration ...
!
hostname Force10
!
enable password level 8 notjohn
enable password force10
!
username admin password 0 admin
username john password 0 john privilege 8
!
privilege exec level 8 configure
privilege configure level 8 snmp-server
!
```

The user john is assigned privilege level 8 and assigned a password.

All other users are assigned a password to access privilege level 8

The command configure is assigned to privilege level 8 since it is needed to reach the CONFIGURATION mode where the snmp-server commands are located.

The snmp-server commands, in the CONFIGURATION mode, are assigned to privilege level 8.

Figure 49 Configuring a Custom Privilege Level

Figure 50 is a screen shot of the Telnet session for user “john”. The **show privilege** command output confirms that “john” is in privilege level 8. In the EXEC privilege mode, “john” can access only the commands listed. In the CONFIGURATION mode, “john” can access only the **snmp-server** commands.

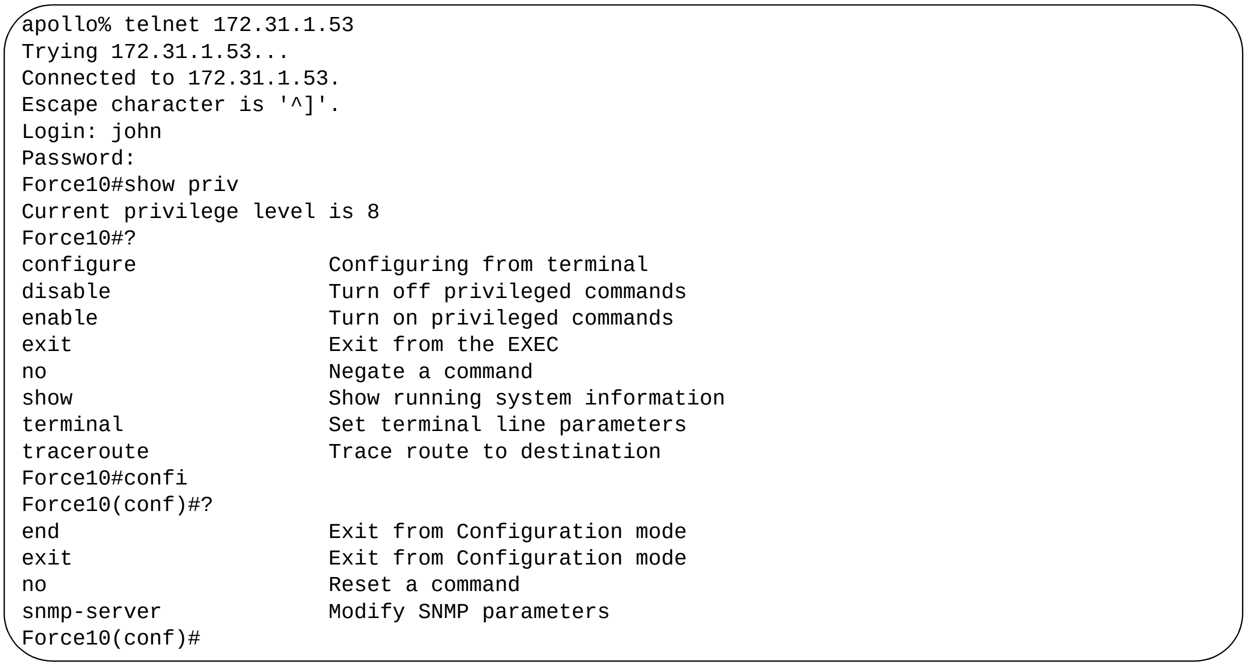

```
apollo% telnet 172.31.1.53
Trying 172.31.1.53...
Connected to 172.31.1.53.
Escape character is '^]'.
Login: john
Password:
Force10#show priv
Current privilege level is 8
Force10#?
configure               Configuring from terminal
disable                 Turn off privileged commands
enable                  Turn on privileged commands
exit                    Exit from the EXEC
no                      Negate a command
show                    Show running system information
terminal                Set terminal line parameters
traceroute              Trace route to destination
Force10#confi
Force10(conf)#?
end                     Exit from Configuration mode
exit                    Exit from Configuration mode
no                      Reset a command
snmp-server             Modify SNMP parameters
Force10(conf)#
```

Figure 50 User john’s Login and the List of Available Commands

specify LINE mode password and privilege

You can specify a password authentication of all users on different *terminal* lines. The user’s privilege level will be the same as the privilege level assigned to the terminal line, unless a more specific privilege level is is assigned to the user.

To specify a password for the terminal line, use the following commands, in any order, in the LINE mode:

Command Syntax	Command Mode	Purpose
privilege level *level*	LINE	Configure a custom privilege level for the terminal lines. • **level** *level* range: 0 to 15. Levels 0, 1 and 15 are pre-configured. Levels 2 to 14 are available for custom configuration.
password [*encryption-type*] *password*	LINE	Specify either a plain text or encrypted password. Configure the following optional and required parameters: • *encryption-type*: Enter 0 for plain text or 7 for encrypted text. • *password*: Enter a text string up to 25 characters long.

To view the password configured for a terminal, use the **show config** command in the LINE mode.

enable and disable privilege levels

Enter the **enable** or **enable privilege-level** command in the EXEC privilege mode to set a user's security level. If you do not enter a privilege level, FTOS sets it to 15 by default.

To move to a lower privilege level, enter the command **disable** followed by the **level-number** you wish to set for the user in the EXEC privilege mode. If you enter **disable** without a level-number, your security level is 1.

RADIUS

Remote Authentication Dial-In User Service (RADIUS) is a distributed client/server protocol. This protocol transmits authentication, authorization, and configuration information between a central RADIUS server and a RADIUS client (the E-Series). The E-Series sends user information to the RADIUS server and requests authentication of the user and password. The RADIUS server returns one of the following responses:

- **Access-Accept**—the RADIUS server authenticates the user
- **Access-Reject**—the RADIUS server does not authenticate the user

If an error occurs in the transmission or reception of RADIUS packets, the error can be viewed by enabling the **debug radius** command.

Transactions between the RADIUS server and the client are encrypted (the users' passwords are not sent in plain text). RADIUS uses UDP as the transport protocol between the RADIUS server host and the client.

For more information on RADIUS, refer to RFC 2865, *Remote Authentication Dial-in User Service*.

RADIUS Authentication and Authorization

FTOS supports RADIUS for user authentication (text password) at login and can be specified as one of the login authentication methods in the **aaa authentication login** command.

When configuring AAA authorization, you can configure to limit the attributes of services available to a user. When authorization is enabled, the network access server uses configuration information from the user profile to issue the user's session. The user's access is limited based on the configuration attributes.

RADIUS exec-authorization stores a user-shell profile and that is applied during user login. You may name the relevant named-lists with either a unique name or the default name. When authorization is enabled by the RADIUS server, the server returns the following information to the client:

- Idle time
- ACL configuration information
- Auto-command
- Privilege level

After gaining authorization for the first time, you may configure these attributes.

Note: RADIUS authentication/authorization is done for every login. There is no difference between first-time login and subsequent logins.

Idle Time

Every session line has its own idle-time. If the idle-time value is not changed, the default value of 30 minutes is used. RADIUS specifies idle-time allow for a user during a session before timeout. When a user logs in, the lower of the two idle-time values (configured or default) is used. The idle-time value is updated if both of the following happens:

- The administrator changes the idle-time of the line on which the user has logged in
- The idle-time is lower than the RADIUS-returned idle-time

ACL

The RADIUS server can specify an ACL. If an ACL is configured on the RADIUS server, and if that ACL is present, user may be allowed access based on that ACL. If the ACL is absent, authorization fails, and a message is logged indicating the this.

RADIUS can specify an ACL for the user if both of the following are true:

- If an ACL is absent
- There is a very long delay for an entry, or a denied entry because of an ACL, and a message is logged

Note: The ACL name must be a string. Only standard ACLs in authorization (both RADIUS and TACACS) are supported. Authorization is denied in cases using Extended ACLs.

Auto-command

You may configure the system through the RADIUS server to automatically execute a command when you connect to a specific line. To do this, use the command **auto-command**. The auto-command is executed when the user is authenticated and before the prompt appears to the user.

Privilege Level

Through the RADIUS server, you can use the command **privilege level** to configure a privilege level for the user to enter into when they connect to a session.This value is configured on the client system.

Configuration Task List for RADIUS

To authenticate users using RADIUS, at least one RADIUS server must be specified so that the E-Series cab communicate with and configure RADIUS as one of your authentication methods.

The following list includes the configuration tasks for RADIUS.

- define an aaa method list to be used for RADIUS on page 141 (mandatory)
- apply the method list to terminal lines on page 142 (mandatory except when using default lists)
- specify a RADIUS server host on page 142 (mandatory)
- set global communication parameters for all RADIUS server hosts on page 143 (optional)
- monitor RADIUS on page 144 (optional)

For a complete listing of all commands related to RADIUS, refer to .

➡ **Note: RADIUS authentication and authorization are done in a single step. Hence, authorization cannot be used independent of authentication. However, if RADIUS authorization is configured and authentication is not, then a message is logged stating this. During authorization, the next method in the list (if present) is used, or if another method is not present, an error is reported.**

To view the configuration, use the **show config** in the LINE mode or the **show running-config** command in the EXEC privilege mode.

define an aaa method list to be used for RADIUS

To configure RADIUS to authenticate or authorize users on the E-Series, you must create an AAA method list. Default-method-lists do not need to be explicitly applied to the line, hence, they are not-mandatory. To create a method list, enter either one of the following commands in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
aaa authentication login *method-list-name* radius	CONFIGURATION	Enter a text string (up to 16 characters long) as the name of the method list you wish to use with the RADIUS authentication method.
aaa authorization exec {*method-list-name* | **default**} **radius tacacs+**	CONFIGURATION	Create methodlist with RADIUS and TACACS+ as authorization methods. Typical order of methods: RADIUS, TACACS+, Local, None. If authorization is denied by RADIUS, the session ends (**radius** should not be the last method specified).

apply the method list to terminal lines

To enable RADIUS AAA login authentication for a method list, you must apply it to a terminal line. To configure a terminal line for RADIUS authentication and authorization, enter the following commands:

Command Syntax	Command Mode	Purpose
line {**aux 0** | **console 0** | **vty** *number* [*end-number*]}	CONFIGURATION	Enter the LINE mode.
login authentication {*method-list-name* | **default**}	LINE	Enable AAA login authentication for the specified RADIUS method list. This procedure is mandatory if you are not using default lists.
authorization exec *methodlist*	CONFIGURATION	To use the methodlist.

specify a RADIUS server host

When configuring a RADIUS server host, you can set different communication parameters, such as the UDP port, the key password, the number of retries, and the timeout.

To specify a RADIUS server host and configure its communication parameters, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
radius-server host {*hostname* | *ip-address*} [**auth-port** *port-number*] [**retransmit** *retries*] [**timeout** *seconds*] [**key** [*encryption-type*] *key*]	CONFIGURATION	Enter the host name or IP address of the RADIUS server host. Configure the optional communication parameters for the specific host: • **auth-port** *port-number* range: 0 to 65335. Enter a UDP port number. The default is 1812. • **retransmit** *retries* range: 0 to 100. Default is 3. • **timeout** *seconds* range: 0 to 1000. Default is 5 seconds. • **key** [*encryption-type*] *key:* Enter 0 for plain text or 7 for encrypted text, and a string for the key. This key must match the key configured on the RADIUS server host. If these optional parameters are not configured, the global default values for all RADIUS host are applied.

To specify multiple RADIUS server hosts, configure the **radius-server host** command multiple times. If multiple RADIUS server hosts are configured, FTOS attempts to connect with them in the order in which they were configured. When FTOS attempts to authenticate a user, the software connects with the RADIUS server hosts one at a time, until a RADIUS server host responds with an accept or reject response.

If you want to change an optional parameter setting for a specific host, use the **radius-server host** command. To change the global communication settings to all RADIUS server hosts, refer to set global communication parameters for all RADIUS server hosts on page 143.

➡ **Note:** You can configure global communication parameters (auth-port, key, retransmit, and timeout parameters) and specific host communication parameters on the same E-Series. However, if both global and specific host parameters are configured, the specific host parameters will override the global parameters for that RADIUS server host.

To view the RADIUS configuration, use the **show running-config radius** command in the EXEC privilege mode.

To delete a RADIUS server host, use the **no radius-server host** {*hostname* | *ip-address*} command.

set global communication parameters for all RADIUS server hosts

You can configure global communication parameters (auth-port, key, retransmit, and timeout parameters) and specific host communication parameters on the same E-Series. However, if both global and specific host parameters are configured, the specific host parameters override the global parameters for that RADIUS server host.

To set global communication parameters for all RADIUS server hosts, use any or all of the following commands in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
radius-server deadtime *seconds*	CONFIGURATION	Set a time interval after which a RADIUS host server is declared dead. • *seconds* range: 0 to 2147483647. Default 0 seconds.
radius-server key [*encryption-type*] *key*	CONFIGURATION	Configure a key for all RADIUS communications between the E-Series and RADIUS server hosts. • *encryption-type:* Enter 7 to encrypt the password. Enter 0 to keep the password as plain text. • *key:* Enter a string. You cannot use spaces in the key.
radius-server retransmit *retries*	CONFIGURATION	Configure the number of times FTOS retransmits RADIUS requests. • *retries* range: 0 to 100. Default is 3 retries.
radius-server timeout *seconds*	CONFIGURATION	Configure the time interval the E-Series waits for a RADIUS server host response. • *seconds* range: 0 to 1000. Default is 5 seconds.

To view the configuration of RADIUS communication parameters, use the **show running-config** command in the EXEC privilege mode.

monitor RADIUS

To view information on RADIUS transactions, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Purpose
debug radius	EXEC privilege	View RADIUS transactions to troubleshoot problems.

TACACS+

FTOS supports Terminal Access Controller Access Control System (TACACS+ client, including support for login authentication.

Configuration Task List for TACACS+

The following list includes the configuration task for TACACS+ functions:

- TACACS+ Remote Authentication and Authorization on page 147
- Specifying a TACACS+ Server Host on page 147
- Choosing TACACS+ as Authentication Method on page 144

For a complete listing of all commands related to TACACS+, refer to the *FTOS Command Line Reference*.

Choosing TACACS+ as Authentication Method

One of the login authentication methods available is TACACS+ and the user’s name and password are sent for authentication to the TACACS hosts specified.To use TACACS+ to authenticate users, you must specify at least one TACACS+ server for the E-Series to communicate with and configure TACACS+ as one of your authentication methods.

To select TACACS as the login authentication method, use these commands in the following sequence in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**tacacs-server host** {*ip-address* | *host*}	CONFIGURATION	Configure a TACACS+ server host. Enter the IP address or host name of the TACACS+ server. Use this command multiple times to configure multiple TACACS+ server hosts.
2	**aaa authentication login** {*method-list-name* | **default**} **tacacs+** [...*method3*]	CONFIGURATION	Create a method-list-name and specify that TACACS+ is the method for login authentication. The **tacacs+** method should not be the last method specified.
3	**line** {**aux 0** | **console 0** | **vty** *number* [*end-number*]}	CONFIGURATION	Enter the LINE mode.
4	**login authentication** {*method-list-name* | **default**}	LINE	Assign the *method-list* to the terminal line.

To view the configuration, use the **show config** in the LINE mode or the **show running-config tacacs+** command in the EXEC privilege mode.

If authentication fails using the primary method, FTOS employs the second method (or third method, if necessary) automatically. For example, if the TACACS+ server is reachable, but the server key is invalid, FTOS proceeds to the next authentication method. In Figure 51, the TACACS+ is incorrect, but the user is still authenticated by the secondary method.

Figure 51 Failed Authentication

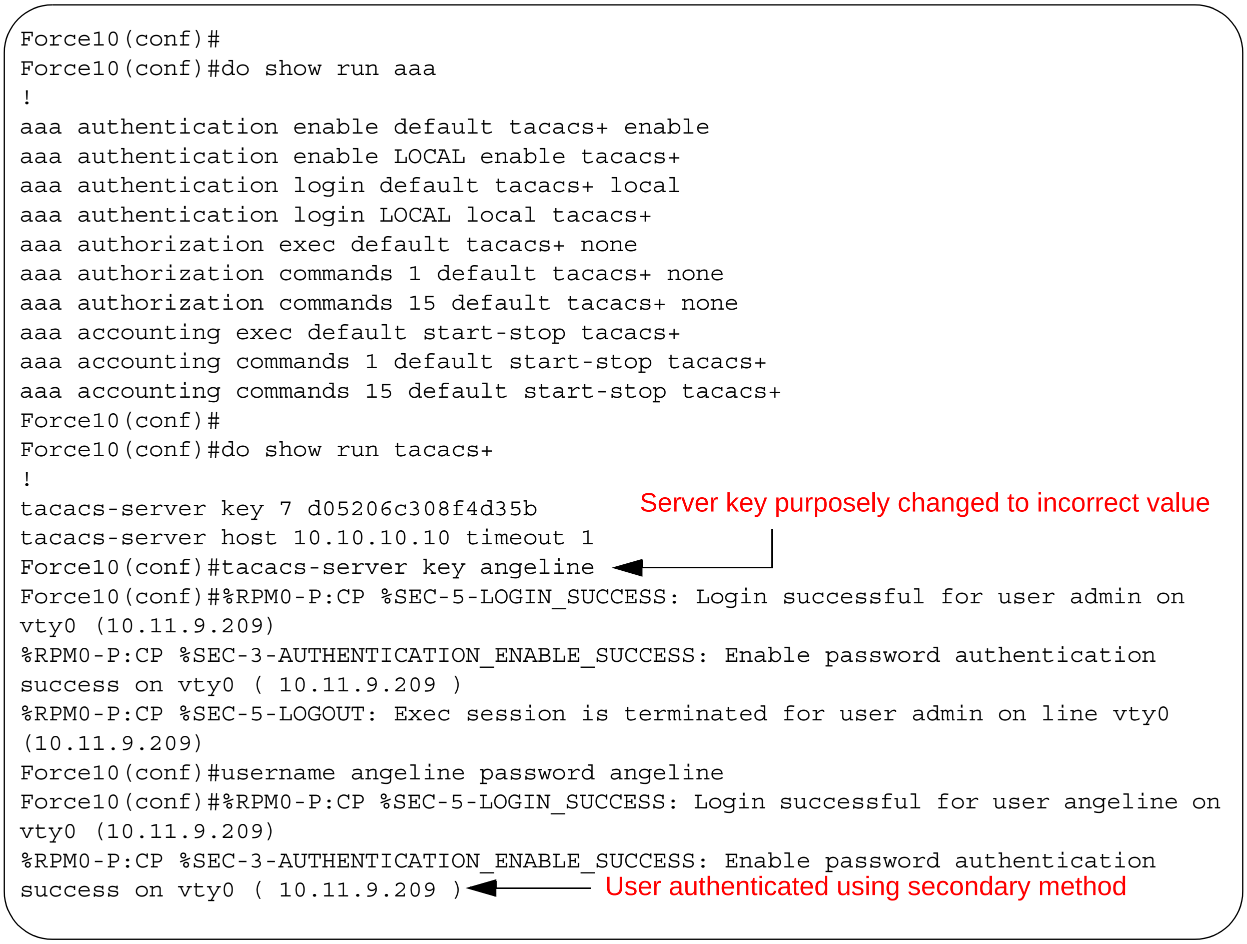

```
Force10(conf)#
Force10(conf)#do show run aaa
!
aaa authentication enable default tacacs+ enable
aaa authentication enable LOCAL enable tacacs+
aaa authentication login default tacacs+ local
aaa authentication login LOCAL local tacacs+
aaa authorization exec default tacacs+ none
aaa authorization commands 1 default tacacs+ none
aaa authorization commands 15 default tacacs+ none
aaa accounting exec default start-stop tacacs+
aaa accounting commands 1 default start-stop tacacs+
aaa accounting commands 15 default start-stop tacacs+
Force10(conf)#
Force10(conf)#do show run tacacs+
!
tacacs-server key 7 d05206c308f4d35b
tacacs-server host 10.10.10.10 timeout 1
Force10(conf)#tacacs-server key angeline
Force10(conf)#%RPM0-P:CP %SEC-5-LOGIN_SUCCESS: Login successful for user admin on
vty0 (10.11.9.209)
%RPM0-P:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable password authentication
success on vty0 ( 10.11.9.209 )
%RPM0-P:CP %SEC-5-LOGOUT: Exec session is terminated for user admin on line vty0
(10.11.9.209)
Force10(conf)#username angeline password angeline
Force10(conf)#%RPM0-P:CP %SEC-5-LOGIN_SUCCESS: Login successful for user angeline on
vty0 (10.11.9.209)
%RPM0-P:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable password authentication
success on vty0 ( 10.11.9.209 )
```

Server key purposely changed to incorrect value

User authenticated using secondary method

monitor TACACS+

To view information on TACACS+ transactions, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Purpose
debug tacacs+	EXEC privilege	View TACACS+ transactions to troubleshoot problems.

TACACS+ Remote Authentication and Authorization

FTOS takes the access class from the TACACS+ server. Access class is the class of service that restricts Telnet access and packet sized. If you have configured remote authorization, then FTOS ignores the access class you have configured for the VTY line. FTOS instead gets this access class information from the TACACS+ server. FTOS needs to know the username and password of the incoming user before it can fetch the access class from the server. A user, therefore, will at least see the login prompt. If the access class denies the connection, FTOS closes the Telnet session immediately.

Specifying a TACACS+ Server Host

When configuring a TACACS+ server host, you can set different communication parameters, such as the the key password.

To specify a TACACS+ server host and configure its communication parameters, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
tacacs-server host {*hostname* | *ip-address*} [**port** *port-number*] [**timeout** *seconds*] [**key** *key*]	CONFIGURATION	Enter the host name or IP address of the TACACS+ server host. Configure the optional communication parameters for the specific host: • **port** *port-number* range: 0 to 65335. Enter a TCP port number. The default is 49. • **timeout** *seconds* range: 0 to 1000. Default is 10 seconds. • **key** *key:* Enter a string for the key. This key must match a key configured on the TACACS+ server host. This parameter should be the last parameter configured. If these optional parameters are not configured, the default global values are applied.

To specify multiple TACACS+ server hosts, configure the **tacacs-server host** command multiple times. If multiple TACACS+ server hosts are configured, FTOS attempts to connect with them in the order in which they were configured.

To view the TACACS+ configuration, use the **show running-config tacacs+** command in the EXEC privilege mode.

To delete a TACACS+ server host, use the **no tacacs-server host** {*hostname* | *ip-address*} command.

```
freebsd2# telnet 2200:2200:2200:2200:2200::2202
Trying 2200:2200:2200:2200:2200::2202...
Connected to 2200:2200:2200:2200:2200::2202.
Escape character is '^]'.
Login: admin
Password:
Force10#
Force10#
!-The prompt is returned as the connection is authenticated.
```

Command Authorization

The AAA command authorization feature configures FTOS to send each configuration command to a TACACS server for authorization before it is added to the running configuration.

By default, the command AAA authorization commands configures the system to check both EXEC level and CONFIGURATION level commands. Use the command **no aaa authorization config-commands** to enable only EXEC-level command checking.

If rejected by the AAA server, the command is not added ot the running configuration, and messages similar to Message 2 are displayed.

Message 2 Configuration Command Rejection

```
04:07:48: %RPM0-P:CP %SEC-3-SEC_AUTHORIZATION_FAIL: Authorization failure  Command
authorization failed for user (denyall) on vty0 ( 10.11.9.209 )
```

VTY Line and Access-Class Configuration

The Force10 Operating System provides several ways to configure access classes for VTY lines, including:

- VTY Line Local Authentication and Authorization on page 148
- VTY Line Remote Authentication and Authorization on page 149

VTY Line Local Authentication and Authorization

FTOS retrieves the access class from the local database. To use this feature:

1. Create a username
2. Enter a password
3. Assign an access class
4. Enter a privilege level

FTOS can assign different access classes to different users by username. Until the user attempts to login, FTOS does not know if they will be assigned a VTY line. This means that an incoming user always sees a login prompt even if you have excluded them from the VTY line with a **deny-all** access class. Once the user identifies themselves, FTOS retrieves the access class from the local database and applies it. (FTOS also subsequently can close the connection if the user is denied access).

.

```
Force10(conf)#user gooduser password abc privilege 10 access-class permitall
Force10(conf)#user baduser password abc privilege 10 access-class denyall
Force10(conf)#
Force10(conf)#aaa authentication login localmethod local
Force10(conf)#
Force10(conf)#line vty 0 9
Force10(config-line-vty)#login authentication localmethod
Force10(config-line-vty)#end
```

Figure 52 Example Access-Class Configuration Using Local Database

VTY Line Remote Authentication and Authorization

FTOS retrieves the access class from the VTY line.

The Force10 OS takes the access class from the VTY line and applies it to ALL users. FTOS does not need to know the identity of the incoming user and can immediately apply the access class. If the authentication method is radius, TACACS+, or line, and you have configured an access class for the VTY line, FTOS immediately applies it. If the access-class is **deny all** or **deny for the incoming subnet**, FTOS closes the connection without displaying the login prompt.

```
Force10(conf)#ip access-list standard deny10
Force10(conf-ext-nacl)#permit 10.0.0.0/8
Force10(conf-ext-nacl)#deny any
Force10(conf)#
Force10(conf)#aaa authentication login tacacsmethod tacacs+
Force10(conf)#tacacs-server host 256.1.1.2 key force10
Force10(conf)#
Force10(conf)#line vty 0 9
Force10(config-line-vty)#login authentication tacacsmethod
Force10(config-line-vty)#
Force10(config-line-vty)#access-class deny10
Force10(config-line-vty)#end
(same applies for radius and line authentication)
```

Figure 53 Example Access Class Configuration Using TACACS+ Without Prompt

VTY MAC-SA Filter Support

FTOS supports MAC access lists which permit or deny users based on their source MAC address. With this approach, you can implement a security policy based on the source MAC address.

To apply a MAC ACL on a VTY line, use the same **access-class** command as IP ACLs (Figure 54).

Figure 54 Example Access Class Configuration Using TACACS+ Without Prompt

```
Force10(conf)#mac access-list standard sourcemac
Force10(config-std-mac)#permit 00:00:5e:00:01:01
Force10(config-std-mac)#deny any
Force10(conf)#
Force10(conf)#line vty 0 9
Force10(config-line-vty)#access-class sourcemac
Force10(config-line-vty)#end
```

SCP and SSH

Secure Shell (SSH) is a protocol for secure remote login and other secure network services over an insecure network. FTOS is compatible with SSH versions 1.5 and 2, both the client and server modes. SSH sessions are encrypted and use authentication. For details on command syntax, see the Security chapter in the *FTOS Command Line Interface Reference*.

SCP is a remote file copy program that works with SSH and is supported by FTOS.

To use the SSH client, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Purpose
ssh {*hostname*} [**-l** *username* | **-p** *port-number* | **-v** {**1** | **2**}	EXEC privilege	Open an SSH connection specifying the hostname, username, port number, and version of the SSH client. *hostname* is the IP address or hostname of the remote device. • Enter an IPv4 address in dotted decimal format (A.B.C.D), • Enter an IPv6 address in hexadecimal format (0000:0000:0000:0000:0000:0000:0000:0000). Elision of zeros is supported.

To enable the SSH server for version 1 and 2, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ip ssh server {enable | port *port-number* }	CONFIGURATION	To configure the E-Series as an SCP/SSH server, use this command.

To enable the SSH server for version 1 or 2 only, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ip ssh server version {1|2}	CONFIGURATION	To configure the Force10 system as an SSH server that uses only version 1 or 2, use this command.

Use the command **show ip ssh** to confirm your settings. Figure show that only version 2 is enabled using the command **ip ssh server version 2**.

```
Force10(conf)#ip ssh server version 2
Force10(conf)#do show ip ssh
SSH server                : disabled.
SSH server version        : v2.
Password Authentication   : enabled.
Hostbased Authentication  : disabled.
RSA       Authentication  : disabled.
   Vty          Encryption      Remote IP
```

Figure 55 Specifying an SSH version

To disable SSH server functions, enter **no ip ssh server enable**.

To view your SSH configuration, use the following command in EXEC privilege mode:

Command Syntax	Command Mode	Purpose
show ip ssh	EXEC privilege	Display SSH connection information.

Other SSH-related commands include:

- **crypto key generate**: Generate keys for the SSH server.
- **debug ip ssh:** Enables collecting SSH debug information.
- **ip scp topdir:** Identify a location for files used in secure copy transfer.
- **ip ssh authentication-retries:** Configure the maximum number of attempts that should be used to authenticate a user.
- **ip ssh connection-rate-limit:** Configure the maximum number of incoming SSH connections per minute.
- **ip ssh hostbased-authentication enable:** Enable hostbased-authentication for the SSHv2 server.
- **ip ssh key-size:** Configure the size of the server-generated RSA SSHv1 key.
- **ip ssh password-authentication enable:** Enable password authentication for the SSH server.
- **ip ssh pub-key-file:** Specify the file to be used for host-based authentication.
- **ip ssh rhostsfile:** Specify the rhost file to be used for host-based authorization.

- **ip ssh rsa-authentication enable:** Enable RSA authentication for the SSHv2 server.
- **ip ssh rsa-authentication:** Add keys for the RSA authentication.
- **show crypto:** Display the public part of the SSH host-keys.
- **show ip ssh client-pub-keys:** Display the client public keys used in host-based authentication.
- **show ip ssh rsa-authentication:** Display the authorized-keys for the RSA authentication.
- **ssh-peer-rpm**: Open an SSH connection to the peer RPM.

SSH with IPv6

FTOS supports both inbound and outbound SSH sessions using IPv6 addressing. Inbound SSH supports accessing the system through the management interface as well as through a physical Layer 2 or Layer 3 interface.

Figure 56 illustrates an outbound SSH connection to a Unix server.

Figure 56 Outbound SSH with IPv6

```
Force10#ssh 2200:2200:2200:2200:2200::2201 -l admin -v 2 -p 22
Trying 2200:2200:2200:2200:2200::2201...
Password:
Last login: Mon Jun  4 04:26:29 2007 from 10.11.9.209
Copyright (c) 1980, 1983, 1986, 1988, 1990, 1991, 1993, 1994
        The Regents of the University of California.  All rights reserved.
FreeBSD 4.8-RELEASE (GENERIC) #0: Thu Apr  3 10:53:38 GMT 2003
Welcome to FreeBSD!
….
freebsd2>
Nice tcsh prompt: set prompt = '%n@%m:%~%# '
> exit
logout
Force10#
The following figure illustrates an inbound IPv6 SSH session:
freebsd2# admin@2200:2200:2200:2200:2200::2202
admin@2200:2200:2200:2200:2200::2202's password:
Force10#
Force10#exit
Connection to 2200:2200:2200:2200:2200::2202 closed by remote host.
Connection to 2200:2200:2200:2200:2200::2202 closed.
```

Telnet with IPv6

You can use the standard **telnet** command to access a system configured with IPv6 addresses. Telnet to IPv6 link local addresses is not supported.

Outbound Telnet will use the physical interface through which the packet is sent as the source IPv6 address. Configuring IP source interface via address on the loopback is not supported.

Figure 57 shows an outbound Telnet connection.

Figure 57 Outbound Telnet with IPv6

```
Force10#telnet 2200:2200:2200:2200:2200::2201
Trying 2200:2200:2200:2200:2200::2201...
Connected to 2200:2200:2200:2200:2200::2201.
Exit character is '^]'.
FreeBSD/i386 (freebsd2.force10networks.com) (ttyp1)
login: admin
Force10#
```

The following sample configuration uses TACACS+ authentication with SSH.

<table>
<tr><th>Step</th><th>Task</th></tr>
<tr><td>1</td><td>Enable Telnet server functionality.</td></tr>
<tr><td>2</td><td>Configure a AAA method list specifying TACACS+ authentication and authorization.
<pre>
Force10#show run aaa
!
aaa authentication login tacmethod tacacs+
aaa authorization exec tacmethod tacacs+
3. Configure the TACACS+ server.
Force10#show run tacacs+
!
tacacs-server key 7 387a7f2df5969da4
tacacs-server host 10.11.197.49
</pre></td></tr>
<tr><td>3</td><td>Apply the method list to the VTY lines.
<pre>
Force10#show run line
!
line console 0
line aux 0
line vty 0
 login authentication tacmethod
 authorization exec tacmethod
line vty 1
 login authentication tacmethod
 authorization exec tacmethod
line vty 2
 login authentication tacmethod
 authorization exec tacmethod
line vty 3
 login authentication tacmethod
 authorization exec tacmethod
line vty 4
 login authentication tacmethod
 authorization exec tacmethod
Force10#
</pre></td></tr>
</table>

Step	Task
4	Attempt an inbound Telnet session using IPv6 addresses.

```
freebsd2# telnet 2200:2200:2200:2200:2200::2202
Trying 2200:2200:2200:2200:2200::2202...
Connected to 2200:2200:2200:2200:2200::2202.
Escape character is '^]'.
Login: admin
Password:
Force10#
Force10#
!-The prompt is returned as the connection is authenticated.
```

Enabling and Disabling the SSH Daemon

By default, the SSH daemon is disabled. To enable the SSH daemon, you must use the command shown below, or enable it in the startup config.

Use the **no ip ssh server enable** command to disable the SSH daemon:

Enabling and Disabling the Telnet Daemon

By default, the Telnet daemon is enabled. To disable the Telnet daemon, you must use the command shown below or disable it in the startup config.

Use the **no ip telnet server enable** command to enable or disable the Telnet daemon:

Trace List

You can log packet activity on a port to confirm the source of traffic attacking a system. Once the Trace list is enabled on the system, you view its traffic log to confirm the source address of the attacking traffic. In FTOS, Trace lists are similar to extended IP ACLs, except that Trace lists are not applied to an interface. Instead, Trace lists are enabled for all switched traffic entering the E-Series.

The number of entries allowed per Trace list is 1K.

In the E-Series, you can create a trace filter based on any of the following criteria:

- Source IP address
- Destination IP address
- Source TCP port number
- Destination TCP port number
- Source UDP port number
- Destination UDP port number

For trace lists, you can match criteria on specific or ranges of TCP or UDP ports or established TCP sessions.

Note: If there are unresolved next-hops and a Trace-list is enabled, there is a possibility that the traffic hitting the CPU will not be rate-limited.

When creating an trace list, the sequence of the filters is important. You have a choice of assigning sequence numbers to the filters as you enter them, or FTOS assigns numbers in the order the filters were created. For more information on sequence numbering, refer to Chapter 17, IP Access Control Lists, Prefix Lists, and Route-maps, on page 301.

Configuration Task List for Trace lists

The following configuration steps include mandatory and optional steps.

- create a trace list on page 155 (mandatory)
- apply trace list on page 159 (mandatory)

For a complete listing of all commands related to Trace lists, refer to .

create a trace list

Trace lists filter and log traffic based on source and destination IP addresses, IP host addresses, TCP addresses, TCP host addresses, UDP addresses, and UDP host addresses. When configuring the Trace list filters, include the **count** and **bytes** parameters so that any hits to that filter are logged.

Since traffic passes through the filter in the order of the filter's sequence, you can configure the trace list by first entering the TRACE LIST mode and then assigning a sequence number to the filter.

To create a filter for packets with a specified sequence number, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip trace-list** *trace-list-name*	CONFIGURATION	Enter the TRACE LIST mode by creating an trace list.

Step	Command Syntax	Command Mode	Purpose
2	**seq** *sequence-number* {**deny** | **permit**} {**ip** | *ip-protocol-number*} {*source mask* | **any** | **host** *ip-address*} {*destination mask* | **any** | **host** *ip-address*} [**count** [**byte**] | **log**]	TRACE LIST	Configure a drop or forward filter. Configure the following required and optional parameters: • *sequence-number* range: 0 to, 4294967290. • **ip**: to specify IP as the protocol to filter for. • *ip-protocol-number* range: 0 to 255. • *source*: An IP address as the source IP address for the filter to match. • *mask:* a network mask • **any**: to match any IP source address • **host** *ip-address:* to match IP addresses in a host. • *destination*: An IP address as the source IP address for the filter to match. • **count:** count packets processed by the filter. • **byte:** count bytes processed by the filter. • **log:** is supported.

To create a filter for TCP packets with a specified sequence number, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip trace-list** *trace-list-name*	CONFIGURATION	Create a trace list and assign it a unique name.
2	**seq** *sequence-number* {**deny** | **permit**} **tcp** {*source mask* | **any** | **host** *ip-address*} [*operator port* [*port*]] {*destination mask* | **any** | **host** *ip-address*} [*operator port* [*port*]] [**established**] [**count** [**byte**] | **log**]	TRACE LIST	Configure a trace list filter for TCP packets. • *source*: An IP address as the source IP address for the filter to match. • *mask:* a network mask • **any**: to match any IP source address • **host** *ip-address:* to match IP addresses in a host. • *destination*: An IP address as the source IP address for the filter to match. • **count:** count packets processed by the filter. • **byte:** count bytes processed by the filter. • **log:** is supported.

To create a filter for UDP packets with a specified sequence number, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip trace-list** *access-list-name*	CONFIGURATION	Create a trace list and assign it a unique name.
2	**seq** *sequence-number* {**deny** | **permit**} **udp** {*source mask* | **any** | **host** *ip-address*} [*operator port* [*port*]] {*destination mask* | **any** | **host** *ip-address*} [*operator port* [*port*]] [**count** [**byte**] | **log**]	TRACE LIST	Configure a trace list filter for UDP packets. • *source*: An IP address as the source IP address for the filter to match. • *mask:* a network mask • **any**: to match any IP source address • **host** *ip-address:* to match IP addresses in a host. • *destination*: An IP address as the source IP address for the filter to match. • **count:** count packets processed by the filter. • **byte:** count bytes processed by the filter. • **log:** is supported.

When you create the filters with a specific sequence number, you can create the filters in any order and the filters are placed in the correct order.

Note: When assigning sequence numbers to filters, keep in mind that you might need to insert a new filter. To prevent reconfiguring multiple filters, assign sequence numbers in multiples of five or another number.

Figure 58 illustrates how the **seq** command orders the filters according to the sequence number assigned. In the example, filter 15 was configured before filter 5, but the **show config** command displays the filters in the correct order.

```
Force10(config-trace-acl)#seq 15 deny ip host 12.45.0.0 any log
Force10(config-trace-acl)#seq 5 permit tcp 121.1.3.45 0.0.255.255 any
Force10(config-trace-acl)#show conf
!
ip trace-list dilling
 seq 5 permit tcp 121.1.0.0 0.0.255.255 any
 seq 15 deny ip host 12.45.0.0 any log
Force10(config-trace-acl)#
```

Figure 58 Trace list Using seq Command Example

If you are creating a Trace list with only one or two filters, you can let FTOS assign a sequence number based on the order in which the filters are configured. FTOS assigns filters in multiples of 5.

To configure a filter for a Trace list without a specified sequence number, use any or all of the following commands in the TRACE LIST mode:

Command Syntax	Command Mode	Purpose
{**deny** | **permit**} {**ip** | *ip-protocol-number*} {*source mask* | **any** | **host** *ip-address*} {*destination mask* | **any** | **host** *ip-address*} [**count** [**byte**] | **log**]	TRACE LIST	Configure a deny or permit filter to examine IP packets. Configure the following required and optional parameters: • **ip**: to specify IP as the protocol to filter for. • *ip-protocol-number* range: 0 to 255. • *source*: An IP address as the source IP address for the filter to match. • *mask:* a network mask • **any**: to match any IP source address • **host** *ip-address:* to match IP addresses in a host. • *destination*: An IP address as the source IP address for the filter to match. • **count:** count packets processed by the filter. • **byte:** count bytes processed by the filter. • **log:** is supported.
{**deny** | **permit**} **tcp** {*source mask* | **any** | **host** *ip-address*} [*operator port* [*port*]] {*destination mask* | **any** | **host** *ip-address*} [*operator port* [*port*]] [**established**] [**count** [**byte**] | **log**]	TRACE LIST	Configure a deny or permit filter to examine TCP packets. Configure the following required and optional parameters: • *source*: An IP address as the source IP address for the filter to match. • *mask:* a network mask • **any**: to match any IP source address • **host** *ip-address:* to match IP addresses in a host. • *destination*: An IP address as the source IP address for the filter to match. • **precedence** *precedence* range: 0 to 7. • **tos** *tos-value* range: 0 to 15 • **count:** count packets processed by the filter. • **byte:** count bytes processed by the filter. • **log:** is supported.

Command Syntax	Command Mode	Purpose
{**deny** | **permit**} **udp** {*source mask* | **any** | **host** *ip-address*} [*operator port* [*port*]] {*destination mask* | **any** | **host** *ip-address*} [*operator port* [*port*]] | **log**]	TRACE LIST	Configure a deny or permit filter to examine UDP packets. Configure the following required and optional parameters: • *source*: An IP address as the source IP address for the filter to match. • *mask:* a network mask • **any**: to match any IP source address • **host** *ip-address:* to match IP addresses in a host. • *destination*: An IP address as the source IP address for the filter to match. • **precedence** *precedence* range: 0 to 7. • **tos** *tos-value* range: 0 to 15 • **count:** count packets processed by the filter. • **byte:** count bytes processed by the filter. • **log:** is supported.

Figure 59 illustrates a Trace list in which the sequence numbers were assigned by the software. The filters were assigned sequence numbers based on the order in which they were configured (for example, the first filter was given the lowest sequence number). The **show config** command in the TRACE LIST mode displays the two filters with the sequence numbers 5 and 10.

```
Force10(config-trace-acl)#deny tcp host 123.55.34.0 any
Force10(config-trace-acl)#permit udp 154.44.123.34 0.0.255.255 host 34.6.0.0
Force10(config-trace-acl)#show config
!
ip trace-list nimule
 seq 5 deny tcp host 123.55.34.0 any
 seq 10 permit udp 154.44.0.0 0.0.255.255 host 34.6.0.0
```

Figure 59 Trace list Example

To view all configured Trace lists and the number of packets processed through the Trace list, use the **show ip accounting trace-list** command (Figure 58) in the EXEC privilege mode.

apply trace list

After you create a Trace list, you must enable it. Without enabling the Trace list, no traffic is filtered.

You can enable one Trace list.

To enable a Trace list, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ip trace-group *trace-list-name*	CONFIGURATION	Enable a configured Trace list to filter traffic.

To remove a Trace list, use the **no ip trace-group** *trace-list-name* command syntax.

Once the Trace list is enabled, you can view its log with the **show ip accounting trace-list** *trace-list-name* [**linecard** *number*] command.

```
Force10#show ip accounting trace-list dilling
Trace List dilling on linecard 0
 seq 2 permit ip host 10.1.0.0 any count (0 packets)
 seq 5 deny ip any any
Force10#
```

Figure 60 show ip accounting trace-list Command Example

Protection Against TCP Tiny and Overlapping Fragment Attack

Tiny and overlapping fragment attack is a class of attack where configured ACL entries—denying TCP port-specific traffic—can be bypassed, and traffic can be sent to its destination although denied by ACL. RFC 1858 and 3128 proposes a countermeasure to the problem. This countermeasure is configured into the line cards and enabled by default.

Chapter 7 Layer 2

This chapter describes the FTOS Layer 2 features:

- VLAN Interfaces on page 161
- Native VLANs on page 168
- MAC Addressing and MAC Access Lists on page 169
- NIC Teaming on page 184
- Microsoft Clustering on page 185

For information on configuring both Layer 2 and Layer 3 ACLs on an interface, see Chapter 17, IP Access Control Lists, Prefix Lists, and Route-maps, on page 301.

VLAN Interfaces

Virtual LANs or VLANs are a logical broadcast domain or logical grouping of interfaces in a LAN in which all data received is kept locally and broadcast to all members of the group. When in Layer 2 mode VLANs move traffic at wire speed and can span multiple devices. FTOS supports up to 4093 port-based VLANs and 1 Default VLAN, as specified in IEEE 802.1Q.

VLANs provide the following benefits:

- Improved security because you can isolate groups of users into different VLANs
- Ability to create one VLAN across multiple devices

For complete information on VLANs, refer to IEEE Standard 802.1Q *Virtual Bridged Local Area Networks*.

Table 9 displays the defaults for VLANs in FTOS.

Table 9 VLAN Defaults on FTOS

Feature	Default
Spanning Tree group ID	All VLANs are part of Spanning Tree group 0
Mode	Layer 2 (no IP address is assigned)
Default VLAN ID	VLAN 1

This section covers the following:

Default VLAN

When interfaces are configured for Layer 2 mode, they are automatically placed in the Default VLAN as untagged interfaces. Only untagged interfaces can belong to the Default VLAN.

Figure 61 displays the outcome of placing an interface in Layer 2 mode. To configure an interface for Layer 2 mode, use the **switchport** command. In Step 1, the **switchport** command places the interface in Layer 2 mode.

In Step 2, you see that the **show vlan** command the in EXEC privilege mode indicates that the interface is now part of the Default VLAN (VLAN 1).

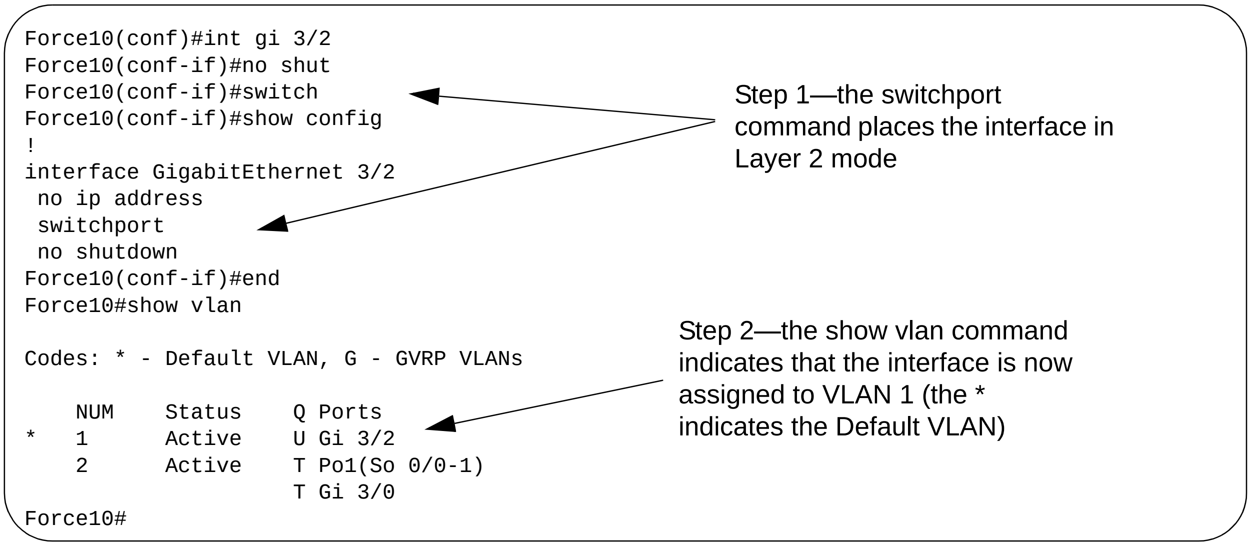

```
Force10(conf)#int gi 3/2
Force10(conf-if)#no shut
Force10(conf-if)#switch
Force10(conf-if)#show config
!
interface GigabitEthernet 3/2
 no ip address
 switchport
 no shutdown
Force10(conf-if)#end
Force10#show vlan

Codes: * - Default VLAN, G - GVRP VLANs

    NUM    Status    Q Ports
*   1      Active    U Gi 3/2
    2      Active    T Po1(So 0/0-1)
                     T Gi 3/0
Force10#
```

Step 1—the switchport command places the interface in Layer 2 mode

Step 2—the show vlan command indicates that the interface is now assigned to VLAN 1 (the * indicates the Default VLAN)

Figure 61 Interfaces and the Default VLAN Example

By default, VLAN 1 is the Default VLAN. To change that designation, use the **default vlan-id** command in the CONFIGURATION mode. You cannot delete the Default VLAN.

Untagged interfaces must be part of a VLAN, so to remove an interface from the Default VLAN, you must create another VLAN and place the interface into that VLAN. The alternative is to enter the **no switchport** command and the FTOS removes the interface from the Default VLAN.

Tagged interfaces require an additional step. Since tagged interfaces can belong to multiple VLANs, you must remove the tagged interface from all VLANs, using the **no tagged** *interface* command. Only after the interface is untagged and a member of the Default VLAN can you use the **no switchport** command to remove the interface from Layer 2 mode. For more information, see .

Port-Based VLANs

Port-based VLANs are a broadcast domain defined by different ports or interfaces. In the FTOS, a port-based VLAN can contain interfaces from different line cards within the chassis. The FTOS supports 4094 port-based VLANs.

Port-based VLANs offer increased security for traffic, conserve bandwidth, and allow switch segmentation. Interfaces in different VLANs do not communicate with each other, adding some security to the traffic on those interfaces. Different VLANs can communicate between each other by means of IP routing. Since traffic is only broadcast or flooded to the interfaces within a VLAN, the E-Series conserves bandwidth. Finally, you can have multiple VLANs configured on one switch, thus segmenting the device.

Interfaces within a port-based VLAN must be in Layer 2 mode and can be tagged or untagged in the VLAN ID.

VLANs and Port Tagging

To add an interface to a VLAN, it must be in Layer 2 mode. After you place an interface in Layer 2 mode, it is automatically placed in the Default VLAN. FTOS supports IEEE 802.1Q tagging at the interface level to filter traffic. When tagging is enabled, a Tag Header is added to the frame after the destination and source MAC addresses and that information is preserved as the frame moves through the network. Figure 62 illustrates the structure of a frame with a Tag Header. The VLAN ID is inserted in the Tag Header.

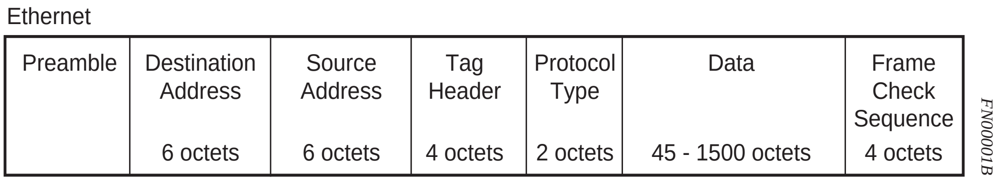
Ethernet

Preamble	Destination Address	Source Address	Tag Header	Protocol Type	Data	Frame Check Sequence
	6 octets	6 octets	4 octets	2 octets	45 - 1500 octets	4 octets

Figure 62 Tagged Frame Format

The Tag Header contains some key information used by FTOS:

- VLAN Protocol Identifier, which identifies the frame as tagged according to the IEEE 802.1Q specifications (2 bytes).
- Tag Control Information (TCI), which includes the VLAN ID (2 bytes total). The VLAN ID has a total of 4,096 values, but 2 are reserved.

Note: The insertion of the Tag Header into the Ethernet frame increases the size of the frame to more than the 1518 bytes specified in the IEEE 802.3 standard. Some devices that are not compliant with IEEE 802.3 may not support the larger frame size.

Information contained in the Tag Header allows the E-Series to prioritize traffic and to forward information to ports associated with a specific VLAN ID. Tagged interfaces can belong to multiple VLANs, while untagged interfaces can belong only to one VLAN.

Configuration Task List for VLANs

The following list includes the configuration tasks for VLANs:

- create a port-based VLAN on page 164 (mandatory)
- assign interfaces to a VLAN on page 165 (optional)
- assign an IP address to a VLAN on page 167 (optional)

For a complete listing of all commands related to VLANs, see .

create a port-based VLAN

The Default VLAN as VLAN 1 is part of the E-Series system startup configuration and does not require configuration. To configure a port-based VLAN, you must create the virtual interface and then add physical interfaces or port channel interfaces to the VLAN.

To create a port-based VLAN, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
interface vlan *vlan-id*	CONFIGURATION	Configure a port-based VLAN if the *vlan-id* is different from the Default VLAN ID. After you create a VLAN, you must assign interfaces in Layer 2 mode to the VLAN to activate the VLAN.

Use the **show vlan** command (Figure 36) in the EXEC privilege mode to view the configured VLANs.

```
Force10#show vlan

Codes: * - Default VLAN, G - GVRP VLANs

    NUM    Status    Q Ports
*   1      Inactive  U So 9/4-11
    2      Active    U Gi 0/1,18
    3      Active    U Gi 0/2,19
    4      Active    T Gi 0/3,20
    5      Active    U Po 1
    6      Active    U Gi 0/12
                     U So 9/0
Force10#
```

Figure 63   show vlan Command Example

A VLAN is active only if the VLAN contains interfaces and those interfaces are operationally up. In Figure 63, VLAN 1 is inactive because it contains the interfaces that are not active. The other VLANs listed in the Figure 63 contain enabled interfaces and are active.

➡ **Note:** In a VLAN, the **shutdown** command stops Layer 3 (routed) traffic only. Layer 2 traffic continues to pass through the VLAN. If the VLAN is not a routed VLAN (that is, configured with an IP address), the **shutdown** command has no affect on VLAN traffic.

When you delete a VLAN (using the **no interface vlan** *vlan-id* command), any interfaces assigned to that VLAN are assigned to the Default VLAN as untagged interfaces.

assign interfaces to a VLAN

Only interfaces in Layer 2 mode can be assigned to a VLAN using the **tagged** and **untagged** commands. Use the **switchport** command to place an interface in Layer 2 mode.

These Layer 2 interfaces can further be designated as tagged or untagged. For more information on interfaces in . When an interface is placed in Layer 2 mode by the **switchport** command, the interface is automatically designated untagged and placed in the Default VLAN.

To view which interfaces are tagged or untagged and to which VLAN they belong, use the **show vlan** command. For example, Figure 63 shows that six VLANs are configured on the E-Series, and two interfaces are assigned to VLAN 2. The Q column in the **show vlan** command example notes whether the interface is tagged (T) or untagged (U). For more information on this command, see .

To just view the interfaces in Layer 2 mode, enter the **show interfaces switchport** command in the EXEC privilege mode and EXEC mode.

To tag frames leaving an interface in Layer 2 mode, you must assign that interface to a port-based VLAN to tag it with that VLAN ID.

On the C-Series, egress mirrored packets are always tagged, even if the egress port is a member of an untagged VLAN.

To tag interfaces, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**interface vlan** *vlan-id*	CONFIGURATION	Configure a port-based VLAN if the *vlan-id* is different from the Default VLAN ID.
2	**tagged** *interface*	INTERFACE	Enable an interface to include the IEEE 802.1Q tag header.

Figure 64 illustrates the steps and commands to add a tagged interface (port channel 1) to VLAN 4.

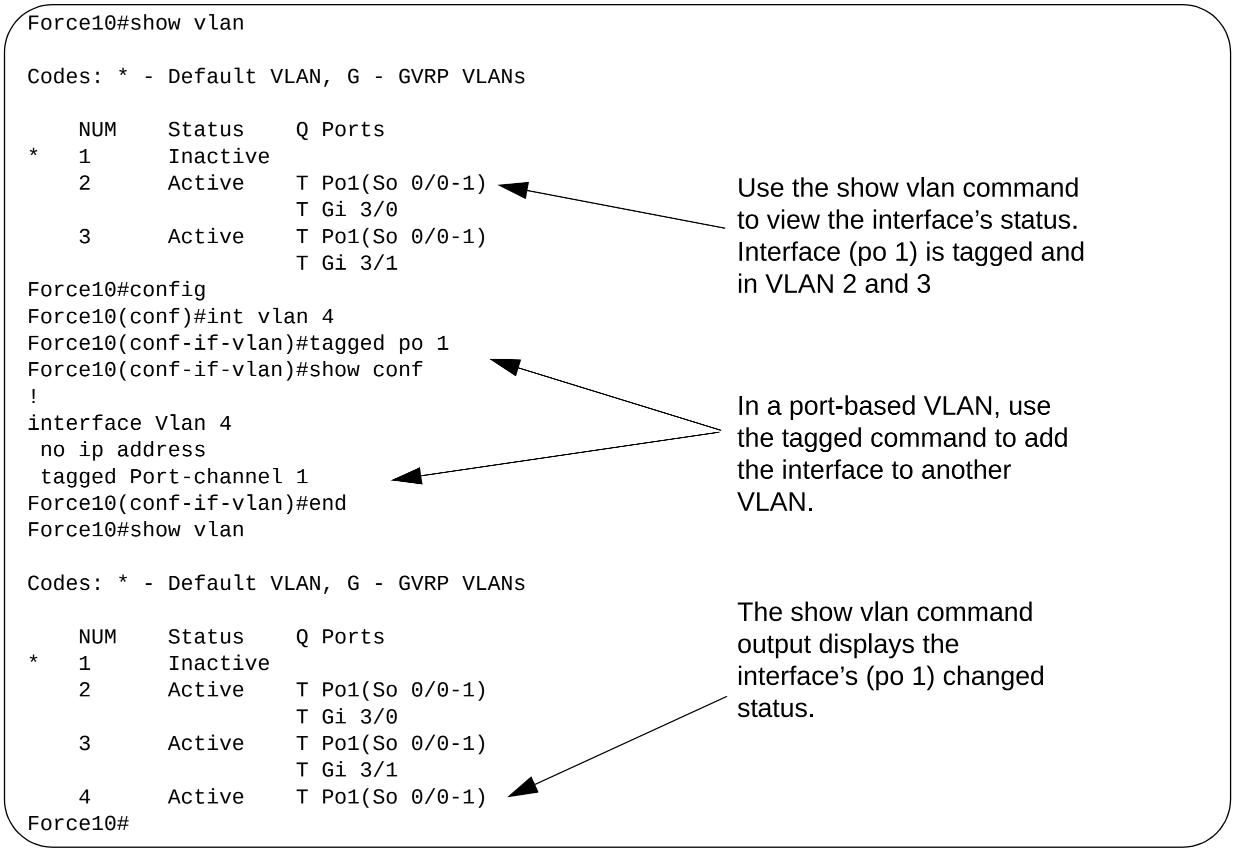

```
Force10#show vlan

Codes: * - Default VLAN, G - GVRP VLANs

    NUM    Status    Q Ports
*   1      Inactive
    2      Active    T Po1(So 0/0-1)
                     T Gi 3/0
    3      Active    T Po1(So 0/0-1)
                     T Gi 3/1
Force10#config
Force10(conf)#int vlan 4
Force10(conf-if-vlan)#tagged po 1
Force10(conf-if-vlan)#show conf
!
interface Vlan 4
 no ip address
 tagged Port-channel 1
Force10(conf-if-vlan)#end
Force10#show vlan

Codes: * - Default VLAN, G - GVRP VLANs

    NUM    Status    Q Ports
*   1      Inactive
    2      Active    T Po1(So 0/0-1)
                     T Gi 3/0
    3      Active    T Po1(So 0/0-1)
                     T Gi 3/1
    4      Active    T Po1(So 0/0-1)
Force10#
```

Use the show vlan command to view the interface's status. Interface (po 1) is tagged and in VLAN 2 and 3

In a port-based VLAN, use the tagged command to add the interface to another VLAN.

The show vlan command output displays the interface's (po 1) changed status.

Figure 64 Example of Adding an Interface to Another VLAN

Only a tagged interface can be a member of multiple VLANs. Hybrid ports are not supported, so the same interface cannot be assigned to two VLANs if the interface is untagged in one VLAN and tagged in the other VLAN.

When you remove a tagged interface from a VLAN (using the **no tagged** *interface* command), it will remain tagged only if it is a tagged interface in another VLAN. If the tagged interface is removed from the only VLAN to which it belongs, the interface is placed in the Default VLAN as an untagged interface.

With the **untagged** command you can move untagged interfaces from the Default VLAN to another VLAN.

To move untagged interfaces, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**interface vlan** *vlan-id*	CONFIGURATION	Configure a port-based VLAN if the *vlan-id* is different from the Default VLAN ID.

Step	Command Syntax	Command Mode	Purpose
2	**untagged** *interface*	INTERFACE	Configure an interface as untagged. This command is available only in VLAN interfaces.

The **no untagged** *interface* command removes the untagged interface from a port-based VLAN and places the interface in the Default VLAN. You cannot use the **no untagged** *interface* command in the Default VLAN. Figure 65 illustrates the steps and commands to move an untagged interface from the Default VLAN to another VLAN.

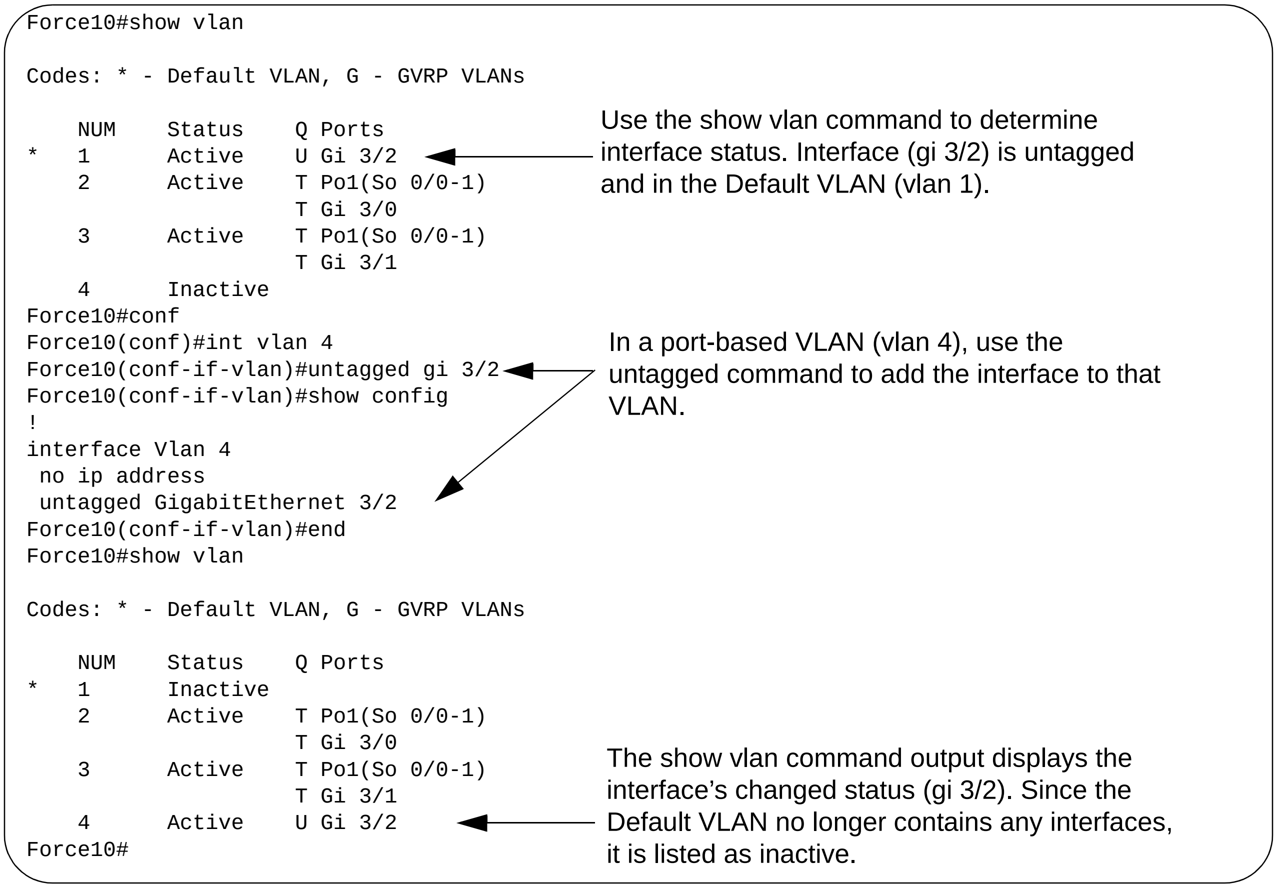

```
Force10#show vlan

Codes: * - Default VLAN, G - GVRP VLANs

    NUM    Status    Q Ports
*   1      Active    U Gi 3/2
    2      Active    T Po1(So 0/0-1)
                     T Gi 3/0
    3      Active    T Po1(So 0/0-1)
                     T Gi 3/1
    4      Inactive
Force10#conf
Force10(conf)#int vlan 4
Force10(conf-if-vlan)#untagged gi 3/2
Force10(conf-if-vlan)#show config
!
interface Vlan 4
 no ip address
 untagged GigabitEthernet 3/2
Force10(conf-if-vlan)#end
Force10#show vlan

Codes: * - Default VLAN, G - GVRP VLANs

    NUM    Status    Q Ports
*   1      Inactive
    2      Active    T Po1(So 0/0-1)
                     T Gi 3/0
    3      Active    T Po1(So 0/0-1)
                     T Gi 3/1
    4      Active    U Gi 3/2
Force10#
```

Use the show vlan command to determine interface status. Interface (gi 3/2) is untagged and in the Default VLAN (vlan 1).

In a port-based VLAN (vlan 4), use the untagged command to add the interface to that VLAN.

The show vlan command output displays the interface's changed status (gi 3/2). Since the Default VLAN no longer contains any interfaces, it is listed as inactive.

Figure 65 Example of Moving an Untagged Interface to Another VLAN

The only way to remove an interface from the Default VLAN is to place the interface in Default mode by entering the **no switchport** command in the INTERFACE mode.

assign an IP address to a VLAN

VLANs are a Layer 2 feature. For two physical interfaces on different VLANs to communicate, you must assign an IP address to the VLANs to route traffic between the two interfaces.

The **shutdown** command in INTERFACE mode does not affect Layer 2 traffic on the interface; the **shutdown** command only prevents Layer 3 traffic from traversing over the interface.

VLAN interfaces do not support SNMP, FTP or TFTP.

To assign an IP address, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip address *ip-address mask* [**secondary**]	INTERFACE	Configure an IP address and mask on the interface. • *ip-address mask:* enter an address in dotted-decimal format (A.B.C.D) and the mask must be in slash format (/24). • **secondary:** the IP address is the interface's backup IP address. You can configure up to eight secondary IP addresses.

In FTOS, VLANs and other logical interfaces can be placed in Layer 3 mode to receive and send routed traffic.

Native VLANs

C-Series	✓
E-Series	NO

Platform Specific Feature: Native VLANs is supported on C-Series only.

The physical ports can be either untagged for membership to one VLAN or tagged for membership to multiple VLANs. An untagged port must be connected to a VLAN unaware station (one that does not understand VLAN tags), and a tagged port must be connected to a VLAN aware station (one that generates and understands VLAN tags). Native VLAN support breaks this barrier so that a port can be connected to both VLAN aware and VLAN unaware stations.

Native VLAN is useful in deployments where a Layer 2 port can receive both tagged and untagged traffic on the same physical port. The classical example is connecting a VOIP phone and a PC on to the same port of the switch. The VOIP phone is configured to generate tagged packets (with VLAN = VOICE VLAN), and the PC attached generates untagged packets.

To configure a port so that it can be a member of an untagged and tagged VLANs, use the command **portmode hybrid** from INTERFACE mode. The port must have no other Layer 2 or Layer 3 configurations when entering this command or a message like Message 3 is displayed.

Message 3 Native VLAN Error

```
% Error: Port is in Layer-2 mode Gi 5/6.
```

MAC Addressing and MAC Access Lists

Media Access Control (MAC) is a sublayer of the data link layer (Layer 2 of the OSI seven-layer model). MAC addresses (machine addresses) are used to interconnect LAN components and dictate how each device accesses and shares the network connection. MAC addresses are displayed in a hexadecimal format. Figure 66 displays the format used for MAC addresses in the E-Series.

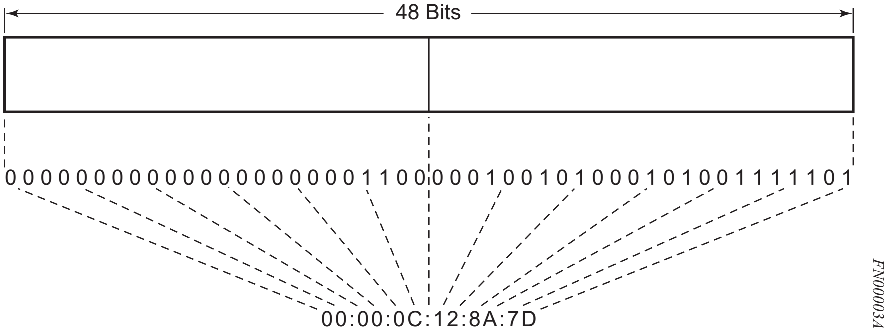

Figure 66 MAC Address Format

MAC addresses are used in Access Control Lists (ACLs) to prevent flooding of multicast traffic and to filter traffic. In the E-Series, you create an ACL to drop or forward traffic from MAC destination or source addresses, and you can filter traffic based on the Ethernet frame format used by the traffic. As soon as you configure the **mac access-list** command on an interface, it is applied to that interface and filters traffic on that interface.

For more information on MAC addresses, refer to IEEE Standard 802.1D *Media Access Control (MAC) Bridges*.

This section covers the following:

- MAC Access Control List Basics on page 169
- MAC ACL Implementation on page 170
- Configuration Task List for MAC ACLs on page 171
- Managing the MAC Address Table on page 178

MAC Access Control List Basics

An ACL is a series of sequential filters that contain a matching criterion (the MAC address) and an action (deny or permit). The filters are processed in sequence; for example, if the traffic does not match the first filter, the second filter is applied. When the MAC address matches a filter, FTOS drops or forwards the traffic based on the filter’s designated action. If the MAC address does not match any of the filters in the ACL, the traffic is forwarded. This default behavior is different from IP ACL, which drops traffic not matching any filters.

MAC ACL Implementation

The maximum size of MAC ACLs is determined by the CAM size of the line card and the Layer 2 CAM allocation between MAC addresses and MAC ACLs. Once you determine the maximum possible for your line card, you must also determine the CAM's allocation of MAC addresses versus MAC ACLs.

In E-Series, you can assign multiple ingress ACLs per interface. For TeraScale line cards you can also assign one egress ACL per interface. If an ACL is not assigned to an interface, it is not used by the software in any other capacity.

In FTOS, you can create two different types of MAC ACLs: standard or extended. A standard MAC ACL filters traffic based on the source MAC address. An extended MAC ACL filters traffic based on any of the following criteria:

- Source MAC address
- Destination MAC address
- Source MAC host address
- Destination MAC host address
- Ethernet frame type of the traffic

Both standard and extended ACLs allow you to filter traffic with any MAC address. Your first decision in configuring MAC access control lists is deciding whether the ACL will filter based solely on the MAC source address or based on additional factors.

The well-known MAC addresses (also known as protocol addresses) 0180c2000000 through 0180c200000f are always permitted, even if you configure a MAC ACL **deny** filter for these addresses. This default prevents Spanning-tree loops when the **mac learning-limit** command is configure on Spanning-tree enabled ports.

Note: (For EF cards only.) When ACL logging and byte counter are enabled simultaneously, the byte counter may show the wrong value. Instead, enable packet counter with logging.
Note: MAC accounting accounts for packets denied by an L2 access list when mirroring is configured.

The following are addition facts about MAC addresses:

- Each system is pre-assigned a block of MAC addresses that are stored in the backplane EEPROM.
- EtherScale platforms have 1k pre-allocated for MAC addresses.
- TeraScale E1200 and E600 platforms pre-allocate 2k for MAC addresses, the E300 pre-allocates 1.5k.
- Port/VLAN MAC addresses do not change after the system reboot.
- The MAC address on the management port of the RPM is not part of the system MAC address allocation pool.

MAC ACLs are supported over VLAN interfaces on TeraScale systems.

Configuration Task List for MAC ACLs

The following list includes the configuration tasks for MAC ACLs and MAC Addressing:

- configure standard MAC access control list on page 171 (mandatory)
- configure extended MAC access control list on page 173 (mandatory)
- assign a MAC ACL to an interface on page 176 (mandatory)
- specify CAM portion for MAC ACLs on page 177 (optional)

For a complete listing of all commands related to MAC addresses and MAC ACLs, refer to

configure standard MAC access control list

Standard MAC ACLs filter traffic based on the source MAC address. Since traffic passes through the ACL in the order of the filters' sequence, you can configure the MAC ACL by first entering the MAC ACCESS LIST mode and then assigning a sequence number to the filter.

To create a filter with a specified sequence number, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**mac access-list standard** *access-list-name*	CONFIGURATION	Enter the MAC ACCESS LIST mode by creating a standard MAC ACL.
2	**seq** *sequence-number* {**deny** | **permit**} {**any** | *source-mac-address*} [**count** [**byte**]] | [**log**]	MAC ACCESS LIST	Configure a MAC ACL filter with a specific sequence number. The **any** keyword filters on any source MAC address. **log** is not supported on C-Series.

When you create the filters with specific sequence numbers, you can create the filters in any order and FTOS orders the filters correctly.

Note: Keep in mind when assigning sequence numbers to filters that you may need to insert a new filter. To prevent reconfiguring multiple filters, assign sequence numbers in multiples of five or another number.

Figure 67 illustrates how the **seq** command orders the filters according to the sequence number assigned. In the example, filter 5 was configured before filter 2, but the **show config** command displays the filters in the correct order.

```
Force10(conf)#mac access-list standard stringbean
Force10(config-std-macl)#seq 5 deny 00:00:00:00:11:22
Force10(config-std-macl)#seq 2 permit any
Force10(config-std-macl)#show config
!
mac access-list standard stringbean
 seq 2 permit any
 seq 5 deny 00:00:00:00:11:22
Force10(config-std-macl)#
```

Figure 67 seq Command Example

To delete a filter, use the **no seq** *sequence-number* command in the MAC ACCESS LIST mode.

If you are creating a standard ACL with only one or two filters, you can let the E-Series software assign a sequence number based on the order in which the filters are configured. The E-Series software assigns filters in multiples of 5.

To configure a filter without a specified sequence number, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**mac access-list standard** *access-list-name*	CONFIGURATION	Create a standard MAC ACL and assign it a unique name.
2	{**deny** | **permit**} {**any** | *source-mac-address mask*} [**count** [**byte**]] [**log**]	MAC ACCESS LIST	Configure a MAC ACL filter. The **any** keyword filters on any source MAC address. **log** is not supported on C-Series.

Figure 68 illustrates a standard MAC ACL in which the sequence numbers were assigned by the E-Series software. The filters were assigned sequence numbers based on the order in which they were configured (for example, the first filter was given the lowest sequence number). The **show config** command in the MAC ACCESS LIST mode displays the two filters with the sequence numbers 5 and 10.

```
Force10(conf)#mac access standard belmont
Force10(config-std-macl)#permit 00:00:00:11:32:00
Force10(config-std-macl)#permit any
Force10(config-std-macl)#show config
!
mac access-list standard belmont
 seq 5 permit 00:00:00:11:32:00
 seq 10 permit any
Force10(config-std-macl)#
```

Figure 68 Standard MAC ACL Example

To view a specific configured MAC ACLs, use the **show mac accounting access-list** *access-list-name* command (Figure 69) in the EXEC privilege mode.

```
Force10#show mac accounting access-list belmont interface gigabitethernet 0/1 in
Standard mac access-list belmont on GigabitEthernet 0/1
 seq 5  permit 00:00:00:11:32:00
 seq 10  permit any
Force10#
```

Figure 69 show mac accounting access-list Command Example

To delete a filter, enter the **show config** in the MAC ACCESS LIST mode and locate the sequence number of the filter you want to delete; then use the **no seq** *sequence-number* command in the MAC ACCESS LIST mode.

configure extended MAC access control list

Extended MAC ACLs filter on source and destination MAC addresses. In addition, you have the option of filtering traffic based on the Ethernet frame structure. The E-Series software offers the option to filter traffic based on one of three Ethernet frame formats.

Table 10 lists the three formats to filter, the keywords used in the CLI, and a description.

Table 10 Three Ethernet Formats

Format	Keyword	Description
IEEE 802.3	llc	The frame format complies with IEEE Standard 802.3 and contains both a Data Link Header and an LLC header.
Ethernet II	ev2	The frame format complies with the original Ethernet II specification, and the Data Link Header contains 14 bytes of information. This format type does not contain an LLC header.
IEEE 802.3 SNAP	snap	The frame format complies with the IEEE Standard 802.3 SNAP (SubNetwork Access Protocol) specification. This format contains both a Data Link Header and an LLC header, in addition to a SNAP field (5 bytes).

Since traffic passes through the filter in order of the filter's sequence, you can configure the MAC ACL by first entering the MAC ACCESS LIST mode and then assigning a sequence number to the filter.

Table 11 Layer 2 ACL Supported Features

	EtherScale		TeraScale		C-Series	
Feature	Ingress	Egress	Ingress	Egress	Ingress	Egress
VLAN	Per-VLAN/ Range of VLANs	Not supported	Per-VLAN/Range of VLANs	Is applied to all VLANs		Not supported

Table 11 Layer 2 ACL Supported Features

	EtherScale		TeraScale		C-Series	
Logging	No supported	Not supported	Supported	Supported		Not supported
EV2/ SNAP/ LLC[-1]	Configurable	No supported	Not configurable [-1]	Not configurable[-1]		Not supported

Note: -1: When a user configures an ACL for an ethertype like “8137” for an encapsulation type like “SNAP”, the ACL applies to all encapsulation types (EV2, SNAP, LLC) for Terascale.

To create a filter with a specified sequence number, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**mac access-list extended** *access-list-name*	CONFIGURATION	Create a extended MAC ACL and assign it a unique name.
2	**seq** *sequence-number* {**deny** | **permit**} {**any** | **host** *mac-address* | *mac-source-address mac-source-address-mask*} {**any** | **host** *mac-address* | *mac-destination-address mac-destination-address-mask*} [*ethertype-operator*] [**count** [**byte**]] [**log**]	MAC ACCESS LIST	Configure a MAC ACL filter. The **any** keyword filters on any source MAC address. The **host** keyword followed by a MAC address filters all MAC addresses with that host. The optional *ethertype-operator* values are discussed in Table 10. **log** not supported on EtherScale line cards or C-Series.

When you create the filters with specific sequence numbers, you can create the filters in any order and FTOS orders the filters correctly.

Note: Keep in mind that when assigning sequence numbers to filters you might need to insert a new filter. To prevent reconfiguring multiple filters, assign sequence numbers in multiples of five or another number.

Figure 70 illustrates how the **seq** command orders the filters according to the sequence number assigned. In the example, filter 15 was configured before filter 5, but the **show config** command displays the filters in the correct order.

```
Force10(conf)#mac access-list extended dunedin
Force10(config-ext-macl)#seq 15 deny 00:00:00:11:ed:00 ff:ff:ff:ff:ff:ff 00:00:00:ab:11:00
ff:ff:ff:ff:ff:ff
Force10(config-ext-macl)#seq 5 permit host 00:00:00:00:45:ef any
Force10(config-ext-macl)#show config
!
mac access-list extended dunedin
 seq 5 permit host  00:00:00:00:45:ef any
 seq 15 deny 00:00:00:00:ec:00 ff:ff:ff:ff:ff:ff 00:00:00:aa:00:00 ff:ff:ff:ff:ff:ff
Force10(config-ext-macl)#
```

Figure 70 Extended MAC ACL Using the seq Command Example

If you are creating a standard ACL with only one or two filters, you can let the E-Series software assign a sequence number based on the order in which the filters are configured. The E-Series software assigns filters in multiples of 5.

To configure a filter without a specified sequence number, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**mac access-list extended** *access-list-name*	CONFIGURATION	Create a extended MAC ACL and assign it a unique name.
2	{**deny** | **permit**} {**any** | **host** *mac-address* | *mac-source-address mac-source-address-mask*} {**any** | **host** *mac-address* | *mac-destination-address mac-destination-address-mask*} [*ethertype-operator*] [**count** [**byte**]] [**log**]	MAC ACCESS LIST	Configure a MAC ACL filter with a specific sequence number. The **any** keyword filters on any source MAC address. The **host** keyword followed by a MAC address filters all MAC addresses with that host. The optional *ethertype-operator* values are discussed in Table 10 on page 107. **log** is not supported on C-Series.

Figure 71 illustrates an extended MAC ACL in which the sequence numbers were assigned by FTOS. The filters were assigned sequence numbers based on the order in which they were configured (for example, the first filter was given the lowest sequence number). The **show config** command in the MAC ACCESS LIST mode displays the two filters with the sequence numbers 5 and 10.

```
Force10(conf)#mac access-list extended auckland
Force10(config-ext-macl)#permit 00:00:00:00:22:ee ff:ff:ff:ff:ff:ff any
Force10(config-ext-macl)#deny host 22:00:00:11:ab:ef 00:00:00:ce:00:00 ff:ff:ff:ff:ff:ff
Force10(config-ext-macl)#show config
!
mac access-list extended auckland
 seq 5 permit 00:00:00:00:22:ee ff:ff:ff:ff:ff:ff any
 seq 10 deny host  22:00:00:11:ab:ef 00:00:00:ce:00:00 ff:ff:ff:ff:ff:ff
Force10(config-ext-macl)#
```

Figure 71 Extended MAC ACL Example

To view all configured MAC ACLs, use the **show mac accounting access-list** [*access-list-name*] **interface** [*interface-name*] **in/out** command in the EXEC mode.

assign a MAC ACL to an interface

To pass traffic through a configured MAC ACL, you must assign that ACL to a Layer 2 interface. The MAC ACL is applied to all traffic entering the Layer 2 interface and the traffic is either forwarded or dropped depending on the criteria and actions specified in the ACL.

For more information on Layer 2 interfaces, see Chapter 9, Interfaces, on page 197.

To apply a MAC ACL (standard or extended) to a physical or port channel interface, use these commands in the following sequence in the INTERFACE mode:

Step	Command Syntax	Command Mode	Purpose
1	**switchport**	INTERFACE	Place the interface in Layer 2 mode.
2	**mac access-group** *access-list-name* {**in** [**vlan** *vlan-range*] | **out**}	INTERFACE	Purpose: Apply a MAC ACL to traffic entering or exiting an interface. • **in**: configure the ACL to filter incoming traffic • **out**: configure the ACL to filter outgoing traffic. Available only with E-Series TeraScale cards. • **vlan** *vlan-range*: (OPTIONAL) specify a range of VLANs.

To view which MAC ACL is applied to an interface, use the **show config** command (Figure 52) in the INTERFACE mode or the **show running-config** command in the EXEC mode.

```
Force10(conf-if-gi-0/4)#show config
!
interface GigabitEthernet 0/4
 no ip address
 switchport
 mac access-group dunedin out
 no shutdown
Force10(conf-if-gi-0/4)#
```

Figure 72 show config Command in the INTERFACE Mode

specify CAM portion for MAC ACLs

C-Series	NO
E-Series	✓

Platform Specific Feature: specify CAM portion for MAC ACLs is supported on E-Series only.

For EtherScale lline cards, you can change the allocation in the Layer 2 CAM between MAC addresses and MAC ACLs. By default, the 75% of the Layer 2 CAM is reserved for MAC addresses and the remaining 25% is reserved for MAC ACLs.

To reallocate the Layer 2 CAM for MAC ACLs, use the following command:

Command Syntax	Command Mode	Purpose
mac cam fib-partition {**25** | **50** | **75** | **100**} *slot-number*	CONFIGURATION	Reapportion the Layer 2 CAM space for MAC addresses. • *slot-number* range: 0 to 13 for E1200; 0 to 6 for E600; 0 to 5 for E300

After you enter this command, the user is prompted with the following message:

```
Line card should be reset for new CAM entries to take effect.
Proceed with reset? [yes/no]:
```

Figure 73 Prompt After issuing the mac cam fib-partition Command

To view the MAC CAM allocation on all line cards, use the **show mac cam** command (Figure 74).

```
Force10#show mac cam
Slot  Type     MAC CAM Size   MAC FIB Entries   MAC ACL Entries
  0   EX2YD     64K entries       48K (75%)          8K (25%)
  9   F12PC     32K entries       24K (75%)          4K (25%)
 12   F12PD     64K entries       48K (75%)          8K (25%)
 13   E24PD     64K entries       48K (75%)          8K (25%)
Note: All CAM entries are per portpipe.
Force10#
```

Figure 74 show mac cam Command Example

Managing the MAC Address Table

The primary commands for this activity are:

- **clear mac-address-table**
- **mac-address-table aging-time**
- **mac-address-table static**
- **mac-address-table station-move**
- **mac learning-limit**
- **show mac-address-table**
- **show mac learning-limit**

For the complete command set for managing MAC addresses, see the “MAC Addressing” section in the Layer 2 chapter of the *FTOS Command Line Interface Reference*.

The MAC Address Table is controlled by:

- **An aging timer**: A time is applied to a dynamic address in the table. A dynamic address is learned by an interface. If no packet from or to the address arrives on the switch within the timer period, the address is removed from the table.

  The default timer is 1800 seconds. Use the **mac-address-table aging-time** command to modify the interval.
- **Manually-added and removed addresses**: A MAC address that you manually add to the table is static and is therefore not affected by the timer. See the section Configure static MAC addresses (optional) on page 179 to manually add a MAC address.
- **MAC Address Learning Limit feature**: You can set a limit on the number of addresses associated with a particular interface. As a default, after you set a limit, a MAC address learned by the interface is added to the MAC Address Table as a static address. You have other options in this feature, for setting address limits and actions involved in violations of the learning limit.

Configuration Task List for MAC Addressing

The following list includes configuration tasks for MAC addressing, all optional:

- Configure static MAC addresses (optional) on page 179
- Configure a MAC address learning limit for an interface (optional) on page 180

- Verify MAC learning-limit settings on page 181
- Configure MAC learning-limit violation actions (optional) on page 181
- Verify MAC learning-limit violation settings on page 182
- Recovering from MAC address learning-limit violations on page 183

Configure static MAC addresses (optional)

Occasionally you want to statically configure some MAC addresses for devices that always remain attached to the E-Series. When a static MAC address is configured and its interface is disabled, the packets destined to this MAC address are not flooded.

To configure static MAC addresses, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
mac-address-table static *mac-address* **output** *interface* **vlan** *vlan-id*	CONFIGURATION	Assign a static MAC address to an interface and a VLAN.

To view the static MAC address and the dynamically-learnt MAC addresses, use the **show mac-address-table static** command in the EXEC mode.

```
Force10#show mac-address static
VlanId     Mac Address           Type   Interface        State
 1      00:00:00:00:11:22       Static  Po 3            Inactive
Force10#
```

Figure 75 show mac-address-table static Command Example

To clear dynamically-learnt MAC addresses, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Purpose
clear mac-address-table dynamic {**address** *mac-address* | **all** | **interface** *interface* | **vlan** *vlan-id*}	EXEC privilege	Clear only dynamically-learnt MAC addresses. Configure one of the following parameters: • **address** *mac-address*: dynamically-learnt MAC address • **all**: all dynamically-learnt MAC addresses • **interface** *interface*: specify an interface. • **vlan** *vlan-id*: enter a VLAN ID.

Configure a MAC address learning limit for an interface (optional)

The MAC address learning limit feature enables you to set an upper limit on the number of MAC addresses that can be entered for a particular interface in the MAC Address Table. After you set a learning limit, and the limit is reached, any MAC addresses encountered by the interface are ignored and the traffic is dropped. To set a limit on the number of MAC addresses that can access a particular port, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
mac learning-limit *1-1000000* [**dynamic**][**no-station-move**]	INTERFACE	• *1-1000000*: Enter a number between 1 and 1 million as the maximum number of MAC addresses that can be entered in the MAC Address Table for the selected interface. • **dynamic** (OPTIONAL): MAC addresses learned for this interface enter the MAC Address Table as dynamic addresses when the learning limit is set. • **no-station-move** (OPTIONAL): Stop addresses learned on this interface from being moved to another interface in the MAC Address Table.

As described in Managing the MAC Address Table on page 178, dynamic addresses are subject to aging out of the table.

The "sticky MAC" option: MAC addresses learned under a learning limit will enter the MAC Address Table either as static or dynamic addresses, depending on whether the **dynamic** option is set, as described above. In either case, however, if the address is then learned on another interface, the address is moved to that other interface in the MAC Address Table. This can cause unnecessary MAC shuffling within the system, causing unnecessary flooding. The **no-station-move** option, also commonly called a "sticky MAC" option, prevents that from happening by causing the first entry to persist in the table even if the address is received on other interface.

Notes:

- For both the **dynamic** and **no-station-move** options, addresses entered before those options are set are not affected.
- MAC learning should be on a per-VLAN basis.
- Egress ACLs can be applied to interfaces that have those features configured on ingress. Only the MAC limit-learned addresses should be permitted on the ingress interfaces. ACLs can be applied on the egress (typically to stop unknown protocol traffic). See MAC ACL Implementation on page 170.

The **dynamic** and **no-station-move** options do not overlap. You can configure either or both, as shown in Figure 76 below. When the original interface gets unplugged, the **no-station-move** setting does not prevent a reassignment of the MAC address. In other words, the table clears all address learnt on the interface (physical port). So, even though **no-station-move** is configured on the unplugged interface, all those MAC addresses can be learned on any other interface (physical port).

This example shows four ways to specify a MAC address learning limit of 10 on port 16:

```
Force10(conf-if-gi-1/16)#mac learning-limit 10
Force10(conf-if-gi-1/16)#mac learning-limit 10 dynamic
Force10(conf-if-gi-1/16)#mac learning-limit 10 no-station-move
Force10(conf-if-gi-1/16)#mac learning-limit 10 dynamic no-station-move
```

Figure 76 mac learning-limit Command Example

Verify MAC learning-limit settings

To display MAC address learning limits, use the **show mac learning-limit** [**interface** *interface* | **detail**] [**detail**]command in the EXEC mode.

```
Force10# show mac learning-limit detail
Interface    Vlan   Station-Move    Type
Gi 0/0          10        Enable  Static
Gi 0/1          10        Disable Dynamic
```

Figure 77 show mac learning-limit Command Example

Configure MAC learning-limit violation actions (optional)

For the condition of when the MAC address learning limit is reached on a specific interface and a new address is received, you can select a violation action with the following command in INTERFACE mode:

Command Syntax	Command Mode	Purpose
mac learning-limit learn-limit-violation {**log** | **shutdown**}	INTERFACE	Enter one of the following keywords: • **log**: A message is entered in the syslog. • **shutdown**: Shut down the selected interface (port or LAG) and enter a message in the syslog.

You can select an action for the switch to take for the case in which a MAC address that is learned under one MAC learning limit is then learned on another interface. You can use the following command whether or not you have set the **no-station-move** option, but the violation action will only occur if:

1. You have used the following command to select a violation action.
2. You have selected the **no-station-move** option for the interface.

3. And a station-move occurs on that interface.

Command Syntax	Command Mode	Purpose
mac learning-limit station-move-violation {**log** | **shutdown-original** | **shutdown-offending** | **shutdown-both**}	INTERFACE	Enter one of the following keywords: • **log**: This option directs the switch to enter a message in the syslog. • **shutdown-original**: The first interface to learn the MAC address is shut down (port or LAG). • **shutdown-offending**: The second interface to learn the address is shut down. • **shutdown-both**: Both interfaces that learn the address are shut down.

Note: For all learning-limit violation actions, at least one syslog message is generated, for example:

```
Force10(conf-if-gi-3/8)#2d21h51m: %RPM0-P:RP2 %MACMGR-1-LEARN LIMIT VIOLATION: Learn limit
violation occurred on Gi 3/8: vlan-1: mac-00:00:03:cc:cc:0a

2d21h51m: %RPM0-P:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Gi 3/8
```

Figure 78 mac learning-limit violation log messages Example

Verify MAC learning-limit violation settings

Use the **show mac learning-limit violate-action** command to display settings for learning-limit violation actions.

```
Force10(conf-if-gi-3/8)# do show mac learning-limit violate-action
Interface      Violation-Type Violate-Action      Status
Gi 0/0        station-move  log                 Normal
Gi 3/6         station-move shutdown-original  Normal
Gi 3/7        station-move  Shutdown-Both      Error-disable
Gi 3/8         learn-limit    Shutdown            Normal
```

Figure 79 show mac learning-limit violation Command Example

Recovering from MAC address learning-limit violations

After a learning-limit violation shuts down an interface, you must manually reset it. Use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
mac learning-limit reset [**learn-limit-violation** {**interface** *interface* | **all**} | **station-move-violation** {**interface** *interface* | **all**}]	CONFIGURATION	Optionally, enter either one of the following keywords: • **learn-limit-violation**: Reset ERR_Disabled state caused by a learn-limit violation. • **station-move-violation**: Reset ERR_Disabled state caused by a station-move violation. After entering one of the first two options, enter one of the following keywords: • **interface** *interface*: Enter the **interface** keyword and the interface ID (port number or LAG) to designate a specific interface to reenable. • **all**: Reenable all interfaces shut down by a learning-limit violation.

Note: Alternatively, you can use the standard commands **shutdown** and then **no shutdown** to reenable a specific interface, but a **no shutdown all** command also administratively enables interfaces that are shut down whether or not they were shut down by a learning-limit violation.

```
Force10(conf-if-gi-3/8)#do mac learning-limit reset
Force10(conf-if-gi-3/8)#2d21h50m: %RPM0-P:CP %IFMGR-5-OSTATE_UP: Changed interface state to
up: Gi 3/8
2d21h50m: %RPM0-P:CP %IFMGR-5-ACTIVE: Changed Vlan interface state to active: Vl 1
```

Figure 80 show mac learning-limit violation Command Example

Note: In this MAC Learning Limit section, above, “interface” refers to either ports or port channels (static LAG or LACP LAG).

NIC Teaming

NIC teaming is a feature that allows multiple network interface cards in a server to be represented by one MAC address and one IP Address in order to provide transparent redundancy, balancing, and to fully utilize network adapter resources.

Figure 81 shows a topology where two NICs have been teamed together. In this case, if the primary NIC fails, traffic is automatically switched to the secondary NIC, since they are represented by the same set of addresses.

Figure 81 Redundant NICs with NIC Teaming

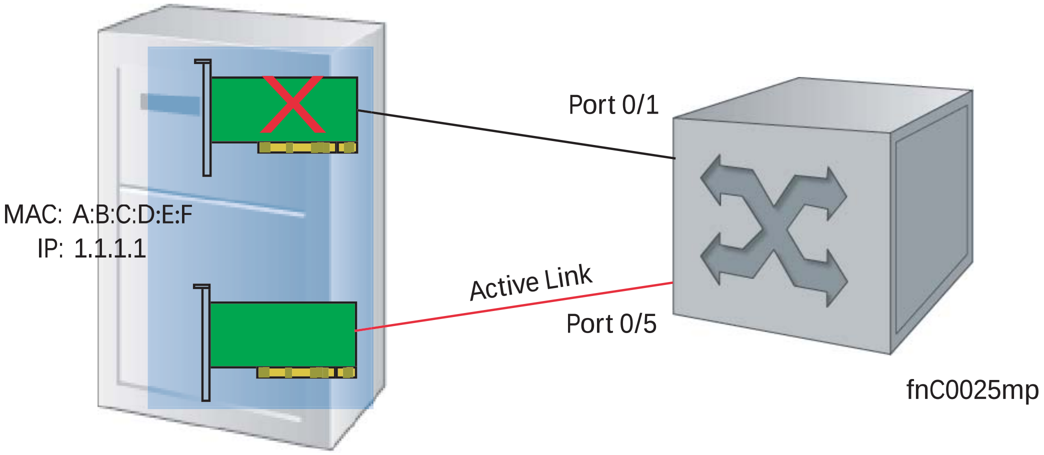

When NIC teaming is employed, consider that the server MAC address is originally learned on Port 0/1 of the switch (Figure 82). When the NIC fails, the same MAC address is learned on Port 0/5 of the switch. The MAC address must be disassociated with the one port and re-associated with another in the ARP table; in other words, the ARP entry must be “moved.” To ensure that this happens, you must configure the command **mac-address-table station-move refresh-arp** on the Force10 switch at the time that NIC teaming is being configured on the server.

Note: If this command is not configured, traffic will continue to be forwarded to the failed NIC until the ARP entry on the switch times-out.

Figure 82 Configuring mac-address-table station-move refresh-arp Command

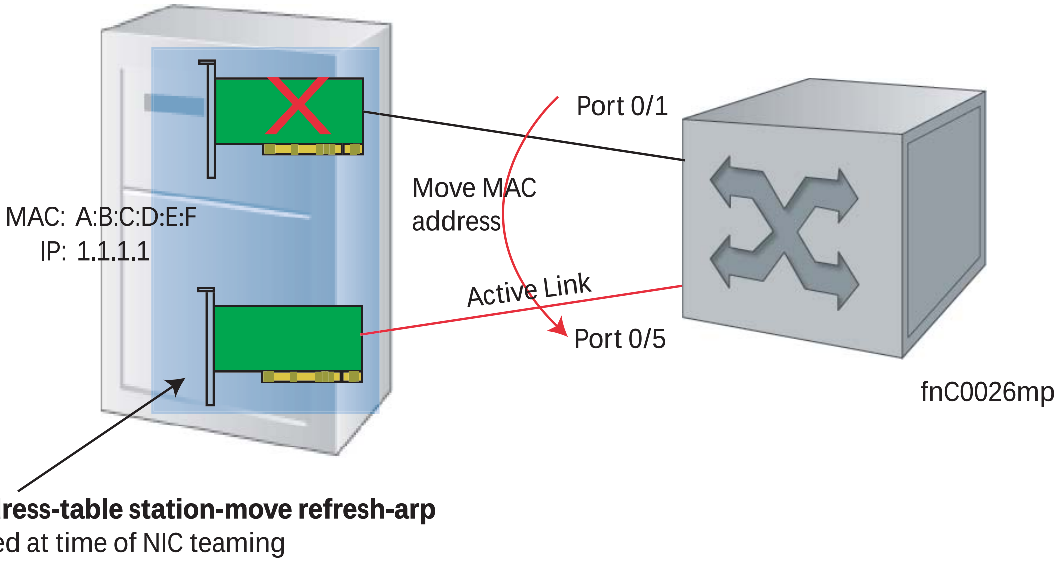
Port 0/1
MAC: A:B:C:D:E:F
IP: 1.1.1.1
Move MAC address
Active Link
Port 0/5
fnC0026mp
mac-address-table station-move refresh-arp
configured at time of NIC teaming

Microsoft Clustering

Microsoft Clustering is a feature that allows multiple servers using Microsoft Windows to be represented by one MAC address and IP address in order to provide transparent failover or balancing. FTOS does not recognize server clusters by default; it must be configured to do so.

Default Behavior

When an ARP request is sent to a server cluster, either the active server or all of the servers send a reply, depending on the cluster configuration. If the active server sends a reply, the Force10 switch learns the active server’s MAC address. If all servers reply, the switch registers only the last received ARP reply, and again the switch learns one server’s actual MAC address (Figure 83); the virtual MAC address is never learned.

Since the virtual MAC address is never learned, traffic is forwarded to only one server rather than the entire cluster, and failover and balancing are not preserved (Figure 84).

Figure 83 Server Clustering: Multiple ARP Replies

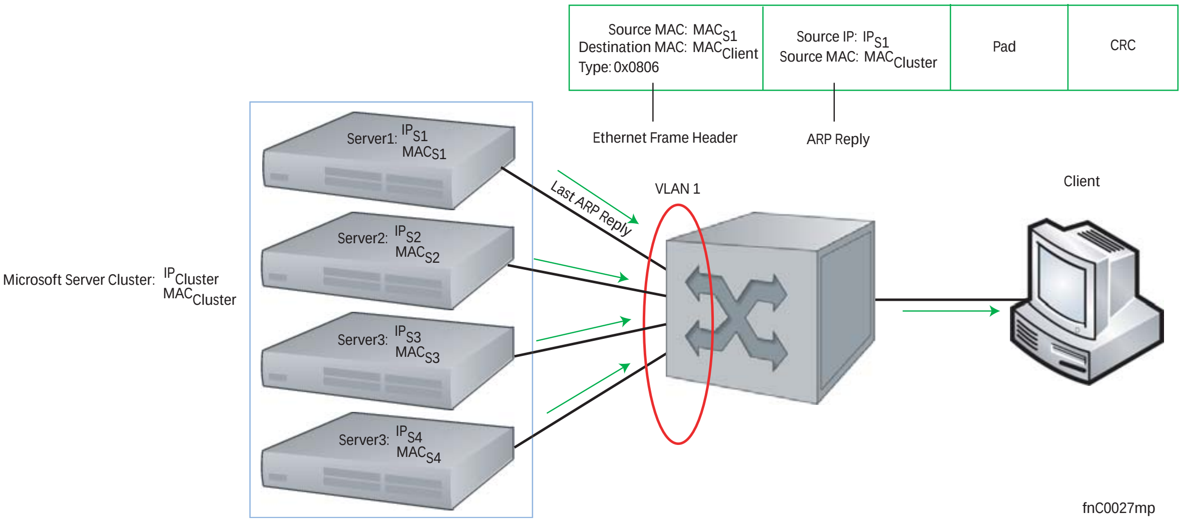
Source MAC: MAC_{S1}
Destination MAC: MAC_{Client}
Type: 0x0806
Source IP: IP_{S1}
Source MAC: $MAC_{Cluster}$
Pad
CRC
Ethernet Frame Header
ARP Reply
Server1: IP_{S1} MAC_{S1}
Server2: IP_{S2} MAC_{S2}
Server3: IP_{S3} MAC_{S3}
Server3: IP_{S4} MAC_{S4}
Microsoft Server Cluster: $IP_{Cluster}$ $MAC_{Cluster}$
Last ARP Reply
VLAN 1
Client
fnC0027mp

Figure 84 Server Clustering: Failover and Balancing Not Preserved

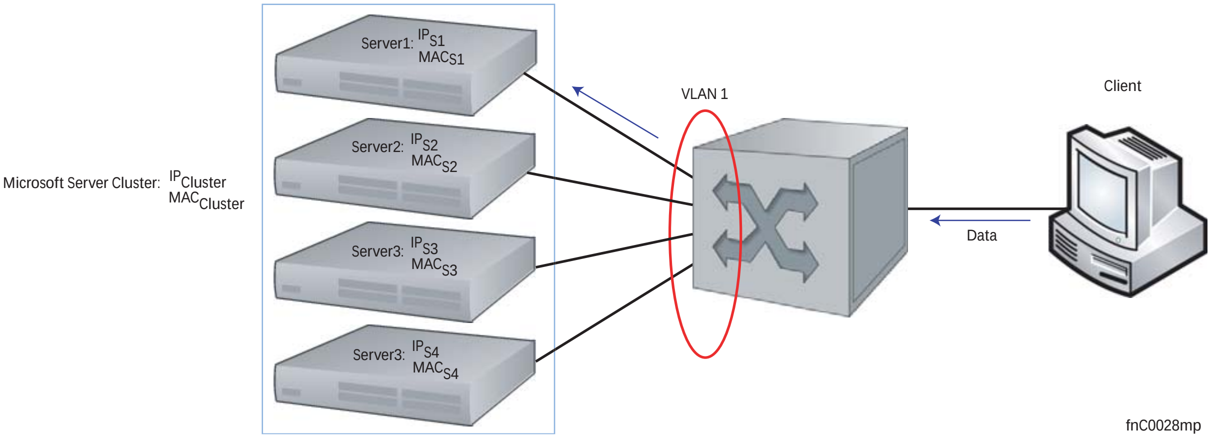
Server1: IP_{S1} MAC_{S1}
Server2: IP_{S2} MAC_{S2}
Server3: IP_{S3} MAC_{S3}
Server3: IP_{S4} MAC_{S4}
Microsoft Server Cluster: $IP_{Cluster}$ $MAC_{Cluster}$
VLAN 1
Client
Data
fnC0028mp

Configuring the Switch for Microsoft Server Clustering

To preserve failover and balancing, the Force10 switch must learn the cluster’s virtual MAC address, and it must forward traffic destined for the server cluster out all member ports in the VLAN connected to the cluster. To ensure that this happens, you must configure the command **vlan-flooding** on the Force10 switch at the time that that the Microsoft cluster is configured (Figure 85).

As shown in Figure 85, the server MAC address is given in the Ethernet frame header of the ARP reply, while the virtual MAC address representing the cluster is given in the pay. The command **vlan-flooding** directs the system to discover that there are different MAC addresses in an APR reply and associate the virtual MAC address with the VLAN connected to the cluster. Then, all traffic destined for the cluster is flooded out of all member ports. Since all of the servers in the cluster receive traffic, failover and balancing are preserved.

Figure 85 Server Cluster: Failover and Balancing Preserved with vlan-flooding

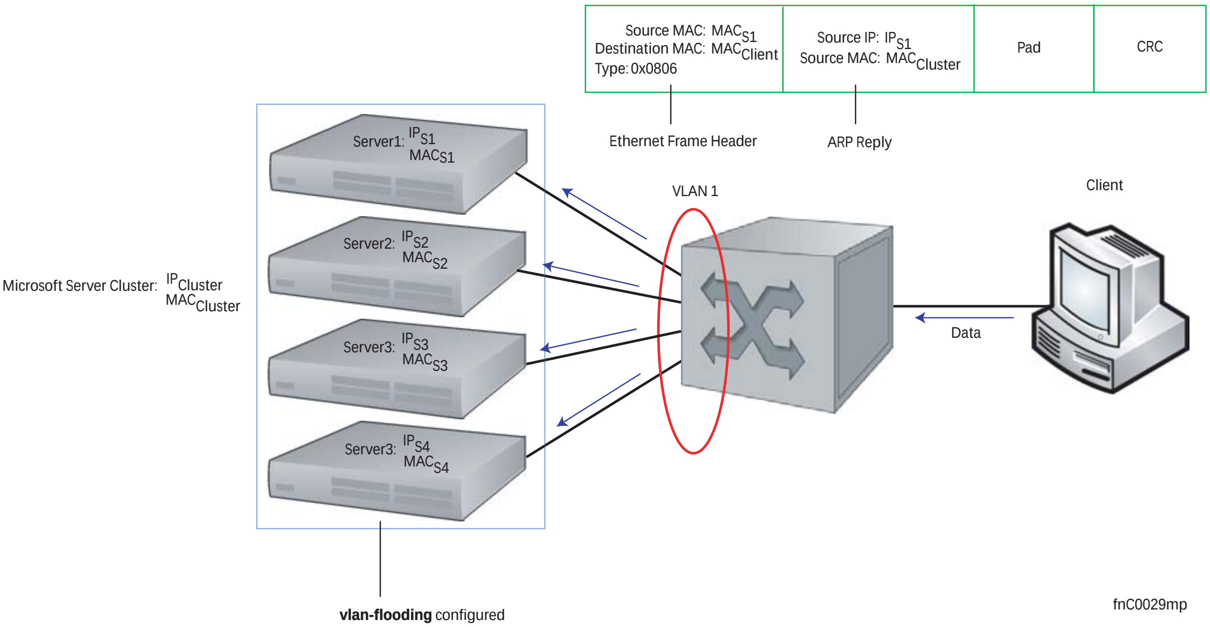

Enabling and Disabling vlan-flooding

- ARP entries already resolved through the VLAN are deleted when the feature is enabled. This ensures that ARP entries across the VLAN are consistent.
- All ARP entries learned after the feature is enabled are deleted when the feature is disabled, and RP2 triggers ARP resolution. The feature is disabled with the command **no vlan-flooding**.
- When a port is added to the VLAN, port automatically receives traffic if the feature is enabled. Old ARP entries are not deleted or updated.
- When a member port is deleted, its ARP entries are also deleted form the CAM.
- Port-channels in the VLAN also receive traffic.
- There is no impact on the configuration from saving the configuration or reing FTOS.
- The feature is not reflected in the output of the **show arp** command but is reflected in the output of the command **show ipf fib**.
- The ARP entries exist in the secondary RPM CAM, so failover has no effect on the feature.

Chapter 8 IPv6

C-Series	NO
E-Series	✓

Platform Specific Feature: IPv6 is supported on E-Series only.

Before using IPv6, change the CAM profile to the CAM **ipv6-extacl.** Once the CAM profile is changed, save the configuration and reboot your router.

CAM-Profile

Figure 86 displays the IPv6 CAM profiles. Enable **ipv6-extacl** before configuring IPv6.

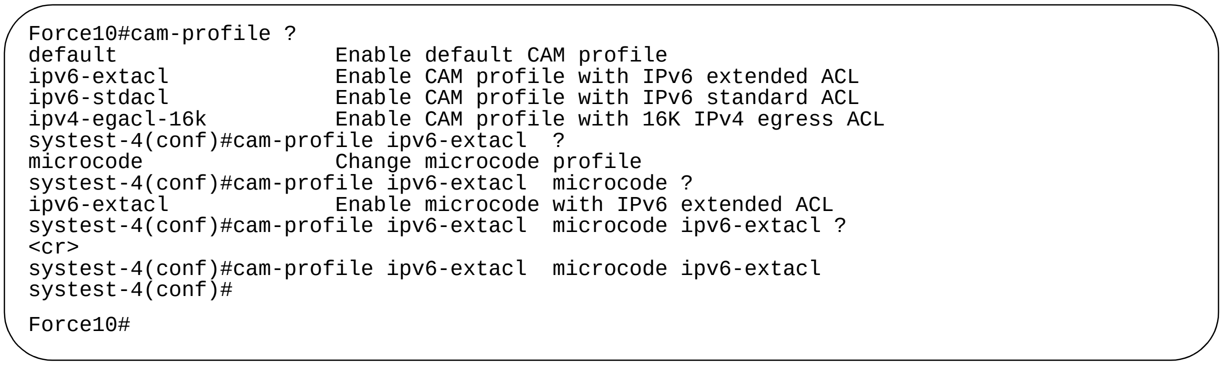

```
Force10#cam-profile ?
default                 Enable default CAM profile
ipv6-extacl             Enable CAM profile with IPv6 extended ACL
ipv6-stdacl             Enable CAM profile with IPv6 standard ACL
ipv4-egacl-16k          Enable CAM profile with 16K IPv4 egress ACL
systest-4(conf)#cam-profile ipv6-extacl  ?
microcode               Change microcode profile
systest-4(conf)#cam-profile ipv6-extacl  microcode ?
ipv6-extacl             Enable microcode with IPv6 extended ACL
systest-4(conf)#cam-profile ipv6-extacl  microcode ipv6-extacl ?
<cr>
systest-4(conf)#cam-profile ipv6-extacl  microcode ipv6-extacl
systest-4(conf)#

Force10#
```

Figure 86 cam-profile Command Example.

Figure 87 displays the IPv6 CAM profiles summary

```
Force10#show cam-profile summary
-- Chassis CAM Profile --
                 : Current Settings : Next Boot
Profile Name     : IPV6-ExtACL      : IPV6-ExtACL
MicroCode Name   : IPv6-ExtACL      : IPv6-ExtACL
-- Line card 11 --
                 : Current Settings : Next Boot
Profile Name     : IPV6-ExtACL      : IPV6-ExtACL
MicroCode Name   : IPv6-ExtACL      : IPv6-ExtACL
-- Line card 13 --
                 : Current Settings : Next Boot
Profile Name     : IPV6-ExtACL      : IPV6-ExtACL
MicroCode Name   : IPv6-ExtACL      : IPv6-ExtACL

Force10#
```

Figure 87 show cam-profile summary

Figure 88 displays the IPv6 CAM entries.

```
Force10#show cam-profile
-- Chassis CAM Profile --
CamSize           : 18-Meg
                  : Current Settings : Next Boot
Profile Name     : IPV6-ExtACL     : IPV6-ExtACL
L2FIB             : 32K entries      : 32K entries
L2ACL             : 1K entries       : 1K entries
IPv4FIB           : 192K entries     : 192K entries
IPv4ACL           : 13K entries      : 13K entries
IPv4Flow          : 8K entries       : 8K entries
EgL2ACL           : 1K entries       : 1K entries
EgIPv4ACL         : 1K entries       : 1K entries
Reserved          : 2K entries       : 2K entries
IPv6FIB           : 6K entries       : 6K entries
IPv6ACL           : 3K entries       : 3K entries
IPv6Flow          : 4K entries       : 4K entries
EgIPv6ACL         : 1K entries       : 1K entries
MicroCode Name   : IPv6-ExtACL     : IPv6-ExtACL

-- Line card 11 --
CamSize           : 18-Meg
                  : Current Settings : Next Boot
Profile Name     : IPV6-ExtACL     : IPV6-ExtACL
L2FIB             : 32K entries      : 32K entries
L2ACL             : 1K entries       : 1K entries
IPv4FIB           : 192K entries     : 192K entries
IPv4ACL           : 13K entries      : 13K entries
IPv4Flow          : 8K entries       : 8K entries
EgL2ACL           : 1K entries       : 1K entries
EgIPv4ACL         : 1K entries       : 1K entries
Reserved          : 2K entries       : 2K entries
IPv6FIB           : 6K entries       : 6K entries
IPv6ACL           : 3K entries       : 3K entries
IPv6Flow          : 4K entries       : 4K entries
EgIPv6ACL         : 1K entries       : 1K entries
MicroCode Name    : IPv6-ExtACL      : IPv6-ExtACL
```

Figure 88 show cam profile command Example

Configuration Task List for IPv6

The following list includes the configuration tasks for:

- Assigning an IPv6 Address to an Interface on page 191
- Assigning a Static IPv6 Route on page 192
- Viewing IPv6 Route Information on page 192
- Clearing IPv6 Routes on page 196

Assigning an IPv6 Address to an Interface

To assign an IPv6 address to an interface, use the **ipv6 address** command:

```
Force10(conf)#interface gigabitethernet 4/0
Force10(conf-if-gi-4/0)#ipv6 address ?
X:X:X:X::X              IPv6 address
Force10(conf-if-gi-4/0)#ipv6 address 3:3:3:3::3 ?
<0-128>                 Prefix length in bits
Force10(conf-if-gi-4/0)#ipv6 address 3:3:3:3::3 /64 ?
<cr>
Force10(conf-if-gi-4/0)#ipv6 address 3:3:3:3::3 /64
```

Figure 89 Configuring a Global IPv6 Address Example

Assigning a Static IPv6 Route

Use the **ipv6 route** command to configure IPv6 static routes. Figure 90 is an example of configuring IPv6 static routes:

```
Force10(conf)#ipv6 route ?
X:X:X:X::X              IPv6 prefix x:x::y
Force10(conf)#ipv6 route 33::0 ?
mask /nn                Mask in slash format
Force10(conf)#ipv6 route 33::0 /64 ?
X:X:X:X::X              Forwarding router's address
gigabitethernet         Gigabit Ethernet interface
loopback                Loopback interface
null                    Null interface
port-channel            Port-Channel interface
sonet                   Sonet Interface
tenGigabitethernet      TenGigabit Ethernet interface
vlan                    Vlan interface
Force10(conf)#ipv6 route 33::0 /64 22::1
Force10(conf)#ipv6 route 44::0 /64 33::1 ?
<1-255>                 Distance metric for this route
permanent               Permanent route
tag                     Set tag for this route
<cr>
Force10(conf)#ipv6 route 55::0 /64 gigabitethernet 9/0 ?
<1-255>                 Distance metric for this route
X:X:X:X::X              Forwarding router's address
permanent               Permanent route
tag                     Set tag for this route
<cr>
Force10(conf)#ipv6 route 55::0 /64 gigabitethernet 9/0 66::1 ?
<1-255>                 Distance metric for this route
permanent               Permanent route
tag                     Set tag for this route
<cr>
```

Figure 90 Configuring IPv6 Static Routes Example

Viewing IPv6 Route Information

The following commands allow you to view specific IPv6 configuration:

Command Syntax	Command Mode	Usage
show running-config [*entity*] [**configured**]	EXEC	Allows you to view the current configuration and indicates changes from the default values.
show ipv6 interface *interface* [**brief**] [**linecard** *slot-number*] [**configuration**]	EXEC	Display the IPv6 physical or loopback interfaces.

Command Syntax	Command Mode	Usage
show ipv6 route [*ipv6-address prefix-length*] [**longer-prefixes**] [**connected**] [**static**] [**summary**]	EXEC	Displays the IPv6 routes.
show ipv6 cam linecard *slot-number* [**summary**]	EXEC	Displays the IPv6 CAM in a specified line card.

Note: ISIS, list, and RIP are not supported in this release.

show running-configuration interface

Figure 91 displays a running configuration for an interface configured with an IPv6 route:

```
Force10#show run int gi 2/3
!
interface GigabitEthernet 2/3
no ip address
ipv6 address 2222:2222:3333:3333:0:3333:0:7777 /32
no shutdown
```

Figure 91 show running-configuration Command Example

show ipv6 interface

Figure 92 displays IPv6 information on a specified interface and line card:

```
Force10#show ipv6 interface gigabitethernet 1/1
GigabitEthernet 1/1 is up, line protocol is up
  IPV6 is enabled
  Link Local address: fe80::201:e8ff:fe04:62c4
  Global Unicast address(es):
    2001::1, subnet is 2001::/64
    2002::1, subnet is 2002::/120
    2003::1, subnet is 2003::/120
    2004::1, subnet is 2004::/32
  Global Anycast address(es):
  Joined Group address(es):
    ff02::1
    ff02::2
    ff02::1:ff00:1
    ff02::1:ff04:62c4
    MTU is 1500
  ICMP redirects are not sent
  DAD is enabled: number of DAD attempts: 1
  ND reachable time is 30 seconds
  ND advertised reachable time is 30 seconds
  ND advertised retransmit interval is 30 seconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
```

Figure 92 show ipv6 cam linecard Command Example

show ipv6 route

Figure 93 displays IPv6 route information in routing tables:

```
Force10#show ipv6 route

Codes: C - connected, L - local, S - static, R - RIP,
       B - BGP, IN - internal BGP, EX - external BGP,LO - Locally Originated,
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,
       L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
Gateway of last resort is not set

       Destination  Dist/Metric, Gateway, Last Change
       ----------------------------------------------------
  C    2001::/64 [0/0]
        Direct, Gi 1/1, 00:28:49
  C    2002::/120 [0/0]
        Direct, Gi 1/1, 00:28:49
  C    2003::/120 [0/0]
        Direct, Gi 1/1, 00:28:49
  C    2004::/32 [0/0]
        Direct, Gi 1/1, 00:28:49
  L    fe80::/10 [0/0]
        Direct, Nu 0, 00:29:09
```

Figure 93 show ipv6 route Command Example

show ipv6 route summary

Figure 94 displays IPv6 route information in summary format:

```
Force10#show ipv6 route summary

Route Source              Active Routes   Non-active Routes
connected                 5               0
static                    0               0
Total                     5               0
Total 5 active route(s) using 952 bytes
```

Figure 94 show ipv6 route summary Command Example

show ipv6 route static

Figure 95 display IPv6 static route information:

```
Force10#show ipv6 route static
Destination Dist/Metric, Gateway, Last Change
-----------------------------------------------------
  S   8888:9999:5555:6666:1111:2222::/96 [1/0]
       via 2222:2222:3333:3333::1, Gi 9/1, 00:03:16
  S   9999:9999:9999:9999::/64 [1/0]
       via 8888:9999:5555:6666:1111:2222:3333:4444, 00:03:16
```

Figure 95 show ipv6 static Command Example

Figure 96 show ipv6 cam linecard n summary Command Example

show running-configuration static

The following example is a running configuration entry for a static route:

```
Force10#show run static

!

ipv6 route 5555:5555:0000:0000:0000:0000:0000:0000 32
4444:4444:6666:6666:0000:3333:0000:7777
```

Figure 97 show running-configuration static Command Example

Clearing IPv6 Routes

Use the clear IPv6 route command to clear routes from the IPv6 routing table.

Command Syntax	Command Mode	Purpose
clear ipv6 route {* | *ipv6-address prefix-length*}	EXEC	Clear (refresh) all or a specific routes from the IPv6 routing table.

Chapter 9 Interfaces

This chapter contains information on configuring interfaces, both physical and logical, with FTOS. This chapter discusses interface types; it also covers the following types of interfaces that are available on the E-Series:

- Interface Modes on page 197
- Viewing Interface Information on page 199
- Using Interface Range Macros on page 203
- Physical Interfaces on page 204
- Management Interface on page 209
- Loopback Interfaces on page 211
- Null Interface on page 211
- Port Channel Interfaces on page 212
- VLAN Interfaces and Layer 3 on page 222
- Bulk Configuration on page 223
- Time Domain Reflectometry on page 226

FTOS supports the MIB-II (RFC 1213) and the Interfaces Group MIB (RFC 2863).

Interface Modes

In the E-Series system, you can place physical, VLANs, and port channel interfaces in two different modes: Layer 2 or Layer 3 mode (Table 12).

Table 12 Interfaces in the E-Series System

Type of Interface	Modes Possible	Require Creation	Default State
10/100/1000 Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet	Layer 2 Layer 3	No	Shutdown (disabled)
SONET (PPP encapsulation)	Layer 3	No	Shutdown (disabled)
Management	n/a	No	Shutdown (disabled)
Loopback	Layer 3	Yes	No shutdown (enabled)
Null interface	n/a	No	Enabled

Table 12 Interfaces in the E-Series System

Type of Interface	Modes Possible	Require Creation	Default State
Port Channel	Layer 2 Layer 3	Yes	Shutdown (disabled)
VLAN	Layer 2 Layer 3	Yes, except for the Default VLAN	No shutdown (active for Layer 2) Shutdown (disabled for Layer 3)

To place a physical or port channel interface in Layer 2 mode, use the **switchport** command; to place an interface in Layer 3 mode, assign an IP address to that interface. These interfaces also contain Layer 2 and Layer 3 commands to configure and modify the interfaces. VLANs are different and, by default, these interfaces are in Layer 2 mode.

Layer 2 Mode

Use the **switchport** command to place an interface in Layer 2 mode. You cannot configure switching or Layer 2 protocols such as Spanning Tree Protocol on the interface unless the interface is in Layer 2 mode.

Figure 98 displays the basic configuration found in a Layer 2 interface.

```
Force10(conf-if)#show config
!
interface Port-channel 1
 no ip address
 switchport
 no shutdown
Force10(conf-if)#
```

Figure 98 show config Command Example of a Layer 2 Interface

Layer 3 Mode

To enable Layer 3 traffic on the interface, add an IP address to the interfaces using the **ip address** command and **no shutdown** command in the INTERFACE mode. In all interface types but VLANs, the **shutdown** command prevents all traffic from passing through the interface. In VLANs, the **shutdown** command prevents Layer 3 traffic from passing through the interface, yet Layer 2 traffic is unaffected by this command. One of the interfaces in the E-Series system must be in Layer 3 mode before you configure or enter a Layer 3 protocol mode (for example, OSPF).

Figure 99 displays the **show config** command example of a Layer 3 interface.

```
Force10(conf-if)#show config
!
interface GigabitEthernet 1/5
 ip address 10.10.10.1 /24
 no shutdown
Force10(conf-if)#
```

Figure 99 show config Command Example of a Layer 3 Interface

When an interface is in either mode, you receive an error message if you try to configure a command that must be in the other mode. For example, in Figure 100, the command **ip address** triggered an error message because the interface is in Layer 2 mode and the **ip address** command is a Layer 3 command.

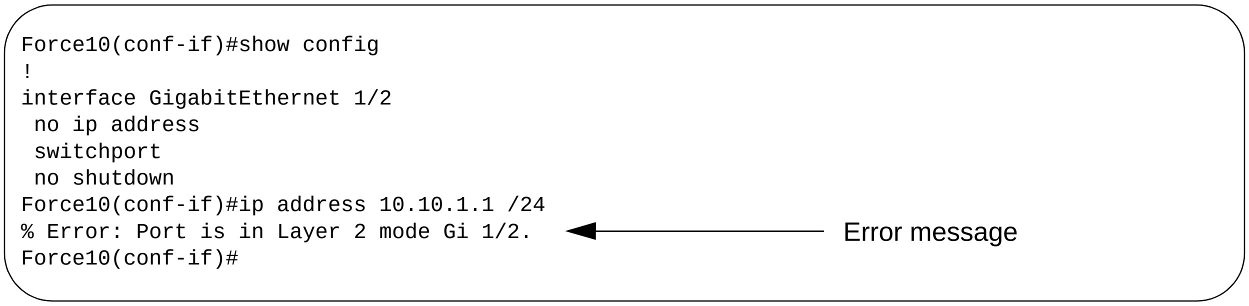

```
Force10(conf-if)#show config
!
interface GigabitEthernet 1/2
 no ip address
 switchport
 no shutdown
Force10(conf-if)#ip address 10.10.1.1 /24
% Error: Port is in Layer 2 mode Gi 1/2.
Force10(conf-if)#
```

Error message

Figure 100 Error Message When Trying to Add an IP Address to Layer 2 Interface

To determine the configuration of an interface, you can use **show config** command in the INTERFACE mode or the various **show interfaces** commands in the EXEC mode.

Viewing Interface Information

To view interface status and configuration, you have multiple choices. The **show interfaces** command in the EXEC mode displays the interface status, IP and MAC addresses, and multiple counters for the amount and type of traffic passing through the interface. If a port channel interface is configured, the **show interfaces** command lists the interfaces configured in the port channel.

Figure 101 displays the configuration and status information for one interface.

```
Force10#show interfaces tengigabitethernet 1/0
TenGigabitEthernet 1/0 is up, line protocol is up
Hardware is Force10Eth, address is 00:01:e8:05:f3:6a
    Current address is 00:01:e8:05:f3:6a
Pluggable media present, XFP type is 10GBASE-LR.
    Medium is MultiRate, Wavelength is 1310nm
    XFP receive power reading is -3.7685
Interface index is 67436603
Internet address is 65.113.24.238/28
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 10000 Mbit, Mode full duplex, Master
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:09:54
Queueing strategy: fifo
Input Statistics:
     0 packets, 0 bytes
     0 Vlans
     0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     0 Multicasts, 0 Broadcasts
     0 runts, 0 giants, 0 throttles
     0 CRC, 0 overrun, 0 discarded
Output Statistics:
     3 packets, 192 bytes, 0 underruns
     3 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     0 Multicasts, 3 Broadcasts, 0 Unicasts
     0 Vlans, 0 throttles, 0 discarded, 0 collisions
Rate info (interval 299 seconds):
     Input 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
     Output 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:00:31
```

Figure 101 show interfaces Command Example

In the EXEC mode, the **show interfaces switchport** command displays only interfaces in Layer 2 mode and their relevant configuration information. The **show interfaces switchport** command (Figure 102) displays the interface, whether the interface supports IEEE 802.1Q tagging or not, and the VLANs to which the interface belongs.

```
Force10#show interfaces switchport
Name: GigabitEthernet 13/0
802.1QTagged: True
Vlan membership:
Vlan    2


Name: GigabitEthernet 13/1
802.1QTagged: True
Vlan membership:
Vlan    2


Name: GigabitEthernet 13/2
802.1QTagged: True
Vlan membership:
Vlan    2


Name: GigabitEthernet 13/3
802.1QTagged: True
Vlan membership:
Vlan    2


--More--
```

Figure 102 show interfaces switchport Command Example

Use the **show ip interfaces brief** command in the EXEC privilege mode to view which interfaces are in Layer 3 mode. In Figure 103, GigabitEthernet 1/5 is in Layer 3 mode since it has an IP address assigned to it and its status is `up`.

```
Force10#show ip interface brief
Interface              IP-Address      OK? Method Status                Protocol
GigabitEthernet 1/0    unassigned      NO  Manual administratively down down
GigabitEthernet 1/1    unassigned      NO  Manual administratively down down
GigabitEthernet 1/2    unassigned      YES Manual up                    up
GigabitEthernet 1/3    unassigned      YES Manual up                    up
GigabitEthernet 1/4    unassigned      YES Manual up                    up
GigabitEthernet 1/5    10.10.10.1      YES Manual up                    up
GigabitEthernet 1/6    unassigned      NO  Manual administratively down down
GigabitEthernet 1/7    unassigned      NO  Manual administratively down down
GigabitEthernet 1/8    unassigned      NO  Manual administratively down down
```

Figure 103 show ip interfaces brief Command Example (Partial)

Use the **show interfaces configured** command in the EXEC privilege mode to view only configured interfaces. In Figure 104, GigabitEthernet 1/5 is in Layer 3 mode since it has an IP address assigned to it and its status is `up`.

Displaying Only Configured Interfaces

The following options have been implemented for **show [ip | running-config] interfaces** commands for (only) linecard interfaces. When the **configured** keyword is used, only interfaces that have non-default configurations are displayed. Dummy linecard interfaces (created with the **linecard** command) are treated like any other physical interface.

Figure 104 lists the possible show commands that have the configured keyword available:

```
Force10#show interfaces configured
Force10#show interfaces linecard 0 configured
Force10#show interfaces gigabitEthernet 0 configured
Force10#show ip interface configured
Force10#show ip interface linecard 1 configured
Force10#show ip interface gigabitEthernet 1 configured
Force10#show ip interface br configured
Force10#show ip interface br linecard 1 configured
Force10#show ip interface br gigabitEthernet 1 configured
Force10#show running-config interfaces configured
Force10#show running-config interface gigabitEthernet 1 configured
```

Figure 104 show Commands with configured Keyword Examples

Rate-interval

The interface rate interval that displays in the output of **show interfaces** can be changed from the default value of 299 to any value between 30 and 299 seconds. Use the **rate-interval** command under interface configuration mode, to configure the desired rate interval.

Though the rate is configurable with any value between 30 and 299, the nearest (floor) multiple of 15 is used. This is because software polling is done once every 15 seconds. So, “30-44” means 30, and “45-59” means 45; etc.

All the LAG members inherit the rate-interval configuration from the LAG.

Figure 105 shows rate interval and configuring to change the rate interval default value:

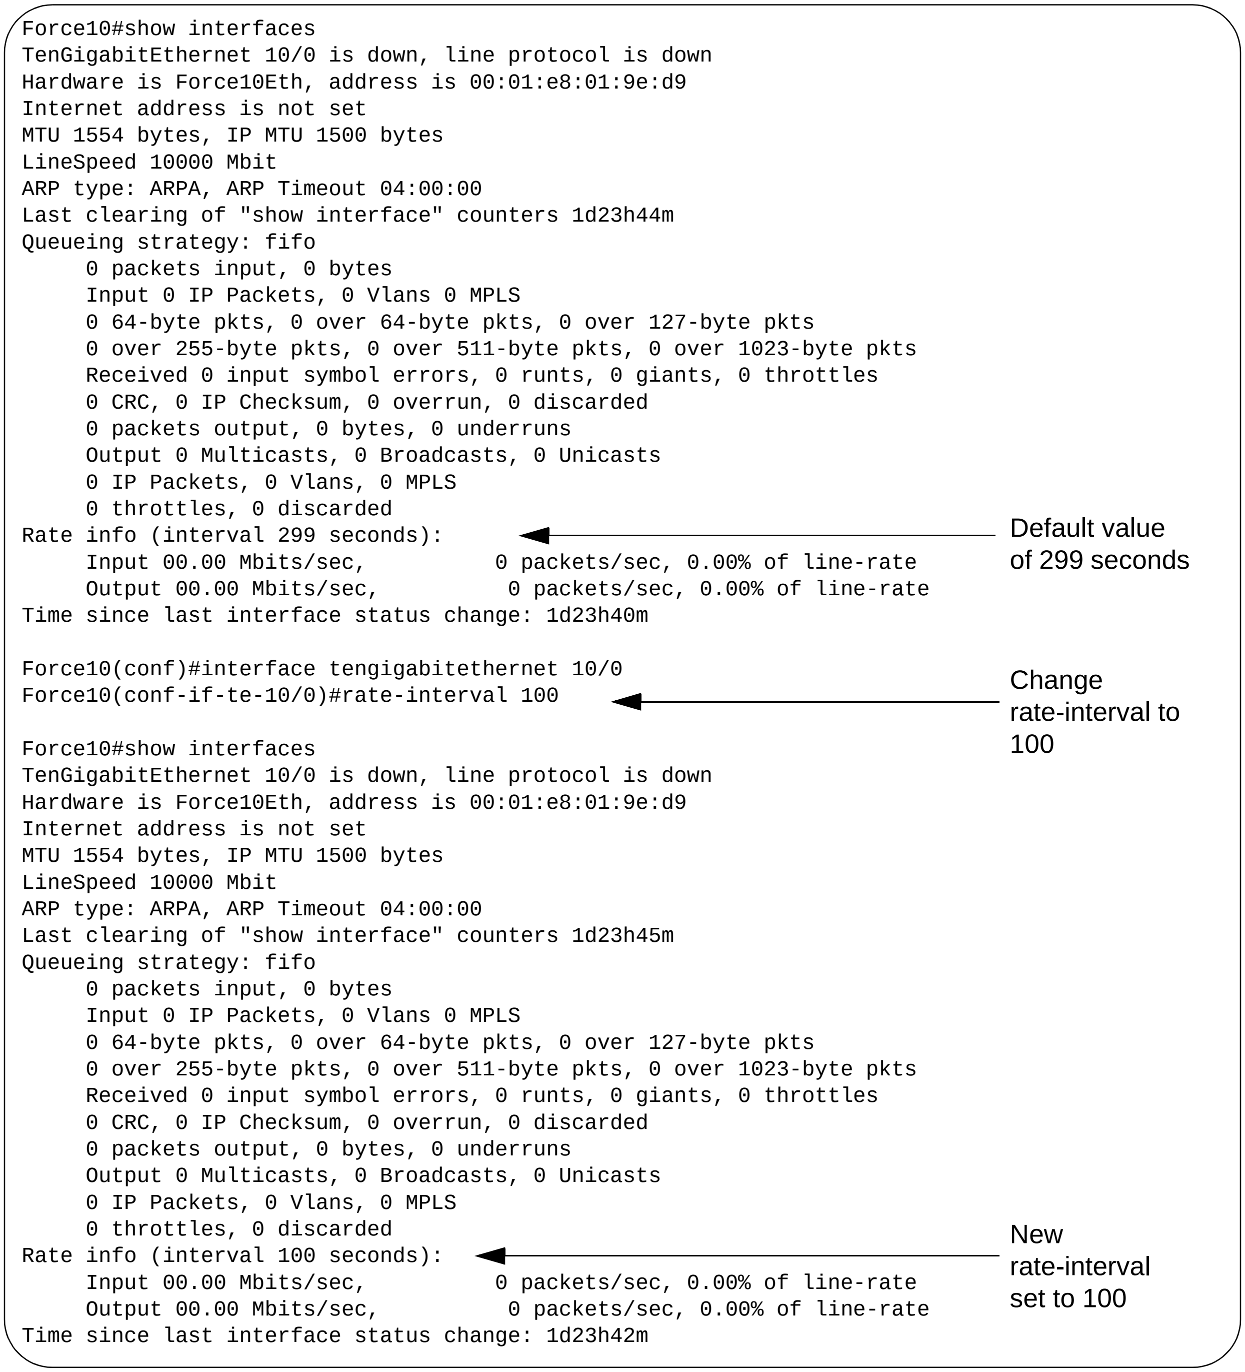

```
Force10#show interfaces
TenGigabitEthernet 10/0 is down, line protocol is down
Hardware is Force10Eth, address is 00:01:e8:01:9e:d9
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 10000 Mbit
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 1d23h44m
Queueing strategy: fifo
     0 packets input, 0 bytes
     Input 0 IP Packets, 0 Vlans 0 MPLS
     0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     Received 0 input symbol errors, 0 runts, 0 giants, 0 throttles
     0 CRC, 0 IP Checksum, 0 overrun, 0 discarded
     0 packets output, 0 bytes, 0 underruns
     Output 0 Multicasts, 0 Broadcasts, 0 Unicasts
     0 IP Packets, 0 Vlans, 0 MPLS
     0 throttles, 0 discarded
Rate info (interval 299 seconds):
     Input 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
     Output 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
Time since last interface status change: 1d23h40m

Force10(conf)#interface tengigabitethernet 10/0
Force10(conf-if-te-10/0)#rate-interval 100

Force10#show interfaces
TenGigabitEthernet 10/0 is down, line protocol is down
Hardware is Force10Eth, address is 00:01:e8:01:9e:d9
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 10000 Mbit
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 1d23h45m
Queueing strategy: fifo
     0 packets input, 0 bytes
     Input 0 IP Packets, 0 Vlans 0 MPLS
     0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
     0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
     Received 0 input symbol errors, 0 runts, 0 giants, 0 throttles
     0 CRC, 0 IP Checksum, 0 overrun, 0 discarded
     0 packets output, 0 bytes, 0 underruns
     Output 0 Multicasts, 0 Broadcasts, 0 Unicasts
     0 IP Packets, 0 Vlans, 0 MPLS
     0 throttles, 0 discarded
Rate info (interval 100 seconds):
     Input 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
     Output 00.00 Mbits/sec,          0 packets/sec, 0.00% of line-rate
Time since last interface status change: 1d23h42m
```

Default value of 299 seconds
Change rate-interval to 100
New rate-interval set to 100

Figure 105 Configuring for Rate Interval Example

Dynamic Counters

By default, counting for the following four applications are enabled:

- IPFLOW
- IPACL
- L2ACL
- L2FIB

For remaining applications, FTOS automatically turns on counting when the application is enabled, and is turned off when the application is disabled. Please note that if more than four counter-dependent applications are enabled on a port pipe, there is an impact on line rate performance.

The following counter-dependent applications are supported by FTOS:

- Egress VLAN
- Ingress VLAN
- Next Hop 2
- Next Hop 1
- Egress ACLs
- ILM
- IP FLOW
- IP ACL
- IP FIB
- L2 ACL
- L2 FIB

Using Interface Range Macros

You can define an interface-range macro to automatically select a range of interfaces for configuration. Before you can use the macro keyword in the interface-range macro command string, you must define the macro.

To define an interface-range macro, enter this command:

Command Syntax	Command Mode	Purpose
Force10 (config)# **define** *interface-range macro_name* {**vlan** *vlan_ID* - *vlan_ID*} **|** {{**gigabitethernet | tengigabitethernet**} *slot/interface* - *interface*} [**,** {**vlan** *vlan_ID* - *vlan_ID*} {{**gigabitethernet | tengigabitethernet**} *slot/interface* - *interface*}]	CONFIGURATION	Defines the interface-range macro and saves it in the running configuration file.

Define the Interface Range

This example shows how to define an interface-range macro named “test” to select Fast Ethernet interfaces 5/1 through 5/4:

```
Force10(config)# define interface-range test gigabitethernet 5/1 - 4
```

To show the defined interface-range macro configuration, use the command **show running-config** in the EXEC mode. The example below shows how to display the defined interface-range macro named “test”:

```
Force10# show running-config | include define

define interface-range test GigabitEthernet5/1 - 4

Force10#
```

Choosing an Interface-range Macro

To use an interface-range macro in the **interface range** command, enter this command:

Command Syntax	Command Mode	Purpose
interface range macro *name*	CONFIGURATION	Selects the interfaces range to be configured using the values saved in a named interface-range macro.

The example below shows how to change to the interface-range configuration mode using the interface-range macro named “test”:

```
Force10(config)# interface range macro test

Force10(config-if)#
```

Physical Interfaces

Four physical interface types are available on the E-Series line cards: 10/100/1000 Ethernet, Gigabit Ethernet and 10 Gigabit Ethernet interfaces, and SONET interfaces. Another interface, found on the Route Processor Module (RPM), is the management Ethernet interface. This interface type is Fast Ethernet and provides management access to the E-Series.

The line card interfaces support Layer 2 and Layer 3 traffic over the 10/100/1000, Gigabit, and 10-Gigabit Ethernet interfaces. SONET interfaces with PPP encapsulation support Layer 3 traffic. These interfaces (except SONET interfaces with PPP encapsulation) also can become part of virtual interfaces such as a VLAN or port channels.

For more information on VLANs, see VLAN Interfaces and Layer 3 on page 222 and for more information on port channels, see Port Channel Interfaces on page 212.

Auto Negotiation on Ethernet Interfaces

By default, auto negotiation is enabled on EtherScale and TeraScale interfaces of 100/1000 Base-T Ethernet line cards and Fiber line cards. Only the 10GE line cards do not support auto negotiation. When using 10GE line cards, verify that the settings on the connecting devices are set to no auto negotiation.

With the 10/100/1000 Ethernet line cards, the **negotiation auto** command is tied to the **speed** command. Auto negotiation is always enabled when the **speed** command is set to 1000 or auto.

The **negotiation auto** command provides a mode option for configuring an individual port to forced master/forced slave once auto negotiation is enabled.

Caution: Ensure that one end of your node is configured as forced-master and one is configured as forced-slave. If both are configured the same (that is forced-master or forced-slave), the show interface command will flap between an auto-neg-error and forced-master/slave states.

Configuration Task List for Physical Interfaces

By default, all 10/100/1000, Gigabit and 10 Gigabit Ethernet and SONET interfaces are disabled to traffic.

The following list includes the configuration task for physical interfaces:

- enable an interface on page 206 (mandatory)
- configure Layer 2 mode on page 207 (optional)
- configure Layer 3 mode on page 207 (optional)
- clear interface counters on page 208 (optional)

For a complete listing of all commands related to physical interfaces, refer to the .

enable an interface

To determine which physical interfaces are available, use the **show running-config** command in the EXEC mode. This command displays all physical interfaces available on the E-Series line cards (Figure 106).

```
Force10#show running
Current Configuration ...
!
interface GigabitEthernet 9/6
 no ip address
 shutdown
!
interface GigabitEthernet 9/7
 no ip address
 shutdown
!
interface GigabitEthernet 9/8
 no ip address
 shutdown
!
interface GigabitEthernet 9/9
 no ip address
 shutdown
```

Figure 106 Interfaces listed in the show running-config Command (Partial)

As soon as you determine the type of physical interfaces, you must enter the INTERFACE mode to enable and configure the interfaces.

To enter the INTERFACE mode, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**interface** *interface*	CONFIGURATION	Enter the keyword **interface** followed by the type of interface and slot/port information: • For a 10/100/1000 Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For the Management interface on the RPM, enter the keyword **ManagementEthernet** followed by the slot/port information. • For a SONET interface, enter the keyword **sonet** followed by slot/port information. • For a 10 Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information.
2	**no shutdown**	INTERFACE	Enable the interface. If the interface is a SONET interface, enter the **encap ppp** command to enable PPP encapsulation. After encapsulation is enabled, enter **no shutdown** to enable the interface.

To confirm that the interface is enabled, use the **show config** command in the INTERFACE mode.

To leave the INTERFACE mode, use the **exit** command or **end** command.

You cannot delete a physical interface.

configure Layer 2 mode

As stated, you must place interfaces in Layer 2 mode to configure Layer 2 protocols on the interface, such as Spanning Tree Protocol.

To configure an interface in Layer 2 mode, use these commands in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
no shutdown	INTERFACE	Enable the interface.
switchport	INTERFACE	Place the interface in Layer 2 (switching) mode.

For information on enabling and configuring Spanning Tree Protocol, see Chapter 7, Layer 2, on page 161. To view the interfaces in Layer 2 mode, use the command **show interfaces switchport** (Figure 102) in the EXEC mode.

configure Layer 3 mode

By assigning an IP address to a physical interface, you place it in Layer 3 mode. Routed traffic now passes through the interface and you can configure routing protocols on that interface.

To assign an IP address, use both of these commands in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
no shutdown	INTERFACE	Enable the interface.
ip address *ip-address mask* [**secondary**]	INTERFACE	Configure a primary IP address and mask on the interface. The *ip-address* must be in dotted-decimal format (A.B.C.D) and the *mask* must be in slash format (/xx). Add the keyword **secondary** if the IP address is the interface's backup IP address.

You can only configure one primary IP address per interface.

To view all interfaces to see which have an IP address assigned, use the **show ip interfaces brief** command (Figure 103).

To view IP information on an interface in Layer 3 mode, use the **show ip interface** command in the EXEC privilege mode (Figure 107).

```
Force10>show ip int vlan 58
Vlan 58 is up, line protocol is up
Internet address is 1.1.49.1/24
Broadcast address is 1.1.49.255
Address determined by config file
MTU is 1554 bytes
Inbound  access list is not set
Proxy ARP is enabled
Split Horizon is enabled
Poison Reverse is disabled
ICMP redirects are not sent
ICMP unreachables are not sent
```

Figure 107 show ip interface Command Example

clear interface counters

The counters in the **show interfaces** command are reset by the **clear counters** command. This command does not clear the counters captured by any SNMP program.

To clear the counters, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Purpose
clear counters [*interface*] [**vrrp** [*vrid*] | **learning-limit**]	EXEC privilege	Clear the counters used in the show interface commands for all VRRP groups, VLANs, and physical interfaces or selected ones. Without an interface specified, the command clears all interface counters. (OPTIONAL) Enter the following interface keywords and slot/port or number information: • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a Loopback interface, enter the keyword **loopback** followed by a number from 0 to 16383. • For a Port Channel interface, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For the management interface on the RPM, enter the keyword **ManagementEthernet** followed by slot/port information. The slot range is 0-1, and the port range is 0. • For a SONET interface, enter the keyword **sonet** followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN, enter the keyword **vlan** followed by a number from 1 to 4094 (OPTIONAL) Enter the keyword **vrrp** to clear statistics for all VRRP groups configured. Enter a number from 1 to 255 as the *vrid*. (OPTIONAL) Enter the keyword **learning-limit** to clear unknown source address (SA) drop counters when MAC learning limit is configured on the interface.

When you enter this command, you must confirm that you want FTOS to clear the interface counters for that interface (Figure 108).

```
Force10#clear counters gi 0/0
Clear counters on GigabitEthernet 0/0 [confirm]
Force10#
```

Figure 108 Clearing an Interface

Management Interface

The management interface is located on the RPM and provides management access to the E-Series system. You can configure this interface with FTOS, but the configuration options on this interface are limited; you cannot configure a gateway address or an IP address that appears in the main routing table of FTOS. In addition, Proxy ARP is not supported on this interface.

To configure a Management interface, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
interface ManagementEthernet *interface*	CONFIGURATION	Enter the slot (0-1) and the port (0). In a system with 2 RPMs, therefore, 2 Management interfaces, the slot number differentiates between the two Management interfaces.

To view the Primary RPM Management port, use the **show interface ManagementEthernet** command in the EXEC privilege mode. If there are 2 RPMs in the system, you cannot view information on that interface.

To configure IP address on a Management interface, use the following command in the MANAGEMENT INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip address *ip-address mask*	INTERFACE	Configure an IP address and mask on the interface. • *ip-address mask:* enter an address in dotted-decimal format (A.B.C.D) and mask must be in /prefix format (/x)

If you have two RPMs in your system, each Management interface must be configured with a different IP address. Unless the **management route** command is configured, you can only access the Management interface from the local LAN. To access the Management interface from another LAN, you must configure the **management route** command to point to the Management interface.

Alternatively, you can use **virtual-ip** to manage a system with one or two RPMs. A virtual IP is an IP address assigned to the system (not to any management interfaces) and is a CONFIGURATION mode command. When a virtual IP address is assigned to the system, the active management interface of the RPM is recognized by the virtual IP address—not by the actual interface IP address assigned to it. During an RPM failover, you do not have to remember the IP address of the new RPM's management interface—the system will still recognizes the virtual-IP address.

Important Things to Remember —virtual-ip

- **virtual-ip** is a CONFIGURATION mode command.
- When applied, the management port on the primary RPM assumes the virtual IP address. Executing **show interfaces** and **show ip interface brief** commands on the primary RPM management interface will display the virtual IP address and not the actual IP address assigned on that interface.
- A duplicate IP address message is printed for management port's virtual IP address on an RPM failover. This is a harmless error that is generated due to a brief transitory moment during failover when both RPMs' management ports own the virtual IP address, but have different MAC addresses.
- The primary management interface will use only the virtual IP address if it is configured. The system can not be accessed through the native IP address of the primary RPM's management interface.

- Once the virtual IP address is removed, the system is accessible through the native IP address of the primary RPM's management interface.
- Primary and secondary management interface IP and virtual IP must be in the same subnet.

Loopback Interfaces

A Loopback interface is a virtual interface in which the software emulates an interface. Packets routed to it are processed locally. Since this interface is not a physical interface, you can configure routing protocols on this interface to provide protocol stability. You can place Loopback interfaces in default Layer 3 mode.

To configure a Loopback interface, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
interface loopback *number*	CONFIGURATION	Enter a number as the loopback interface. Range: 0 to 16383.

To view Loopback interface configurations, use the **show interface loopback** *number* command in the EXEC mode.

To delete a Loopback interface, use the **no interface loopback** *number* command syntax in the CONFIGURATION mode.

Many of the same commands found in the physical interface are found in Loopback interfaces.

See also Configuring ACLs to Loopback on page 313.

Null Interface

The Null interface is another virtual interface created by the E-Series software. There is only one Null interface. It is always up, but no traffic flows on this interface.

To enter the INTERFACE mode of the Null interface, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
interface null 0	CONFIGURATION	Enter the INTERFACE mode of the Null interface.

The only configurable command in the INTERFACE mode of the Null interface is the **ip unreachable** command.

Port Channel Interfaces

Port channel interfaces support link aggregation, as described in IEEE Standard 802.3ad.

This section covers the following topics:

Port Channel Definition and Standards

Link aggregation is defined by IEEE 802.3ad as a method of grouping multiple physical interfaces into a single logical interface—a Link Aggregation Group (LAG) or port channel. A LAG is “a group of links that appear to a MAC client as if they were a single link” according to IEEE 802.3ad. In FTOS, a LAG is referred to as a port channel interface.

This logical interface provides redundancy by allowing the aggregation of up to 16 physical interfaces into one logical interface. If one physical interface goes down in the port channel, another physical interface carries the traffic.

Port Channel Benefits

In the E-Series, a port channel interface provides many benefits, including easy management, link redundancy, and sharing.

Port channels are transparent to the network and can be configured and managed as one interface. For example, you configure one IP address for the group and that IP address is used for all routed traffic on the port channel.

With this feature, you can get larger-capacity interfaces with lower-speed links. For example, you can build a 5-Gigabit interface by aggregating five 1-Gigabit Ethernet interfaces together. If one of the five interfaces fails, traffic is redistributed across the four remaining interfaces.

Port Channel Implementation

FTOS supports two types of port channels:

Static—Port channels that are statically configured

Dynamic—Port channels that are dynamically configured using Link Aggregation Control Protocol (LACP).

The user can configure up to 255 port channels per E-Series (255 for TeraScale, 1 to 32 for EtherScale). As soon as a port channel is configured, the FTOS treats it like a physical interface. For example, the IEEE 802.1Q tagging is maintained while the physical interface is in the port channel.

As of FTOS Version 7.4.1.0, member ports of a LAG are added and programmed into hardware in a predictable order based on the port ID, instead of in the order in which the ports come up. With this implementation, load balancing yields predictable results across line card resets and chassis reloads.

A physical interface can belong to only one port channel at a time.

Each port channel can contain up to 16 Ethernet interfaces of the same interface type/speed, but located on different line cards.

Port channels can contain a mix of 10/100/1000 Ethernet interfaces and Gigabit Ethernet interfaces, and the interface speed (100 or 1000 Mb/s) used by the port channel is determined by the first port channel member that is physically up. FTOS disables the interfaces that do match the interface speed set by the first channel member. That first interface may be the first interface that is physically brought up or was physically operating when interfaces were added to the port channel. For example, if the first operational interface in the port channel is a Gigabit Ethernet interface, all interfaces at 1000 Mb/s are kept up, and all 10/100/1000 interfaces that are not set to 1000 speed or auto negotiate are disabled.

FTOS brings up 10/100/1000 interfaces that are set to auto negotiate so that their speed is identical to the speed of the first channel member in the port channel.

10/100/1000 interfaces in port channels

When both 10/100/1000 interfaces and GigE interfaces are added to a port channel, the interfaces must share a common speed. When interfaces have a configured speed different from the port channel speed, the software disables those interfaces.

The common speed is determined when the port channel is first enabled. At that time, the software checks the first interface listed in the port channel configuration. If that interface is enabled, its speed configuration becomes the common speed of the port channel. If the other interfaces configured in that port channel are configured with a different speed, FTOS disables them.

For example, if four interfaces (Gi 0/0, 0/1, 0/2, 0/3) in which Gi 0/0 and Gi 0/3 are set to speed 100 Mb/s and the others are set to 1000 Mb/s, with all interfaces enabled, and you add them to a port channel by entering **channel-member gigabitethernet 0/0-3** while in the port channel interface mode, and FTOS determines if the first interface specified (Gi 0/0) is up. Once it is up, the common speed of the port channel is 100 Mb/s. FTOS disables those interfaces configured with speed 1000 or whose speed is 1000 Mb/s as a result of auto-negotiation.

In this example, you can change the common speed of the port channel by changing its configuration so the first enabled interface referenced in the configuration is a 1000 Mb/s speed interface. You can also change the common speed of the port channel in this example, by setting the speed of the Gi 0/0 interface to 1000 Mb/s.

balancing

balancing may be applied to IPv4, switched IPv6, and non-IP traffic. For these traffic types, the IP-header-based hash and MAC-based hash may be applied to packets by using the following methods:

Table 13 Hash Methods as Applied to Port Channel Types

Hash (Header Based)	Layer 2 Port Channel	Layer 3 Port Channel
5-tuple	X	X
3-tuple	X	X
Packet-based	X	X
MAC source address (SA) and destination address (DA)	X	

The command -**balance**, allows the user to change the 5-tuple default to 3-tuple, MAC, or packet-based. To configure IP traffic, use the following command in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
[no] -balance [ip-selection {3-tuple | packet-based}] [mac]	CONFIGURATION	To designate a method to balance traffic over a port channel. By default, IP 5-tuple is used to distribute traffic over members port channel. **ip-selection 3-tuple**—to distribute IP traffic based on IP source address, IP destination address, and IP protocol type. **ip-selection packet-based**—to distribute IPV4 traffic based on the IP Identification field in the IPV4 header. **mac**—to distribute traffic based on the MAC source address, and the MAC destination address. See Table 15 for more information.

See also the command -**balance** in the *FTOS Command Line Interface Reference* guide.

To distribute IP traffic over a port channel member, FTOS uses the default IP 5-tuple hash. The 5-tuple and the 3-tuple hash use the following keys:

Table 14 5-tuple and 3-tuple Keys

Keys	5-tuple	3-tuple
IP source address (lower 32 bits)	X	X
IP destination address (lower 32 bits)	X	X
Protocol type	X	X
TCP/UDP source port	X	
TCP/UDP destination port	X	

Note: For IPV6, only the first 32 bits (LSB) of IP Source Address and IP Destination Address are used for hash generation.

IPv4, IPv6, and Non-IP Traffic Handling

The table below presents the different combinations of the **balance** command and its effect on the different port channel types.

Table 15 The -balance Commands and Port Channel Types

Configuration Commands	Switched IP Traffic	Routed IP Traffic (IPv4 only)	Switched Non-IP Traffic
Default (IP 5-tuple)	IP 5-tuple (lower 32 bits)	IP 5-tuple	MAC based
-**balance ip-selection 3-tuple**	IP 3-tuple (lower 32 bits)	IP 3-tuple	MAC based
-**balance ip-selection mac**	MAC based	IP 5-tuple	MAC based
-**balance ip-selection 3-tuple** -**balance ip-selection mac**	MAC based	IP 3-tuple	MAC based
-**balance ip-selection packet-based**	Packet based: IPV4 No distribution: IPV6	Packet based	MAC based
-**balance ip-selection packet-based** -**balance ip-selection mac**	MAC based	Packet based	MAC based

Configuration Task List for Port Channel Interfaces

To configure a port channel, you use the commands similar to those found in physical interfaces. By default, no port channels are configured in the startup configuration on the E-Series.

The following list includes the configuration tasks for port channel interfaces:

- create a port channel on page 216 (mandatory)
- add a physical interface to a port channel on page 216 (mandatory)
- change the criteria used to distribute traffic on page 219
- reassign an interface to a new port channel on page 220 (optional)
- configure the minimum oper up links in a port channel (LAG) on page 221 (optional)
- add or remove a port channel from a VLAN on page 221 (optional)
- assign an IP address to a port channel on page 222 (optional)
- delete or disable a port channel on page 222 (optional)

For a complete listing of all commands related to port channels and other interfaces, refer to .

create a port channel

You can create up to 255 port channels (LAGs) on an E-Series (255 for TeraScale, 1 to 32 for EtherScale). To create a port channel, you must be in the CONFIGURATION mode.

To configure a port channel, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**interface port-channel** *channel-number*	CONFIGURATION	Create a port channel.
2	**no shutdown**	INTERFACE PORT-CHANNEL	Ensure that the port channel is active.

The port channel is now enabled and you can place the port channel in Layer 2 or Layer 3 mode. Use the **switchport** command to place the port channel in Layer 2 mode or configure an IP address to place the port channel in Layer 3 mode.

You can configure a port channel as you would a physical interface by enabling or configuring protocols or assigning access control lists.

add a physical interface to a port channel

A port channel can contain up to 16 physical interfaces that are the same type. The physical interfaces in a port channel can be on any line card in the chassis, but must be the same physical type.

Note: Port channels can contain a mix of Gigabit Ethernet and 10/100/1000 Ethernet interfaces, but FTOS disables the interfaces that are not the same speed of the first channel member in the port channel (see 10/100/1000 interfaces in port channels on page 213).

You can add any physical interface to a port channel if the interface configuration is minimal. Only the following commands can be configured on an interface if it is a member of a port channel:

- **description**
- **shutdown**/**no shutdown**
- **mtu** (if the interface is on a Jumbo-enabled line card and the chassis is in Jumbo mode.)
- **ip mtu** (if the interface is on a Jumbo-enabled line card and the chassis is in Jumbo mode.)

To view the interface’s configuration, enter the INTERFACE mode for that interface and enter the **show config** command or from the EXEC privilege mode, enter the **show running-config interface** *interface* command.

When an interface is added to a port channel, FTOS recalculates the hash algorithm.

To add a physical interface to a port channel, use these commands in the following sequence in the INTERFACE mode of a port channel:

Step	Command Syntax	Command Mode	Purpose
1	**channel-member** *interface*	INTERFACE PORT-CHANNEL	Add the interface to a port channel. The *interface* variable is the physical interface type and slot/port information.
2	**show config**	INTERFACE PORT-CHANNEL	Double check that the interface was added to the port channel.

To view the port channel's status and channel members in a tabular format, use the **show interfaces port-channel brief** (Figure 109) command in the EXEC privilege mode.

```
Force10#show int port brief

LAG Mode  Status       Uptime    Ports
1   L2L3  up           00:06:03  Gi 13/6    (Up) *
                                 Gi 13/12   (Up)
2   L2L3  up           00:06:03  Gi 13/7    (Up) *
                                 Gi 13/8    (Up)
                                 Gi 13/13   (Up)
                                 Gi 13/14   (Up)
Force10#
```

Figure 109 show interfaces port-channel brief Command Example

Figure 109 displays the port channel's mode (L2 for Layer 2 and L3 for Layer 3 and L2L3 for a Layer 2 port channel assigned to a routed VLAN), the status, and the number of interfaces belonging to the port channel.

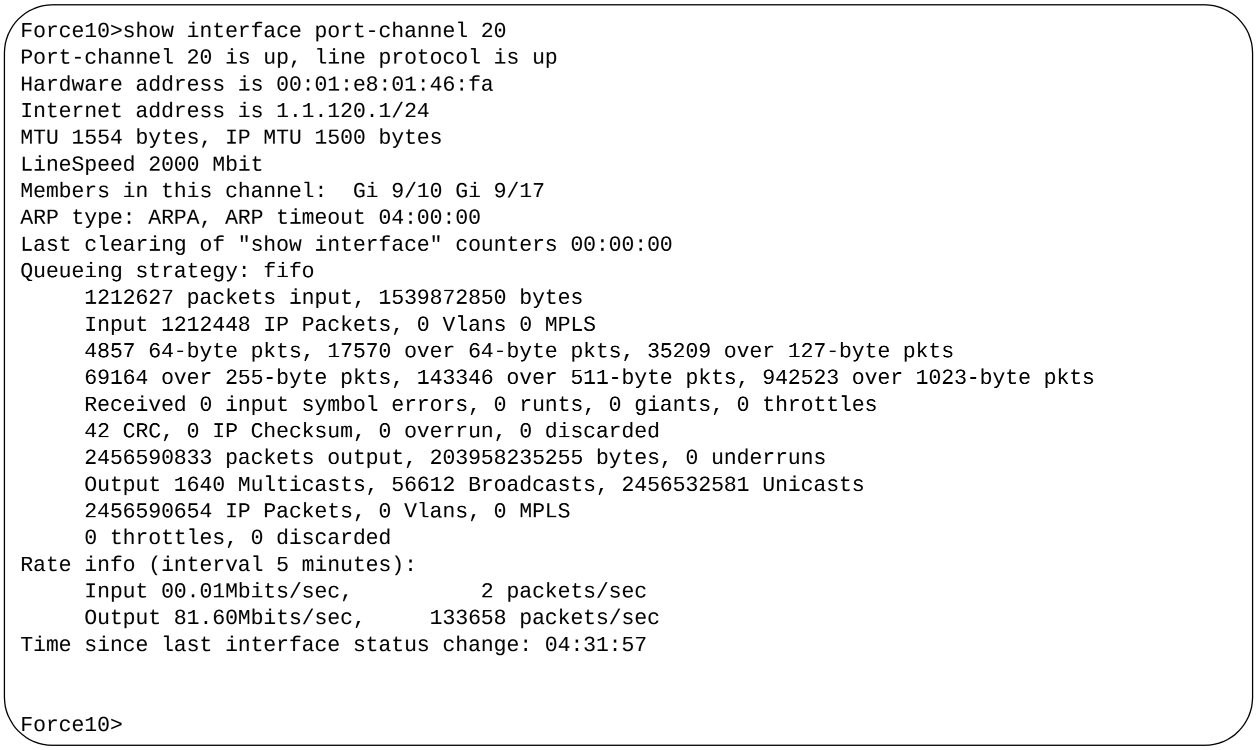

```
Force10>show interface port-channel 20
Port-channel 20 is up, line protocol is up
Hardware address is 00:01:e8:01:46:fa
Internet address is 1.1.120.1/24
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 2000 Mbit
Members in this channel:  Gi 9/10 Gi 9/17
ARP type: ARPA, ARP timeout 04:00:00
Last clearing of "show interface" counters 00:00:00
Queueing strategy: fifo
     1212627 packets input, 1539872850 bytes
     Input 1212448 IP Packets, 0 Vlans 0 MPLS
     4857 64-byte pkts, 17570 over 64-byte pkts, 35209 over 127-byte pkts
     69164 over 255-byte pkts, 143346 over 511-byte pkts, 942523 over 1023-byte pkts
     Received 0 input symbol errors, 0 runts, 0 giants, 0 throttles
     42 CRC, 0 IP Checksum, 0 overrun, 0 discarded
     2456590833 packets output, 203958235255 bytes, 0 underruns
     Output 1640 Multicasts, 56612 Broadcasts, 2456532581 Unicasts
     2456590654 IP Packets, 0 Vlans, 0 MPLS
     0 throttles, 0 discarded
Rate info (interval 5 minutes):
     Input 00.01Mbits/sec,          2 packets/sec
     Output 81.60Mbits/sec,     133658 packets/sec
Time since last interface status change: 04:31:57


Force10>
```

Figure 110   show interface port-channel Command Example

When more than one interface is added to a Layer 2 port channel, FTOS selects one of the active interfaces in the port channel to be the Primary Port. The primary port replies to flooding and sends protocol PDUs. An asterisk in the **show interfaces port-channel brief** command indicates the primary port.

As soon as a physical interface is added to a port channel, the properties of the port channel determine the properties of the physical interface. The configuration and status of the port channel are also applied to the physical interfaces within the port channel. For example, if the port channel is in Layer 2 mode, you cannot add an IP address or a static MAC address to an interface that is part of that port channel. As Figure 111 illustrates, interface GigabitEthernet 1/6 is part of port channel 5, which is in Layer 2 mode, and an error message appeared when an IP address was configured.

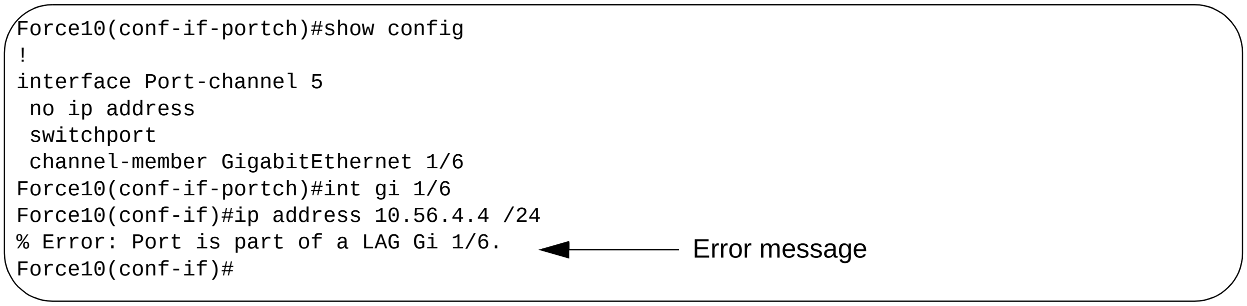

```
Force10(conf-if-portch)#show config
!
interface Port-channel 5
 no ip address
 switchport
 channel-member GigabitEthernet 1/6
Force10(conf-if-portch)#int gi 1/6
Force10(conf-if)#ip address 10.56.4.4 /24
% Error: Port is part of a LAG Gi 1/6.
Force10(conf-if)#
```

Figure 111   Error Message

change the criteria used to distribute traffic

By default, FTOS use a 5-tuple IP selection to distribute traffic over channel members in a port channel. The default criteria is as follows:

- IP source address
- IP destination address
- Protocol type
- TCP/UDP source port
- TCP/UDP destination port

To change the criteria, use the following command in the INTERFACE (Port Channel) mode:

Command Syntax	Command Mode	Purpose
load-balance	CONFIGURATION	Change the traffic

E-Series uses one of 47 possible hash algorithms (16 on EtherScale).

On C-Series, the default ECMP hash configuration is crc-lower. This takes the lower 32 bits of the hash key to compute the egress port. Other options for ECMP hash-algorithms are **crc-upper**, **lsb**, and **dest-ip**.

- **crc-upper** uses the upper 32 bits of the hash key to compute the egress port.
- **lsb** always uses the the least significant bit of the hash key to compute the egress port.
- **dest-ip** uses destination IP address as part of the hash key.

For LAG hashing on C-Series, the source IP, destination IP, source TCP/UDP port and destination TCP/UDP port for hash computation by default. For packets without an Layer 3 header, FTOS automatically uses **load-balance mac source-dest-mac**.

IP hashing or MAC hashing should not be configured at the same time. If you configure an IP and MAC hashing scheme at the same time, the MAC hashing scheme takes precedence over the IP hashing scheme.

To change one of the other, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
hash-algorithm *number*	CONFIGURATION	Change to another algorithm.

reassign an interface to a new port channel

An interface can be a member of only one port channel. If the interface is a member of a port channel, you must remove it from the first port channel and then add it to the second port channel.

Each time you add or remove a channel member from a port channel, FTOS recalculates the hash algorithm for the port channel.

To reassign an interface to a new port channel, use these commands in the following sequence in the INTERFACE mode of a port channel:

Step	Command Syntax	Command Mode	Purpose
1	**no channel-member** *interface*	INTERFACE PORT-CHANNEL	Remove the interface from the first port channel.
2	**interface port-channel** *number*	INTERFACE PORT-CHANNEL	Change to the second port channel INTERFACE mode.
3	**channel-member** *interface*	INTERFACE PORT-CHANNEL	Add the interface to the second port channel.

Figure 112 displays an example of moving the GigabitEthernet 1/8 interface from port channel 4 to port channel 3.

```
Force10(conf-if-portch)#show config
!
interface Port-channel 4
 no ip address
 channel-member GigabitEthernet 1/8
 no shutdown
Force10(conf-if-portch)#no chann gi 1/8
Force10(conf-if-portch)#int port 5
Force10(conf-if-portch)#channel gi 1/8
Force10(conf-if-portch)#sho conf
!
interface Port-channel 5
 no ip address
 channel-member GigabitEthernet 1/8
 shutdown
Force10(conf-if-portch)#
```

Figure 112 Command Example from Reassigning an Interface to a Different Port Channel

configure the minimum oper up links in a port channel (LAG)

You can configure the minimum links in a port channel (LAG) that must be in “oper up” status for the port channel to be considered to be in “oper up” status. Use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
minimum-links *number*	INTERFACE	Enter the number of links in a LAG that must be in “oper up” status. Range: 1 to 16 Default: 1

Figure 113 displays an example of configuring five minimum “oper up” links in a port channel.

```
Force10#config t
Force10(conf)#int po 1
Force10(conf-if-po-1)#minimum-links 5
Force10(conf-if-po-1)#
```

Figure 113 Example of using the minimum-links Command

add or remove a port channel from a VLAN

As with other interfaces, you can add Layer 2 port channel interfaces to VLANs. To add a port channel to a VLAN, you must place the port channel in Layer 2 mode (by using the **switchport** command).

To add a port channel to a VLAN, use either of the following commands in the INTERFACE mode of a VLAN:

Command Syntax	Command Mode	Purpose
tagged port-channel *number*	INTERFACE VLAN	Add the port channel to the VLAN as a tagged interface. An interface with tagging enabled can belong to multiple VLANs.
untagged port-channel *number*	INTERFACE VLAN	Add the port channel to the VLAN as an untagged interface. An interface without tagging enabled can belong to only one VLAN.

To remove a port channel from a VLAN, use either of the following commands in the INTERFACE mode of a VLAN:

Command Syntax	Command Mode	Purpose
no tagged port-channel *number*	INTERFACE VLAN	Remove the port channel with tagging enabled from the VLAN.

Command Syntax	Command Mode	Purpose
no untagged port-channel *number*	INTERFACE VLAN	Remove the port channel without tagging enabled from the VLAN.

To see which port channels are members of VLANs, enter the **show vlan** command in the EXEC privilege mode.

assign an IP address to a port channel

You can assign an IP address to a port channel and use port channels in Layer 3 routing protocols.

To assign an IP address, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip address *ip-address mask* [**secondary**]	INTERFACE	Configure an IP address and mask on the interface. • *ip-address mask:* enter an address in dotted-decimal format (A.B.C.D) and the mask must be in slash format (/24). • **secondary:** the IP address is the interface's backup IP address. You can configure up to eight secondary IP addresses.

delete or disable a port channel

To delete a port channel, you must be in the CONFIGURATION mode and use the **no interface portchannel** *channel-number* command.

When you disable a port channel (using the **shutdown** command) all interfaces within the port channel are operationally down also.

VLAN Interfaces and Layer 3

VLANs are logical interfaces and are, by default, in Layer 2 mode. Physical interfaces and port channels can be members of VLANs. For more information on VLANs and Layer 2, refer to Chapter 7, Layer 2, on page 161.

➡ **Note:** To monitor VLAN interfaces, use the Management Information Base for Network Management of TCP/IP-based internets: MIB-II (RFC 1213).

FTOS supports Inter-VLAN routing (Layer 3 routing in VLANs). You can add IP addresses to VLANs and use them in routing protocols in the same manner that physical interfaces are used. For more information on configuring different routing protocols, refer to the chapters on the specific protocol.

A consideration for including VLANs in routing protocols is that the **no shutdown** command must be configured. (For routing traffic to flow, the VLAN must be enabled.)

To assign an IP address, use the following command in the VLAN INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip address *ip-address mask* [**secondary**]	INTERFACE	Configure an IP address and mask on the interface. • *ip-address mask:* enter an address in dotted-decimal format (A.B.C.D) and the mask must be in slash format (/24). • **secondary:** the IP address is the interface's backup IP address. You can configure up to eight secondary IP addresses.

Figure 114 shows a sample configuration of a VLAN participating in an OSPF process.

```
interface Vlan 10
 ip address 1.1.1.2/24
 tagged GigabitEthernet 2/2-13
 tagged TenGigabitEthernet 5/0
 tagged SONET 12/0
 ip ospf authentication-key force10
 ip ospf cost 1
 ip ospf dead-interval 60
 ip ospf hello-interval 15
 no shutdown
!
```

Figure 114 Sample Layer 3 Configuration of a VLAN

Bulk Configuration

Bulk configuration enables you to determine if interfaces are present, for physical interfaces, or, configured, for logical interfaces.

Interface Range

An interface range is a set of interfaces to which other commands may be applied, and may be created if there is at least one valid interface within the range. Bulk configuration excludes from configuration any non-existing interfaces from an interface range. A default VLAN may be configured only if the interface range being configured consists of only VLAN ports.

The **interface range** command allows you to create an interface range allowing other commands to be applied to that range of interfaces.

The interface range prompt offers the interface (with slot and port information) for valid interfaces. The maximum size of an interface range prompt is 32. If the prompt size exceeds this maximum, it displays (...) at the end of the output.

Note: Non-existing interfaces are excluded from interface range prompt. In the following example, Tengigabit 3/0 and VLAN 1000 do not exist.

Note: When creating an interface range, interfaces appear in the order they were entered and are not sorted.

The **show range** command is available under interface range mode. This command allows you to display all interfaces that have been validated under the interface range context.

The **show configuration** command is also available under the interface range mode. This command allows you to display the running configuration only for interfaces that are part of interface range.

Bulk Configuration Examples

The following are examples of using the **interface range** command for bulk configuration:

- creating a single-range on page 224
- creating a multiple-range on page 224
- duplicate entries on page 225
- excluding a smaller port range on page 225
- overlapping port ranges on page 225
- using commas on page 225
- adding ranges on page 226

creating a single-range

```
Force10(config)# interface range gigabitethernet 5/1 - 23
Force10(config-if-range-gi-5/1-23)# no shutdown
Force10(config-if-range-gi-5/1-23)#
```

Figure 115 Creating a Single-Range Bulk Configuration

creating a multiple-range

```
Force10(conf)#interface range tengigabitethernet 3/0 , gigabitethernet 2/1 - 47 , vlan 1000 , sonet 5/0
Force10(conf-if-range-gi-2/1-47,so-5/0)#
```

Figure 116 Creating a Multiple-Range Prompt

duplicate entries

Duplicate single interfaces and port ranges are excluded from the resulting interface range prompt:

```
Force10(conf)#interface range vlan 1 , vlan 1 , vlan 3 , vlan 3
Force10(conf-if-range-vl-1,vl-3)#
Force10(conf)#interface range gigabitethernet 2/0 - 23 , gigabitethernet 2/0 - 23 , gigab 2/0 - 23
Force10(conf-if-range-gi-2/0-23)#
```

Figure 117 Interface Range Prompt Excluding Duplicate Entries

excluding a smaller port range

If interface range has multiple port ranges, the smaller port range is excluded from prompt:

```
Force10(conf)#interface range gigabitethernet 2/0 - 23 , gigab 2/1 - 10
Force10(conf-if-range-gi-2/0-23)#
```

Figure 118 Interface Range Prompt Excluding a Smaller Port Range

overlapping port ranges

If overlapping port ranges are specified, the port range is extended to the smallest start port number and largest end port number:

```
Force10(conf)#inte ra gi 2/1 - 11 , gi 2/1 - 23
Force10(conf-if-range-gi-2/1-23)#
```

Figure 119 Interface Range Prompt Including Overlapping Port Ranges

using commas

The example below shows how to use commas to add different interface types to the range, enabling all Gigabit Ethernet interfaces in the range 5/1 to 5/23 and both Ten Gigabit Ethernet interfaces 1/1 and 1/2.

```
Force10(config-if)# interface range gigabitethernet 5/1 - 23, tengigabitethernet 1/1 - 2
Force10(config-if-range-gi-5/1-23)# no shutdown
Force10(config-if-range-gi-5/1-23)#
```

Figure 120 Multiple-Range Bulk Configuration Gigabit Ethernet and Ten-Gigabit Ethernet

adding ranges

The example below shows how to use commas to add SONET, VLAN, and port-channel interfaces to the range.

```
Force10(config-ifrange-gi-5/1-23-te-1/1-2)# interface range Sonet 6/1 – 10 , Vlan 2 – 100 , Port 1 – 25
Force10(config-if-range-gi-5/1-23-te-1/1-2-so-5/1-vl-2-100-po-1-25)# no shutdown
Force10(config-if-range)#
```

Figure 121 Multiple-Range Bulk Configuration with SONET, VLAN, and Port-channel

Time Domain Reflectometry

C-Series	NO
E-Series	✓

Platform Specific Feature: Time Domain Reflectometry is supported on E-Series only.

The Time Domain Reflectometer (TDR) is an assistance tool to resolve link issues that helps detect obvious open or short conditions within any of the four copper pairs. TDR sends a signal onto the physical cable and examines the reflection of the signal that returns. By examining the reflection, TDR is able to indicate whether there is a cable fault (when the cable is broken, becomes un-terminated, or if a transceiver is unplugged).

TDR is useful for troubleshooting an interface that is not establishing a link, that is, when the link is flapping or not coming up. TDR is not intended to be used on an interface that is passing traffic. When a TDR test is run on a physical cable, it is important to shut down the port on the far end of the cable otherwise it may lead to incorrect test results.

Note: TDR is an intrusive test. Do not run TDR on a link that is up and passing traffic.

To test the condition of cables on 10/100/1000 BASE-T modules, use the **tdr-cable-test** command:

Step	Command Syntax	Command Mode	Usage
1	**tdr-cable-test gigabitethernet** *<slot>/<port>*	EXEC privilege	To test for cable faults on the GigabitEthernet cable. • Between two ports, you must not start the test on both ends of the cable. • You must enable the interface before starting the test. • The port should be enabled to run the test or the test prints an error message.

Step	Command Syntax	Command Mode	Usage
2	**show tdr gigabitethernet** *<slot>/<port>*	EXEC privilege	Displays TDR test results.

Chapter 10 SONET

C-Series	NO
E-Series	✓

Platform Specific Feature: SONET is supported on E-Series only.

FTOS supports RFC 2558 “Definitions of Managed Objects for the SONET/SDH Interface” and RFC 2615 “PPP-over-SONET/SDH.” FTOS supports two line cards with SONET—Packet-Over-SONET (POS) and 10GE WAN PHY.

Important Points to Remember—POS

Force10 Network’s Packet-Over-SONET line card does not support:

- the E-300 chassis
- S0S1
- Layer 2
- VRRP
- IPv6
- LAG
- APS and protection switching
- POS (Packet over SONET) ports can not be mirrored ports
- SONET alarm reporting cannot be disabled

10GE WAN PHY Interface

10 GE interfaces support LAN and WAN modes. When in WAN mode, the 10 GE interface operates as a SONET interface. Use the **wanport** command in the INTERFACE mode to configure a 10 GE interface to be in WAN mode.

Note that the port must be in shutdown state before the **wanport** command can be executed successfully (Figure 122).

Step	Task	Command Syntax	Command Mode
1	Place the port in shutdown state	**shutdown**	INTERFACE
2	Place the port in WAN mode	**wanport**	INTERFACE
3	Display the active/defective alarms	**show controllers tengigabitethernet** *slot/port*	EXEC

Figure 122 wanport command example

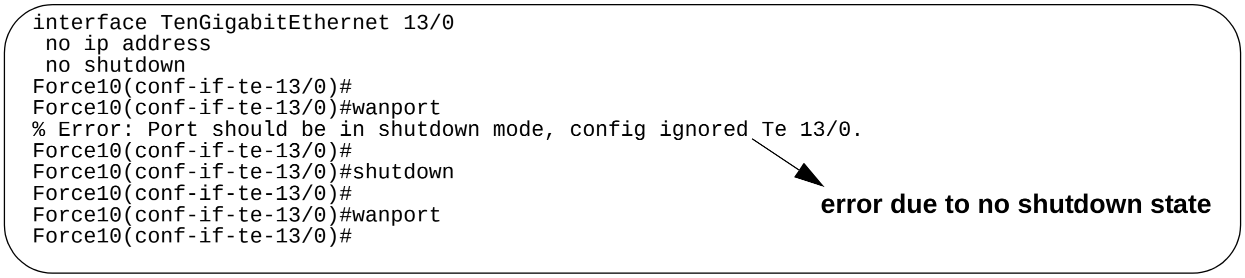

```
interface TenGigabitEthernet 13/0
 no ip address
 no shutdown
Force10(conf-if-te-13/0)#
Force10(conf-if-te-13/0)#wanport
% Error: Port should be in shutdown mode, config ignored Te 13/0.
Force10(conf-if-te-13/0)#
Force10(conf-if-te-13/0)#shutdown
Force10(conf-if-te-13/0)#
Force10(conf-if-te-13/0)#wanport
Force10(conf-if-te-13/0)#
```

error due to no shutdown state

Figure 123 displays the active alarms for the interface.

Figure 123 show controllers tengigabitEthernet command example

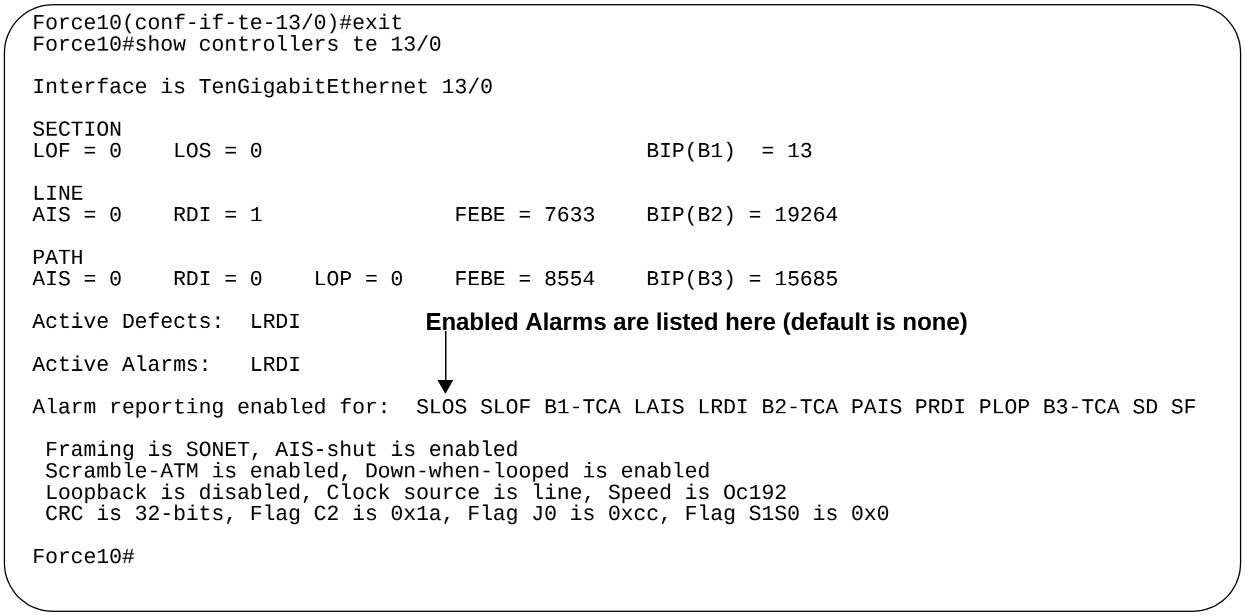

```
Force10(conf-if-te-13/0)#exit
Force10#show controllers te 13/0

Interface is TenGigabitEthernet 13/0

SECTION
LOF = 0    LOS = 0                             BIP(B1)  = 13

LINE
AIS = 0    RDI = 1               FEBE = 7633    BIP(B2) = 19264

PATH
AIS = 0    RDI = 0    LOP = 0    FEBE = 8554    BIP(B3) = 15685

Active Defects:  LRDI

Active Alarms:   LRDI

Alarm reporting enabled for:  SLOS SLOF B1-TCA LAIS LRDI B2-TCA PAIS PRDI PLOP B3-TCA SD SF

 Framing is SONET, AIS-shut is enabled
 Scramble-ATM is enabled, Down-when-looped is enabled
 Loopback is disabled, Clock source is line, Speed is Oc192
 CRC is 32-bits, Flag C2 is 0x1a, Flag J0 is 0xcc, Flag S1S0 is 0x0

Force10#
```

Enabled Alarms are listed here (default is none)

Alarm Reporting

SONET equipment detects events and alarms at each of SONET's three layers—section, line, and path. Typically, a SONET device sends alarms both upstream and downstream to notify other devices of the problem condition. The GR-253-CORE Synchronous Optical Network (SONET) Transport Systems Common Generic Criteria specification defines several alarms:

- Section Loss of Signal (SLOS)
- Section Loss of Frame (SLOF)
- Alarm Indication Signal - Line (AIS-L)
- Signal Degrade Bit Error Rate (SD-BER)
- Signal Failure Bit Error Rate (SF-BER)
- Remote Defect Indication - Line (RDI-L)

The 10 GE WAN PHY and POS interfaces support the following alarms:

- Section alarms—SLOS, SLOF, B1
- Line alarms—AIS, RDI, FEBE(REI), B2
- Path Alarms—AIS, RDI, FEBE(REI), B3

While performance monitoring provides advanced alert of link degradation, alarms indicate a failure. Fault management involves alarm monitoring and generation, reporting, logging, correlation, and clearing. Since E-Series is Terminal Equipment (TE), it must support the alarms in Table 16.

Table 16 Alarm Definitions

SONET/SDH Layer	Alarm	Description
Section/Regenerator	LOF	Loss of Frame condition—when a severely errored frame (SEF) defect on the incoming SONET signal and persists for 3 milliseconds
	LOS	Loss of Sync condition—when an all-zero pattern on the incoming SONET signal last 19 (+/-3) microseconds or longer. This defect might also be reported if the received signal level drops below the specified threshold.

Table 16 Alarm Definitions

SONET/SDH Layer	Alarm	Description
Line/Multiplexing	AIS	Line Alarm Indication Signal is sent by the section terminating equipment (STE) to alert the downstream line terminating equipment (LTE) that a LOS or LOF defect has been detected on the incoming SONET section.
	RDI	Line Remote Defect Indication is reported by the downstream LTE when it detects LOF, LOS, or AIS.
	FEBE	Line Far End Block Errors (accumulated from the M0 or M1 byte) is reported when the downstream LTE detects BIP (B2) errors.
	SD	Signal Degrade is sourced from B2 BIP (BER)
	SF	Signal Failure is sourced from B2 BIP (BER)
Path/Section	AIS	Path Alarm Indication Signal is sent by the LTE to alert the downstream path terminating equipment (PTE) that it has detected a defect on its incoming line signal.
	RDI	Path Remote Defect Indication is reported by the downstream PTE when it detects a defect on the incoming signal.
	FEBE	Path Far End Block Errors (accumulated from G1 byte) is reported when the downstream PTE detects BIP (B3) errors.
	LOP	Loss of pointer is a result of an invalid pointer (H1,H2) or an excess number of new data flag (NDF) enable indications.

Use the **alarm-report** command to configure the alarms that the 10 GE WAN or POS interface can activate.

Task	Command Syntax	Command Mode
Specify which POS/SDH alarms to report to the remote SNMP server.	**alarm-report** {**b1-tca** | **b2-tca** | **b3-tca** | **lais** | **lrdi** | **pais** | **plop** | **prdi** | **sd-ber** | **sf-ber** | **slof** | **slos**}	INTERFACE

SNMP traps are available; however, syslogs are not generated. To view active alarms and defects, use the **show controllers** command. Table 17 defines the alarms that can be enabled by this command. The first three alarms (b1, b2, b3) do not generate reports. The remaining alarms, if enabled for reporting, will generate reports on a trap receiver.

Table 17 Alarm Definitions

Alarm	Description
b1-tca	**B1 BER threshold crossing alarm**
b2-tca	**B2 BER Threshold crossing alarm**
b3-tca	**B3 BER threshold crossing alarm**
lais	**Line Alarm Indication Signal**
lrdi	**Line Remote Defect Indication**
pais	**Path Alarm Indication Signal**
plop	**Path loss of Pointer**
prdi	**Path Remote Defect Indication**
sd-ber	**LBIP BER in excess of Signal Degradation threshold**
sf-ber	**LBIP BER in excess of Signal Failure threshold**
slof	**Section Loss of Frame**
slos	**Section Loss of Signal**

Events that Bring Down a SONET Interface

You can configure the SONET interface to change to a "down state" when certain SONET events are reported. When the event (or trigger) occurs, FTOS brings down the SONET interface. You can use the **delay triggers** command to indicate a 100ms delay in bringing down the SONET interface once the event or trigger is detected.

Task	Command Syntax	Command Mode
Delay triggering the line or path alarms with a 100ms delay.	**delay triggers** {**line** [**lrdi** | **sd-ber** | **sf-ber**] | **path** [**pais** | **prdi**] }	INTERFACE

By default, certain alarms (LOS, LOF, LAIS, PLOP) bring the line protocol down immediately. Use this command, with the **line** option, to delay that trigger event by 100ms.

By default, path alarms (AIS, RDI, LOP) *do not* cause (or trigger) the interface line protocol to go down. This command, with the **path** option, can be used to trigger this action with a delay of 100ms.

SONET MIB

Table 18 lists the supported OIDs of the SONET MIB, as defined in RFC 2558.

Table 18 SONET MIBs

SONET MIB	Description
sonetMediumType	Sonet or SDH depending on the configuration
sonetMediumTimeElapsed	Time in seconds (up to 900 seconds) since the line card is up. Resets after 900 seconds has elapsed
sonetMediumValidIntervals	The number of previous intervals for which valid data has been stored.
sonetMediumLineCoding	This variable describes the line coding for this interface—Non-Return to Zero (NRZ).
sonetMediumCircuitIdentifier	This variable contains the transmission vendor's circuit identifier to facilitate troubleshooting. Note that the circuit identifier, if available, is also represented by ifPhysAddress.
sonetMediumInvalidIntervals	Displays seconds in current 15 minute intervals when data could not be collected.
sonetMediumLoopbackConfig	Displays if loopback is line or internal
sonetSESthresholdSet	Displays which recognized set of SES thresholds is supported.

SONET Traps

This is a Force10-specific enterprise MIB.

Table 19 SONET Traps and OID

Trap	OID	Trap Object
SONET_S_LOS **Section Loss of Signal**	**1.3.6.1.4.1.6027.3.3.2.2.0.1**	**alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)**
SONET_S_LOF **Section Loss of Frame**	**1.3.6.1.4.1.6027.3.3.2.2.0.2**	**alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)**
SONET_S_B1TCA **Section B1 Threshold Crossing Alert**	**1.3.6.1.4.1.6027.3.3.2.2.0.3**	**alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)**

Table 19 SONET Traps and OID (continued)

Trap	OID	Trap Object
SONET_L_AIS **Line Alarm Indication Signal**	**1.3.6.1.4.1.6027.3.3.2.2.0.9**	**alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6**
SONET_L_RDI **Line Remote Defect Indication**	**1.3.6.1.4.1.6027.3.3.2.2.0.10**	**alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)**
SONET_L_FEBE **Line Far-end Background Block Errors**	**1.3.6.1.4.1.6027.3.3.2.2.0.11**	**alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)**
SONET_L_B2TCA **Line B2 Threshold Crossing Alert**	**1.3.6.1.4.1.6027.3.3.2.2.0.12**	**alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)**
SONET_P_AIS **Path Alarm Indication Signal**	**1.3.6.1.4.1.6027.3.3.2.2.0.17**	**alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)**
SONET_P_RDI **Path Remote Defect Indication**	**1.3.6.1.4.1.6027.3.3.2.2.0.18**	**alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)**
SONET_P_FEBE **Path Far-end Background Block Errors**	**1.3.6.1.4.1.6027.3.3.2.2.0.19**	**alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)**
SONET_P_LOP **Path Loss of Pointer**	**1.3.6.1.4.1.6027.3.3.2.2.0.20**	**alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)**

Table 19 SONET Traps and OID (continued)

Trap	OID	Trap Object
SONET_P_NEWPTR Path New Pointer	1.3.6.1.4.1.6027.3.3.2.2.0.21	alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)
SONET_P_PSE	1.3.6.1.4.1.6027.3.3.2.2.0.22	alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)
SONET_P_NSE	1.3.6.1.4.1.6027.3.3.2.2.0.23	alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)
SONET_P_B3TCA Line B3 Threshold Crossing Alert	1.3.6.1.4.1.6027.3.3.2.2.0.24	alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)
SONET_SD_BER Signal Degrade Bit Error Rate	1.3.6.1.4.1.6027.3.3.2.2.0.27	alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)
SONET_SF_BER Signal Failure Bit Error Rate	1.3.6.1.4.1.6027.3.3.2.2.0.28	alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)
SONET_LOC Loss of Cell Delineation	1.3.6.1.4.1.6027.3.3.2.2.0.29	alarm state (1.3.6.1.4.1.6027.3.3.1.2.1.1.3), alarm type(1.3.6.1.4.1.6027.3.3.1.2.1.1.2), ifindex(1.3.6.1.4.1.6027.3.3.1.2.1.1.4), slot(1.3.6.1.4.1.6027.3.3.1.2.1.1.5), port(1.3.6.1.4.1.6027.3.3.1.2.1.1.6)

Packet Over SONET Interfaces

Packet Over SONET interfaces require some configuration considerations. When you remove encapsulation (**no encap** command) on a SONET interface command, you administratively shutdown the interface and configuration information (such as IP address) is deleted from the interface. A SONET interface without encapsulation is always operationally down.

When you enable encapsulation on the interface (**encap hdlc/ppp**), PPP negotiation begins after you enable the interface (**no shutdown** command). You can enable authentication and other related commands once negotiation is completed.

Note: Encapsulation must be configured before the interface is enabled for traffic.

Different equipment vendors have set different defaults for PPP encapsulation; therefore when you configure the E-Series to use PPP encapsulation between the E-Series and another vendor's equipment, verify the following settings:

- Set one side of the link to **clock source internal**.
- E-Series's SONET interface defaults are ATM scrambling disabled; flag is c2 0xCF (207) and j0 is 0xCF (207).
- FTOS supports Challenge-Handshake Authentication Protocol (CHAP) and/or Password Authentication Protocol (PAP) authentication.
- Confirm that the MTU settings are the same on both end of the link. If you configure the **ip mtu** command with a different value on the far end of the link, the interface on the E-Series goes down.

Note: SONET uses synchronous transport signal (STS) framing. When framing is configured on an interface, it should only be done when the interface is shut down.

Chapter 11 sFlow

C-Series	NO
E-Series	✓

Platform Specific Feature: sFlow is supported on E-Series only.

This chapter contains information on configuring sFlow globally, on an interface, and on a line card using FTOS on E-Series.

Important Points to Remember

- Force10 Networks recommends that the sFlow Collector be connected to the Force10 chassis through a line card port rather than the RPM Management Ethernet port.
- Community list and local preference fields are not filled up in extended gateway element in sFlow datagram.
- 802.1P source priority field is not filled up in extended switch element in sFlow datagram.
- Only Destination and Destination Peer AS number are packed in the dst-as-path field in extended gateway element
- If packet being sampled is redirected using PBR (Policy-Based Routing), sFlow datagram may contain incorrect extended gateway/router information.
- sFlow does not support packing extended information for IPv6 packets. Only the first 128 bytes of the IPv6 packet is shipped in the datagram.
- Source VLAN field in the extended switch element will not be packed in case of routed packet.
- Destination VLAN field in the extended switch element will not be packed in case of Multicast packet.

- The maximum number of packets that can be sampled and processed per second is:
  - 7500 packets when no extended information packing is enabled
  - 7500 packets when only extended-switch information packing is enabled.
  - 1600 packets when extended-router and/or extended-gateway information packing is enabled

FTOS supports sFlow version 5. sFlow is a standard-based sampling technology embedded within switches and routers which is used to monitor network traffic. It is designed to provide traffic monitoring for high speed networks with many switches and routers. sFlow uses two types of sampling:

- Statistical packet-based sampling of switched or routed packet flows
- Time-based sampling of interface counters

The sFlow monitoring system consists of an sFlow Agent (embedded in the switch/router) and an sFlow collector. The sFlow Agent resides anywhere within the path of the packet, and combines the flow samples and interface counters into sFlow datagrams and forwards them to the sFlow Collector at regular intervals. The datagrams consists of information on, but not limited to, packet header, ingress and egress interfaces, sampling parameters, and interface counters.

Packet sampling is typically done by the ASIC. sFlow Collector analyses the sFlow datagrams received from different devices and produces a network-wide view of traffic flows.

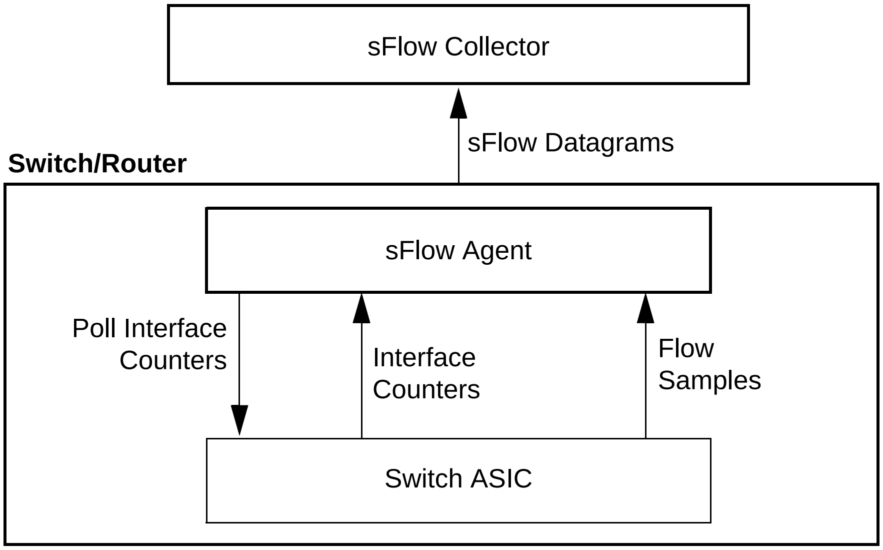

Figure 124 sFlow Traffic Monitoring System

Enabling and Disabling sFlow

By default, sFlow is *disabled* globally on the system. To enable sFlow globally, use the **sflow enable** command in CONFIGURATION mode. Use the **no** version of this command to disable sFlow globally.

Command Syntax	Command Mode	Usage
[**no**] **sflow enable**	CONFIGURATION	Enables sFlow globally.

Enabling and Disabling on an Interface

By default, sFlow is *disabled* on all interfaces. To enable sFlow on a specific interface, use the **sflow enable** command in INTERFACE mode. Use the **no** version of this command to disable sFlow on an interface. This CLI is supported on physical ports and LAG ports.

Command Syntax	Command Mode	Usage
[**no**] **sflow enable**	INTERFACE	This enables sFlow on an interface.

sFlow Show Commands

FTOS includes the following sFlow display commands:

- Show sFlow Globally on page 242
- Show sFlow on an Interface on page 243
- Show sFlow on a Line Card on page 243

Show sFlow Globally

Use the following command to view sFlow statistics:

Command Syntax	Command Mode	Purpose
show sflow	EXEC	Displays sFlow configuration information and statistics.

Figure 125 is a sample output from the **show sflow** command:

```
Force10#show sflow
sFlow services are enabled
Global default sampling rate: 32768
Global default counter polling interval: 20
1 collectors configured
Collector IP addr: 133.33.33.53, Agent IP addr: 133.33.33.116, UDP port: 6343
77 UDP packets exported
0 UDP packets dropped
165 sFlow samples collected
69 sFlow samples dropped due to sub-sampling

Linecard 1 Port set 0 H/W sampling rate 8192
  Gi 1/16: configured rate 8192, actual rate 8192, sub-sampling rate 1
  Gi 1/17: configured rate 16384, actual rate 16384, sub-sampling rate 2

Linecard 3 Port set 1 H/W sampling rate 16384
  Gi 3/40: configured rate 16384, actual rate 16384, sub-sampling rate 1
```

Figure 125 show sflow Command Example

Show sFlow on an Interface

Use the following command to view sFlow information on a specific interface:

Command Syntax	Command Mode	Purpose
show sflow interface *interface-name*	EXEC	Displays sFlow configuration information and statistics on a specific interface.

Figure 126 is a sample output from the **show sflow interface** command:

```
Force10#show sflow interface gigabitethernet 1/16
Gi 1/16
Configured sampling rate         :8192
Actual sampling rate             :8192
Sub-sampling rate                :2
Counter polling interval         :15
Samples rcvd from h/w            :33
Samples dropped for sub-sampling :6
```

Figure 126 show sflow interface Command Example

The configuration, shown in Figure 126, is also displayed in the running configuration (Figure 127):

```
Force10#show running-config interface gigabitethernet 1/16
!
interface GigabitEthernet 1/16
 no ip address
 mtu 9252
 ip mtu 9234
 switchport
 sflow enable
 sflow sample-rate 8192
 no shutdown
```

Figure 127 show running-config interface *interface* Command Example

Show sFlow on a Line Card

Use the following command to view sFlow statistitics on a specified line card:

Command Syntax	Command Mode	Purpose
show sflow linecard *slot-number*	EXEC	Displays sFlow configuration information and statistics on the specified interface.

Figure 128 is a sample output from the **show sflow linecard** command:

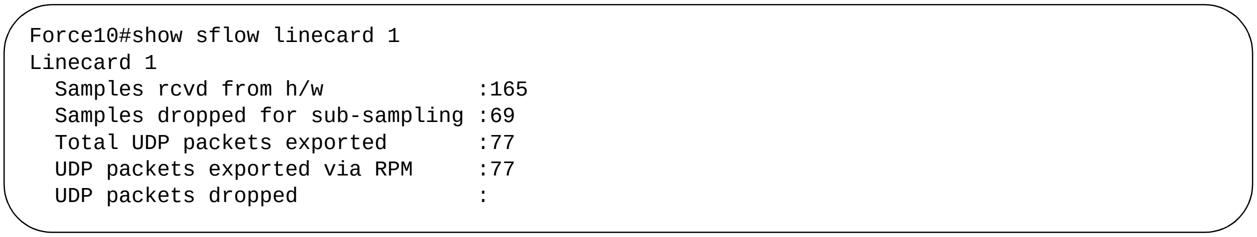

```
Force10#show sflow linecard 1
Linecard 1
  Samples rcvd from h/w            :165
  Samples dropped for sub-sampling :69
  Total UDP packets exported       :77
  UDP packets exported via RPM     :77
  UDP packets dropped              :
```

Figure 128 show sflow linecard Command Example

Specifying Collectors

The **sflow collector** command allows identification of sFlow Collectors to which sFlow datagrams are forwarded. The user can specify up to two sFlow collectors. If two Collectors are specified, the samples are sent to both.

Command Syntax	Command Mode	Usage
sflow collector *ip-address* **agent-addr** *ip-address* [*number* [**max-datagram-size** *number*]] | [**max-datagram-size** *number*]	CONFIGURATION	Allows identification of sFlow collectors to which sFlow datagrams are forwarded. Default UDP port: 6343 Default max-datagram-size: 1400

Polling Intervals

The **sflow polling-interval** command configures the polling interval for an interface in the maximum number of seconds between successive samples of counters to be sent to the collector. This command changes the global default counter polling (20 seconds) interval. You can configure an interface to use a different polling interval.

The polling interval can be configured globally (in CONFIGURATION mode) or by interface (in INTERFACE mode) by executing the interval command:

Command Syntax	Command Mode	Usage
sflow polling-interval *interval value*	CONFIGURATION or INTERFACE	Changes the global default counter polling interval. *interval value*—in seconds. Range: 15 to 86400 seconds Default: 20 seconds

Sampling Rate

The sFlow sampling rate is the number of packets that are skipped before the next sample is taken. sFlow does not have time-based packet sampling.

The **sflow sample-rate** command, when issued in CONFIGURATION mode, changes the default sampling rate. By default, the sampling rate of an interface is set to the same value as the current global default sampling rate.If the value entered is not a correct power of 2, the command generates an error message with the previous and next power-of-2 value. Select one of these two number and re-enter the command. (For more information on values in power-of-2, see Sub-sampling on page 245.)

The sample rate can be configured globally or by interface using the sample rate command:

Command Syntax	Command Mode	Usage
[**no**] **sflow sample-rate** *sample-rate*	CONFIGURATION or INTERFACE	To change the global or interface sampling rate. Sample-rate value: Range 0-8388608 (only in powers of two, for example, 4096, 8192, 16384) Configuration Default value: 32768 Interface Default value: The global sampling rate

Sub-sampling

The sFlow sample rate is not the frequency of sampling, but the number of packets that are skipped before the next sample is taken. Although a sampling rate can be configured for each port, TeraScale line cards can support only a single sampling rate per port pipe.

Therefore, sFlow Agent uses sub-sampling to create multiple sampling rates per port pipe. To achieve different sampling rates for different ports in a port-pipe, sFlow Agent takes the lowest numerical value of the sampling rate of all the ports within the port pipe, and configures all ports to this value. sFlow Agent is then able to skip samples on ports where you require a larger sampling rate value.

Sampling rates are configurable in powers of two. This allows the smallest sampling rate possible to be configured on the hardware, and also allows all other sampling rates to be available through sub-sampling.

For example, if Gig 1/0 and 1/1 are in a port pipe, and they are configured with a sampling rate of 4096 on interface Gig 1/0, and 8192 on Gig 1/1, sFlow Agent does the following:

1. Configures the hardware to a sampling rate of 4096 for all ports with sFlow enabled on that port pipe.
2. Configure interface Gig 1/0 to a sub-sampling rate of 1 to achieve an actual rate of 4096.
3. Configure interface Gig 1/1 to a sub-sampling rate of 2 to achieve an actual rate of 8192.

Note: Sampling rate backoff can change the sampling rate value that is set in the hardware. This equation shows the relationship between actual sampling rate, sub-sampling rate, and the hardware sampling rate for an interface:

*Actual sampling rate = sub-sampling rate * hardware sampling rate*

Note the absence of a configured rate in the equation. That is because when the hardware sampling rate value on the port-pipe exceeds the configured sampling rate value for an interface, the actual rate changes to the hardware rate. The sub-sampling rate never goes below a value of one.

Back-off Mechanism

If the sampling rate for an interface is set to a very low value, the CPU can get overed with flow samples under high-traffic conditions. In such a scenario, a binary back-off mechanism gets triggered, which doubles the sampling-rate (that is, halves the number of samples per second) for all interfaces. The backoff mechanism continues to double the sampling-rate until CPU over condition is cleared. This is as per sFlow version 5 draft. Once the back-off changes the sample-rate, users must manually change the sampling rate to the desired value.

As a result of back-off, the actual sampling-rate of an interface may differ from its configured sampling rate. The actual sampling-rate of the interface and the configured sample-rate can be viewed by using the **show sflow** command.

sFlow on LAG ports

When a physical port becomes a member of a LAG, it inherits the sFlow configuration from the LAG port.

Extended sFlow

Extended sFlow packs additional information in the sFlow datagram depending on the type of sampled packet. The following options can be enabled:

- **extended-switch** — 802.1Q VLAN ID and 802.1p priority information
- **extended-router** — Next-hop and source and destination mask length.
- **extended-gateway** — Source and destination AS number and the BGP next-hop.

Note: The entire AS path is not included. BGP community-list and local preference information are not included. These fields are assigned default values and are not interpreted by the collector.

Use the command **sflow** [**extended-switch**] [**extended-router**] [**extended-gateway**] **enable** command. By default packing of any of the extended information in the datagram is disabled.

Use the command **show sflow** to confirm that extended information packing is enabled, as shown in Figure 129.

Figure 129 Confirming that Extended sFlow is Enabled

```
Force10#show sflow
sFlow services are disabled
Global default sampling rate: 32768
Global default counter polling interval: 20
Global extended information enabled: gateway, router, switch
0 collectors configured
0 UDP packets exported
0 UDP packets dropped
0 sFlow samples collected
0 sFlow samples dropped due to sub-sampling
```

If none of the extended information is enabled, the **show** output is as shown in Figure 130.

Figure 130 Confirming that Extended sFlow is Disabled

```
Force10#show sflow
sFlow services are disabled
Global default sampling rate: 32768
Global default counter polling interval: 20
Global extended information enabled: none
0 collectors configured
0 UDP packets exported
0 UDP packets dropped
0 sFlow samples collected
0 sFlow samples dropped due to sub-sampling
```

Important Points to Remember

- The IP destination address has to be learned via BGP in order to export extended-gateway data.

- If the IP destination address is not learned via BGP the Force10 system does not export extended-gateway data.
- If the IP source address is learned via IGP then *srcAS* and *srcPeerAS* are zero.
- srcAS and srcPeerAS might be zero even though the IP source address is learned via BGP. The Force10 system packs the srcAS and srcPeerAS information only if the route is learned via BGP and it is reachable via the ingress interface of the packet.

The previous points are summarized in following table:

IP SA	IP DA	srcAS and srcPeerAS	dstAS and dstPeerAS	Description
static/connected/IGP	static/connected/IGP	—	—	Extended gateway data is not exported because there is no AS information.
static/connected/IGP	BGP	0	Exported	src_as & src_peer_as are zero because there is no AS information for IGP.
BGP	static/connected/IGP	—	—	Extended gateway data is not be exported because IP DA is not learned via BGP.
BGP	BGP	Exported	Exported	Extended gateway data is not packed.

Chapter 12 Port Monitoring

Port Monitoring permits monitoring of network traffic by forwarding a copy of each incoming or outgoing packet from one port to another port. This section contains:

- Important Points to Remember on page 249
- Configuring Port Monitoring on page 249
- Port Monitoring Configuration Examples on page 250

Important Points to Remember

- Port Monitoring is supported on TeraScale E-Series platforms.
- A SONET port can only be configured as a monitored port.
- The Monitored and Monitoring ports must be on the same switch.
- Port Monitoring supports only one Monitored and one Monitoring from a single port pipe.
- Monitored ports from different sessions can set destination as the same Monitoring port.
- FTOS supports as many monitor sessions on a system as the number of port-pipes.
- Port Monitoring is supported on physical ports only. Logical interfaces, such as Port-Channel and VLANs are not supported.

Note: The Monitoring port should not be a part of any other configuration. Avoid assigning VLAN membership to this port. If you attempt to configure a Layer 2 interface assigned to a VLAN as a destination monitoring port, the configuration is rejected with an error message similar to:

```
"% Error: Port is in Layer-2 mode Gi 7/13."
```

In addition, note that policy-based routing is not supported on the monitoring port. Generally, the monitoring port should have “no ip address" and "no shutdown" as the only configuration.

Configuring Port Monitoring

create a monitoring session

To enable Port Monitoring, create a monitoring session:

Command Syntax	Command Mode	Purpose
monitor session *session-ID*	CONFIGURATION	Create a monitoring sessionand enter a session identification number. Range: 0 to 65535

Only one monitor session per port-pipe is allowed, therefore at most, only 28 monitor sessions can be enabled at any time.

The monitoring command is saved in the running configuration at the Monitor Session mode level and can be restored after a chassis re.

To display the monitor session, use the **show monitor session** command.

```
Force10#show monitor session 11

SessionID       Source       Destination     Direction      Mode
-----------     --------     -------------   ------------   -------
   11           Gi 10/0      Gi 10/47        rx             interface

Force10#
```

Figure 131 show monitor session Command Example

Port Monitoring Configuration Examples

Port Monitoring is configured on Core-2 such that Gi 0/0 is the monitored port and Gi 1/0 is the monitoring Port. Traffic A is destined towards Server A and Server B is sending Traffic B to switch Core-1. Therefore, both incoming and outgoing traffic on Gi 0/0 (Traffic A and Traffic B) are monitored to Gi 1/0.

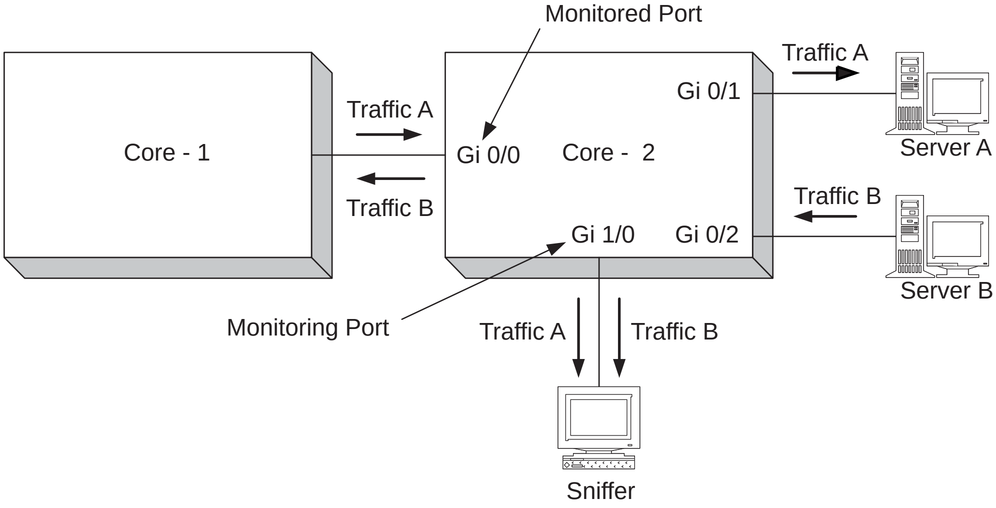

Figure 132 Port Monitoring Example

core-2 configuration example

The following example is for switching traffic; similar applies when routing traffic.

```
Core-2#sh running-config interface gigabitethernet 0/0
!
interface GigabitEthernet 0/0
 no ip address
 switchport
 no shutdown

Core-2#sh running-config interface gigabitethernet 1/0
!
interface GigabitEthernet 1/0
 no ip address
 no shutdown

Core-2#sh running-config interface gigabitethernet 0/1
!
interface GigabitEthernet 0/1
 no ip address
 switchport
 no shutdown

Core-2#sh running-config interface gigabitethernet 0/2
!
interface GigabitEthernet 0/2
 no ip address
 switchport
 no shutdown

Core-2#sh running-config interface vlan 2
!
interface Vlan 2
 description Engineering
 ip address 192.168.1.253/24
 untagged GigabitEthernet 0/0-2
 shutdown

Core-2#show running-config monitor session
!
monitor session 8
 source GigabitEthernet 0/0 destination GigabitEthernet 1/0 direction rx
!
monitor session 11
 source GigabitEthernet 1/0 destination GigabitEthernet 0/0 direction rx


Core-2#
```

Flow-based Monitoring

Flow-based monitoring permits monitoring of a specific set of Layer 2 or Layer 3 flows defined by an ACL, instead of monitoring all packets which enter or exit an interface. This method reduces bandwidth utilization. Flow-based monitoring is particularly useful to inspect suspicious packets, such as after a DOS attack.

Important Points to Remember

- Both switched and routed ports support flow-based monitoring.
- ACLs can be applied in the Ingress and Egress directions on E-Series.
- ACLs can be applied only in the Ingress direction on C-Series.
- The monitoring port can capture traffic in the following directions:
    - inbound only
    - outbound only
    - inbound and outbound

To configure Flow-based monitoring:

Step	Task	Command Syntax	Command Mode
1	Create a monitoring session and assign an ID number.	**monitor session** *number*	CONFIGURATION
2	Enable flow-based monitoring.	**flow-based enable**	MONITOR SESSION
3	Configure an ACL that specifes the **monitor** keyword. See configure a standard IP ACL on page 303.		
4	Apply the ACL to the monitored port. See Assign an IP ACL to an Interface on page 309.		

View monitoring information using the **show ip accounting access-list** and **show monitor session** commands, as shown in Figure 133

Figure 133 Configuring Flow-based Monitoring

```
Force10(conf)#monitor session 0
Force10#flow-based enable
Force10#show ip accounting access-list testflow
!
Extended Egress IP access list testflow on interface 11/13
Total cam count 4
seq 5 permit icmp any any monitor count bytes (100 packets 6800 bytes)
seq 10 permit ip 102.1.1.0/24 any monitor count bytes (100 packets 6800 bytes)
seq 15 deny udp any any count bytes (100 packets 6800 bytes)
seq 20 deny tcp any any count bytes (100 packets 6800 bytes)
Force10#show monitor session
SessionID Source Destination Direction Mode Type
--------- ------ ----------- --------- ---- ----
10 Gi 11/13 Gi 11/35 both interface Flow-based
```

Chapter 13 LACP

The Link Aggregation Control Protocol (LACP) provides a standardized means of exchanging information between two systems (also called Partner Systems) and automatically establishes Link Aggregation Groups between the two partner systems. LACP is supported on the E-Series and C-Series chassis. LACP permits the exchange of messages on a link to allow their Link Aggregation Control instances to:

- reach agreement on the identity of the Link Aggregation Group to which the link belongs
- move the link to that Link Aggregation Group and
- enable the transmission and reception functions in an orderly manner.

The FTOS implementation of LACP is based on the standards specified in the IEEE 802.3: “Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.”

LACP functions by constantly exchanging custom MAC PDUs across LAN Ethernet links. The protocol packets are only exchanged between ports that are configured as LACP capable.

Important Points to Remember

- **The port-channel number mode number** command is not permitted if the **port channel** already has a statically defined member in it (**channel-member** command).
- Static **port-channel number** cannot be created if a dynamic port channel already exists.
- A dynamic port channel can be created with any type of configuration.
- LACP enables the user to add members to the LACP global port channel as long as there are no static members configured using the **channel-member** command.
- Behavior of issuing **shutdown** or **no interface port-channel**.
  - The command **no interface port-channel number** deletes the specified port channel including when the port channel is dynamically created. When the user deletes this port channel using this command, all of the LACP specific commands on the member interfaces are automatically removed and the interfaces are restored to a state that is ready to be configured.

    Note: There will be no configuration on the interface since that is required for an interface to be part of a LAG.
  - Issuing the command **shutdown** on port-channel “xyz” disables the port channel and retains the user commands. However, the system does not allow the channel number “xyz” to be statically created.
- If a physical interface is a part of static LAG, then the command **port-channel-protocol lacp** will be rejected on that interface
- If a physical interface is a part of a dynamic LAG, it can not be added as a member of a static port-channel. The command **channel-member gigabitethernet x/y** will be rejected in the static port-channel interface for that physical interface.

- LACP does not add an interface to a LAG when one of the LAG members is shutdown on the remote interface. If a remote LAG member is shutdown, Message 4 appears on the local system when you attempt to add a member. In this case, enable all members of the LAG on the remote system, and then add any new members on the local system.

Message 4 LACP Remote Port Down Error Message

```
% Error: This port property does not match with other LAG member.
```

The LACP configuration tasks are:

- LACP modes on page 256
- LACP Monitoring using Syslog on page 256
- LACP Configuration on page 257
- Debugging LACP on page 261

LACP modes

FTOS shall provide multiple modes for configuration of LACP. The following are the modes an interface can exist:

• **Off**—in this mode, an interface is not capable of being part of a dynamic port channel group. LACP shall not run on any port that is configured to be in the OFF state.

• **Active**—in this mode, the interface is said to be in the “active negotiating state.” LACP runs on any link that is configured to be in the active state. The port in an active state also automatically initiates negotiations with other ports by initiating LACP packets.

• **Passive**—in this mode, the interface is not in an active negotiating state. LACP runs on any link that is configured to be in the passive state. The port in a passive state also responds to negotiations requests from other ports that are in active state. Ports in passive state respond to LACP packets.

FTOS shall provide the following:

- A port in an active state to set up a port channel (LAG group) with another port in an active state.
- A port in an active state to set up a port channel (LAG group) with another port in a passive state.

A port in a passive state cannot set up a port channel (LAG group) with another port in a passive state.

LACP Monitoring using Syslog

Syslog generates appropriate events to trigger faulty actions in LACP.

LACP Configuration

FTOS enables the user to configure LACP characteristics for a physical port. If the aggregated ports are configured with compatible LACP modes (Off, Active, Passive), FTOS LACP can automatically link them as defined in IEE 802.3 specification Section 43. Use the following commands to configure LACP:

Command Syntax	Command Mode	Purpose
[**no**] **lacp system-priority** *priority-value*	CONFIGURATION interface	To configure the system priority. Priority value shall range from 1 to 65535 (higher the number lower shall be the priority). Default is 32768

Command Syntax	Command Mode	Purpose
[**no**] **port-channel-protocol lacp**	CONFIGURATION interface	To enable or disable LACP on any LAN port: • Default is "LACP disabled" • This command creates a new context.

Command Syntax	Command Mode	Purpose
[**no**] **port-channel** *number* **mode** [**active** | **passive** | **off**]	CONFIGURATION interface	To configure LACP mode. • Default is "LACP active" • **number** cannot statically contain any links

Command Syntax	Command Mode	Purpose
[**no**] **lacp port-priority** *priority-value*	CONFIGURATION interface	To configure port priority. • Priority value shall range from 1 to 65535 (higher the number lower shall be the priority). • Default is 32768

LACP Long-timeout

PDUs are exchanged between port-channel interfaces to maintain LACP sessions. PDUs are transmitted at either a slow or fast transmission rate, depending upon the LACP timeout value. The timeout value is the amount of time that a port-channel interface waits for a PDU from the remote system before bringing the LACP session down. The default time out value is 1 second, but it can be configured to be 30 seconds. Configuring a longer timeout might prevent the port-channel from flapping if the remote system is up, but temporarily unable to transmit PDUs due to a system interruption.

Note: LACP long-timeout is available for dynamic port-channel interfaces only. The command **lacp long-timeout** can be entered for static port-channels, but it has no effect.

To configure LACP long-timeout:

Step	Task	Command Syntax	Command Mode
1	Set the LACP timeout value to long timeout.	**lacp long-timeout**	INTERFACE PORT-CHANNEL

Figure 134 Configuring LACP Long-timeout

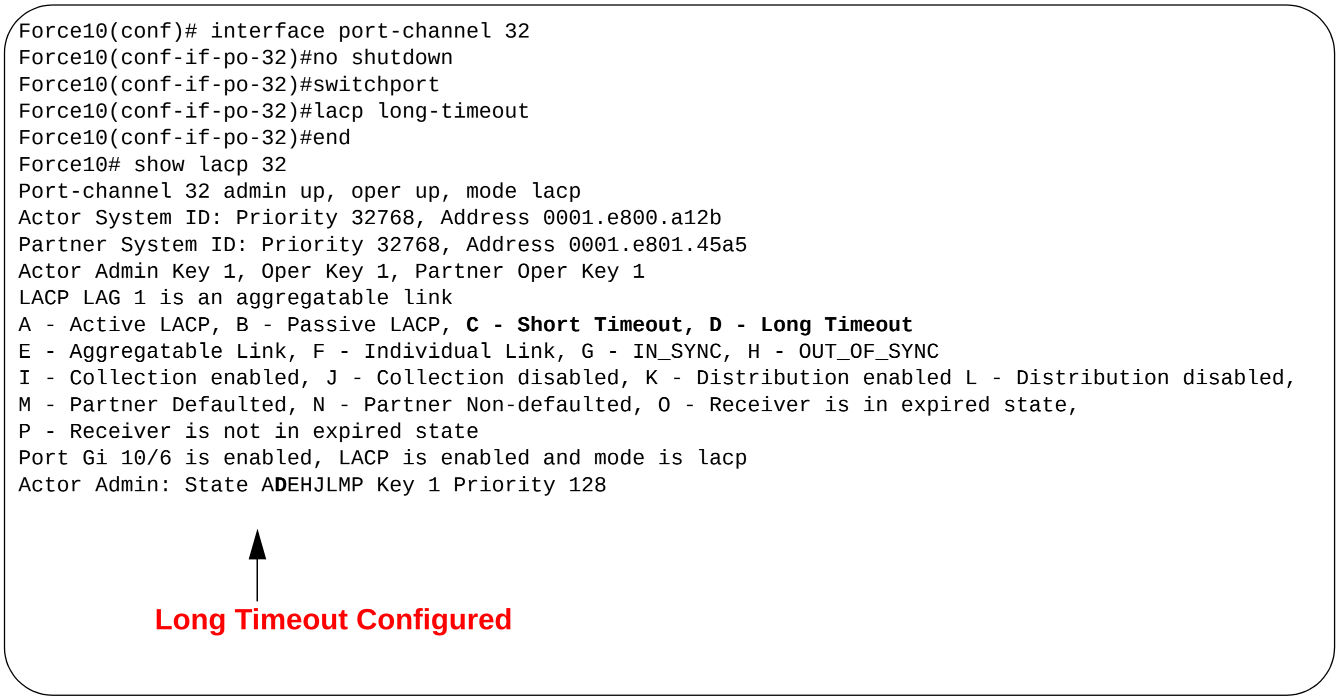

```
Force10(conf)# interface port-channel 32
Force10(conf-if-po-32)#no shutdown
Force10(conf-if-po-32)#switchport
Force10(conf-if-po-32)#lacp long-timeout
Force10(conf-if-po-32)#end
Force10# show lacp 32
Port-channel 32 admin up, oper up, mode lacp
Actor System ID: Priority 32768, Address 0001.e800.a12b
Partner System ID: Priority 32768, Address 0001.e801.45a5
Actor Admin Key 1, Oper Key 1, Partner Oper Key 1
LACP LAG 1 is an aggregatable link
A - Active LACP, B - Passive LACP, C - Short Timeout, D - Long Timeout
E - Aggregatable Link, F - Individual Link, G - IN_SYNC, H - OUT_OF_SYNC
I - Collection enabled, J - Collection disabled, K - Distribution enabled L - Distribution disabled,
M - Partner Defaulted, N - Partner Non-defaulted, O - Receiver is in expired state,
P - Receiver is not in expired state
Port Gi 10/6 is enabled, LACP is enabled and mode is lacp
Actor Admin: State ADEHJLMP Key 1 Priority 128
```

Long Timeout Configured

Note: View PDU exchanges and the timeout value using the command **debug lacp**.

creating a LAG

To create a LAG, define the LAG configuration first, then create a port channel and place it into the default and a non-default VLANs.

Define the interface configuration to dynamically create the LAG port channel (32):

```
Force10(conf)#interface Gigabitethernet 3/15
Force10(conf-if-gi-3/15)#no shutdown
Force10(conf-if-gi-3/15)#port-channel-protocol LACP
Force10(conf-if-gi-3/15)#port-channel 32 mode active
...
Force10(conf)#interface Gigabitethernet 3/16
Force10(conf-if-gi-3/16)#no shutdown
Force10(conf-if-gi-3/16)#port-channel-protocol LACP
Force10(conf-if-gi-3/16)#port-channel 32 mode active
...
Force10(conf)#interface Gigabitethernet 4/15
Force10(conf-if-gi-4/15)#no shutdown
Force10(conf-if-gi-4/15)#port-channel-protocol LACP
Force10(conf-if-gi-4/15)#port-channel 32 mode active
...
Force10(conf)#interface Gigabitethernet 4/16
Force10(conf-if-gi-4/16)#no shutdown
Force10(conf-if-gi-4/16)#port-channel-protocol LACP
Force10(conf-if-gi-4/16)#port-channel 32 mode active
```

Figure 135 Dynamically Creating a LAG Port Channel Example

Use the **switchport** command to automatically place the port channel 32 in the default VLAN: Define the interface configuration to dynamically create the LAG port channel (32):

```
Force10(conf)#interface port-channel 32
Force10(conf-if-po-32)#no shutdown
Force10(conf-if-po-32)#switchport
```

Figure 136 Placing Port Channel into the Default VLAN

Placing the port channel into a non-default VLAN: Define the interface configuration to dynamically create the LAG port channel (32):

```
Force10(conf)#interface vlan 10
Force10(conf-if-vl-10)#tagged port-channel 32
```

Figure 137 Placing a Port Channel into a Non-default VLAN

creating a LAG and interface configurations

The user may first create a LAG configuration, then create the interface configuration. To do so, first use the **switchport** command to automatically place the port channel 32 into the default VLAN:

```
Force10(conf)#interface port-channel 32
Force10(conf-if-po-32)#no shutdown
Force10(conf-if-po-32)#switchport
```

Figure 138 Placing a Port Channel into the Default VLAN

Then create an interface configuration by selecting a port, specifying the LACP protocol, then activating the port channel.:

```
Force10(conf)#interface gigabitethernet 3/15
Force10(conf-if-gi-3/15)#no shutdown
Force10(conf-if-gi-3/15)#port-channel-protocol LACP
Force10(conf-if-gi-3/15)#port-channel 32 mode active
...
Force10(conf)#interface gigabitethernet 3/16
Force10(conf-if-gi-3/16)#no shutdown
Force10(conf-if-gi-3/16)#port-channel-protocol LACP
Force10(conf-if-gi-3/16)#port-channel 32 mode active
...
Force10(conf)#interface gigabitethernet 4/15
Force10(conf-if-gi-4/15)#no shutdown
Force10(conf-if-gi-4/15)#port-channel-protocol LACP
Force10(conf-if-gi-4/15)#port-channel 32 mode active
...
Force10(conf)#interface gigabitethernet 4/16
Force10(conf-if-gi-4/16)#no shutdown
Force10(conf-if-gi-4/16)#port-channel-protocol LACP
Force10(conf-if-gi-4/16)#port-channel 32 mode active
```

Figure 139 Creating an Interface Configuration

The **port-channel 32 mode active** command shown above may be successfully issued as long as there is no existing static channel-member configuration in interface port-channel 32.

To place the port channel into a VLAN.:

```
Force10(conf)#interface vlan 10
Force10(conf-if-vl-10)#tagged port-channel 32
```

Figure 140 Creating an Interface Configuration

Debugging LACP

Command Syntax	Command Mode	Purpose
[no] **debug lacp** [**config | events | pdu** [**in | out |** [*interface* [**in | out**]]]]	EXEC mode	To debug LACP including configuration and events.

Chapter 14 VLAN-Stack VLANs

C-Series	NO
E-Series	✓

Platform Specific Feature: VLAN-Stack Vlans is supported on E-Series only.

This chapter covers the following topics:

- Implementation Information on page 264
- Configuration Task List for VLAN-Stack VLANs on page 265
- VLAN-Stack Configuration Example on page 268

With VLAN-Stack VLANs, you can assign a VLAN ID to untagged frames or frames that already contain a customer VLAN tag. All Customer frames (whether tagged or untagged) are tagged at ingress with a VLAN tag, which is used to forward traffic through the VLAN-Stack aware network. By using a single VLAN tag for multiple VLANs, the customer's VLAN tags are preserved and increase the number of unique VLANs supported in the network because the customer VLAN tags are hidden inside the new VLAN-Stack VLAN tag.

Note: VLAN-Stack VLAN feature is available on Force10 Networks ED, EE and EF series line cards.

A VLAN-Stack tag (with a different Protocol Type) and a new CRC is inserted in every frame at the ingress edge device. These are removed at the egress edge device and the original VLAN tagging is preserved. The intermediate devices treat the frame as a regular Ethernet frame, however the insertion of VLAN-Stack tag increases the maximum frame size by 4 bytes, making it a Baby Giant frame.

Figure 141 illustrates where the VLAN-stack Tag is added (after the Source Address and before the VLAN ID tag). The first part of the tag is the user-configurable protocol type value (default 0x9100) and the second part is the VLAN ID you assign to the VLAN-stack (0007).

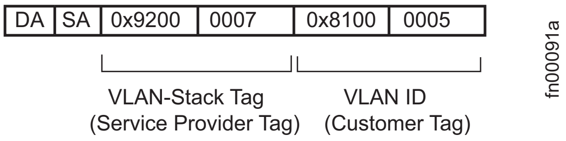

Figure 141 Location of VLAN-Stack Tag in Packet Header

Implementation Information

The VLAN-Stack tag uses a configurable Protocol Type. The default is 0x9100, but you can set it to any value. Intermediate devices in a VLAN-Stack network recognize this Protocol Type and switch packets based on it.

To create a VLAN-Stack aware network, you must designate interfaces as either VLAN-Stack access ports or VLAN-Stack trunk ports. You must assign these interfaces to a VLAN-Stack enabled VLAN.

The following interface types can be VLAN-Stack access or trunk ports:

- Ethernet (Gigabit Ethernet and 10 Gigabit Ethernet)
- Port Channels

Interfaces in the default VLAN, with VLAN-stack access or VLAN-stack trunk configuration, do not switch untagged traffic. These interfaces will switch traffic only when they are added to a non-default VLAN-STACK enabled VLAN.

With VLAN-Stack VLANs, STP traffic can be forwarded or tunneled across the VLAN-Stack network depending on if it is enabled in the network. If STP is enabled on a VLAN-Stack network, then any STP traffic sent by the customer's network is accepted and forwarded natively across the VLAN-Stack network. If STP is disabled on the VLAN-Stack network, then any STP traffic sent by the customer network is tunneled across the VLAN-Stack network.

Important Points to Remember

- Spanning-tree BPDU from the customer's networks are tunneled across the VLAN-Stack network if STP is *not* enabled on VLAN-Stack network. However, if STP is enabled in VLAN-Stack network, STP BPDU from the customer's networks are consumed and not tunneled across the network.
- Layer-3 protocols are not supported on a VLAN-Stack network.

- Assigning an IP address to a VLAN-Stack VLAN is supported when all the members are only VLAN-Stack truck ports. IP addresses on a VLAN-Stack enabled VLAN is not supported if the VLAN contains VLAN-Stack access ports. This facility is provided for SNMP management over a VLAN-Stack enabled VLAN containing only VLAN-Stack trunk interfaces. Layer-3 routing protocols on such a VLAN are not supported.
- It is recommended that you do not use the same MAC address, on different customer VLANs, on the same VLAN-Stack VLAN.

Configuration Task List for VLAN-Stack VLANs

The following list includes the configuration tasks for VLAN-Stack VLANs.

- configure VLAN-Stack access ports on page 265 (mandatory)
- configure a VLAN-Stack trunk port on page 266 (mandatory)
- configure VLAN-Stack VLAN on page 266 (mandatory)
- set the protocol type for VLAN-Stack VLANs on page 267 (optional)

For a complete listing of all commands related to VLAN-Stacking, refer to .

configure VLAN-Stack access ports

VLAN-Stack access ports can only belong to one VLAN-Stack VLAN.

To configure an interface as a VLAN-Stack access port, use these commands:

Step	Command Syntax	Command Mode	Purpose
1	**switchport**	INTERFACE	Designate the interface as a Layer 2 interface.
2	**vlan-stack access**	INTERFACE	Specify as a VLAN-Stack access port.
3	**no shutdown**	INTERFACE	Enable the interface.

Use the **show config** command in the INTERFACE mode or the **show running-config interface** *interface* command to view the interface's configuration.

```
Force10#sh run int gi 7/0
!
interface GigabitEthernet 7/0
 no ip address
 switchport
 vlan-stack access
 no shutdown
Force10#
```

Figure 142 show running-config interface on the E1200-1

To remove the VLAN-Stack access port designation, you must first remove the port from the VLAN-Stack VLAN, using the **no member** *interface* command.

configure a VLAN-Stack trunk port

A VLAN-Stack trunk port is a Layer-2 port that can be a member of multiple VLAN-Stack VLANs.

To configure a VLAN-Stack trunk port, use these commands in the following sequence, starting in the INTERFACE mode:

Step	Command Syntax	Command Mode	Purpose
1	**switchport**	INTERFACE	Designate the interface as a Layer 2 interface.
2	**vlan-stack trunk**	INTERFACE	Specify as a VLAN-Stack trunk port.
3	**no shutdown**	INTERFACE	Enable the interface.

Use the **show config** command in the INTERFACE mode or the **show running-config interface** *interface* command in the EXEC privilege mode to view the configuration.

```
E1200-1#sh run int gi 7/12
!
interface GigabitEthernet 7/12
 no ip address
 switchport
 vlan-stack trunk
 no shutdown
E1200-1#
```

Figure 143 show running-config interface

To remove the VLAN-Stack trunk port designation, you must first remove the port from the VLAN-Stack VLAN, using the **no member** *interface* command.

configure VLAN-Stack VLAN

After you configure interfaces as VLAN-Stack access or trunk ports, add them to a VLAN-Stack VLAN. If you do not add them to a VLAN-Stack VLAN, the ports will be part of the Default VLAN as untagged ports.

To configure a VLAN-Stack VLAN, use these commands in the following sequence starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**vlan-stack compatible**	VLAN	Place the VLAN in VLAN-Stack mode.

Step	Command Syntax	Command Mode	Purpose
2	**member** *interface*	VLAN	Add a VLAN-Stack access port or VLAN-Stack trunk port to the VLAN.

Use the **show vlan** command in the EXEC privilege mode to view the members of a VLAN-Stack VLAN. Members of the VLAN-Stack VLAN are identified by M in the Q column.

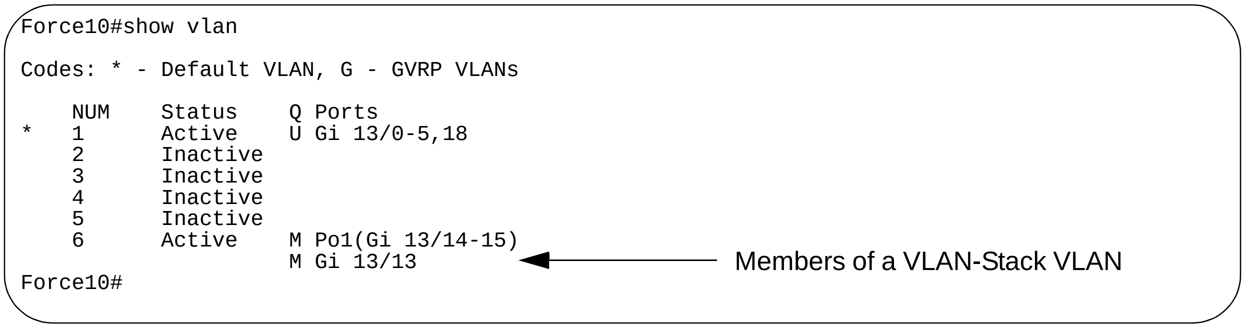

```
Force10#show vlan

Codes: * - Default VLAN, G - GVRP VLANs

    NUM    Status    Q Ports
*   1      Active    U Gi 13/0-5,18
    2      Inactive
    3      Inactive
    4      Inactive
    5      Inactive
    6      Active    M Po1(Gi 13/14-15)
                     M Gi 13/13
Force10#
```

Members of a VLAN-Stack VLAN

Figure 144 show vlan Command Example

set the protocol type for VLAN-Stack VLANs

By default, the VLAN-stack protocol tag is set at 0x9100. In the packet header, the VLAN-stack protocol tag is added after the Destination Address and before the VLAN ID (see Figure 141).

To change the protocol number for VLAN-Stack VLAN, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
vlan-stack protocol-type *value*	CONFIGURATION	Configure the Protocol Type to differentiate it from VLANs. • default value: 9100

To view the non-default VLAN-Stack protocol type configuration, use the **show running-config** command in the EXEC privilege mode. If you do not change the protocol-type value, the default value 0x9100 is used and it does not appear in the running-config file.

VLAN-Stack Configuration Example

Figure 145 is an example of a VLAN-stack network. In this network, customer traffic enters the network on VLAN-stack access port in E1200-1 and E1200-2. The traffic is assigned a VLAN-stack VLAN and those VLAN-stack VLANs are switched in E1200-3. In this example, traffic from Customer 3 entering E1200-2 is switched through E1200-3 to the Customer 3 port on E1200-1.

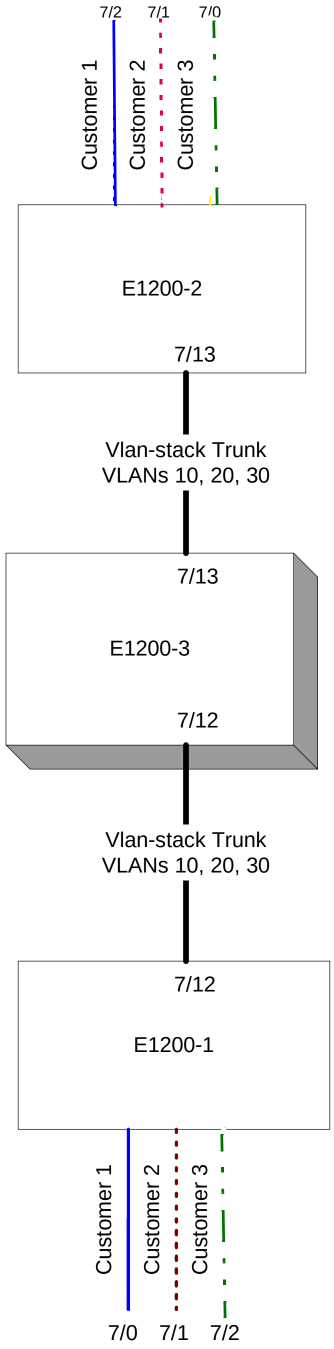

Figure 145 VLAN-Stack Network Example Diagram

E1200-1 Configuration

```
E1200-1#sh run int gi 7/0
!
interface GigabitEthernet 7/0
 no ip address
 switchport
 vlan-stack access
 no shutdown
E1200-1#sh run int gi 7/1
!
interface GigabitEthernet 7/1
 no ip address
 switchport
 vlan-stack access
 no shutdown
E1200-1#sh run int gi 7/2
!
interface GigabitEthernet 7/2
 no ip address
 switchport
 vlan-stack access
 no shutdown
E1200-1#
E1200-1#sh run int gi 7/12
!
interface GigabitEthernet 7/12
 no ip address
 switchport
 vlan-stack trunk
 no shutdown
E1200-1#
E1200-1#sh run int vlan 10
!
interface Vlan 10
 no ip address
 vlan-stack compatible
 member GigabitEthernet 7/0,12
 shutdown
E1200-1#sh run int vlan 20
!
interface Vlan 20
 no ip address
 vlan-stack compatible
 member GigabitEthernet 7/1,12
 shutdown
E1200-1#sh run int vlan 30
!
interface Vlan 30
 no ip address
 vlan-stack compatible
 member GigabitEthernet 7/2,12
 shutdown
E1200-1#
E1200-2 Configuration
E1200-2#sh run int gi 7/0
!
interface GigabitEthernet 7/0
 no ip address
 switchport
 vlan-stack access
 no shutdown
E1200-2#sh run int gi 7/1
!
```

```
interface GigabitEthernet 7/1
 no ip address
 switchport
 vlan-stack access
 no shutdown
E1200-2#sh run int gi 7/2
!
interface GigabitEthernet 7/2
 no ip address
 switchport
 vlan-stack access
 no shutdown
E1200-2#
E1200-2#sh run int gi 7/13
!
interface GigabitEthernet 7/13
 no ip address
 switchport
 vlan-stack trunk
 no shutdown

E1200-2#sh run int vlan 10
!
interface Vlan 10
 no ip address
 vlan-stack compatible
 member GigabitEthernet 7/0,13
 shutdown
E1200-2#sh run int vlan 20
!
interface Vlan 20
 no ip address
 vlan-stack compatible
 member GigabitEthernet 7/1,13
 shutdown
E1200-2#sh run int vlan 30
!
interface Vlan 30
 no ip address
 vlan-stack compatible
 member GigabitEthernet 7/2,13
 shutdown
```

E1200-3 Configuration

```
E1200-3#sh run int gi 7/12
!
interface GigabitEthernet 7/12
 no ip address
 switchport
 vlan-stack trunk
 no shutdown
E1200-3#sh run int gi 7/13
!
interface GigabitEthernet 7/13
 no ip address
 switchport
 vlan-stack trunk
 no shutdown
E1200-3#E1200-3#show run int vlan 10
!
interface Vlan 10
 no ip address
 vlan-stack compatible
 member GigabitEthernet 7/12,13
 shutdown
E1200-3#sh run int vlan 20
!
interface Vlan 20
 no ip address
 vlan-stack compatible
 member GigabitEthernet 7/12,13
 shutdown
E1200-3#sh run int vlan 30
!
interface Vlan 30
 no ip address
 vlan-stack compatible
 member GigabitEthernet 7/12,13
 shutdown
```

Chapter 15 FVRP

C-Series	NO
E-Series	✓

Platform Specific Feature: FVRP is supported on E-Series only.

Force10 VLAN Redundancy Protocol (FVRP) is a proprietary Layer-2 feature that provides rapid failover of links by using VLAN redundancy. With FVRP enabled, one link is carrying active traffic to the core for that VLAN and the other links are in standby mode.

This chapter contains the following sections:

Definitions

FVRP VLAN—a VLAN with FVRP enabled and contains tagged interfaces. FVRP-aware VLAN must contain both core and access ports, and may contain uplinks.

Core switches—an E-Series that participates in the Master election process, including tracking the uplink (if configured). You can have two Core switches, where one is the Master Core switch and the other is the Standby Core switch. The Core switches generate FVRP Configuration Messages to track the FVRP topology, including monitoring and avoiding network loops.

Master—a FVRP VLAN with forwarding links. The Master is a VLAN on a core switch that wins the FVRP master election process. It has the most number of active access ports per VLAN, the lowest FVRP VLAN priority, and the lowest ID (the MAC address and VLAN ID).

Standby—a VLAN with blocked links. This VLAN may become Master if a link goes down on the Master or the FVRP priority on the VLAN changes.

Access switch—an edge switch that does not participate in the FVRP Master election process and need not be an E-Series. The switch responds to configuration messages to avoid loops. If the switch is FVRP-aware (that is, for the hardware platform), it performs quick MAC address aging in response to a Flush Message from the Core switch.

Access link—any downstream link connecting a FVRP core switch with an edge (Access) switch. The link must be a tagged member of a FVRP VLAN.

core link—a Layer-2 interface that is a member of a FVRP VLAN and connects FVRP core switches.

uplink—a Layer-2 interface that is a member of a FVRP VLAN, is tracked by the Master, and is connected to different network.

FVRP Domain—a group of a master VLAN and member VLANs. All VLANs in an FVRP Domain share the same Master Core switch and Standby Core switch. A domain allows faster failover, protocol scalability, and configuration simplicity.

FVRP Region—is a group of access switches which are multi-homed to the same set of core switches. The core switches within an FVRP region exchange FVRP control messages. You configure a Region name. Multiple FVRP regions can be stacked together or in a hierarchy to achieve end-to-end redundancy and per VLAN loop free topology in a switched network.

Control-VLAN—a VLAN that sends FVRP Control messages between FVRP Core and Access switches. This VLAN contains no FVRP commands, but does contain all the interfaces on FVRP-aware switches and all interfaces on non-FVRP-aware switches that connect to an FVRP core switch.

Benefits

FVRP uses VLAN redundancy to provide rapid fail over in Layer-2 networks. This feature works with tagged VLANs and stacked VLANs, Port-Channels, and different Ethernet interfaces. With multiple VLANs participating, FVRP provides the following benefits:

- loop-free topology
- fail over in approximately 2 seconds (depending on the network topology)
- recovery if a access link, switch or uplink fails
- per-VLAN redundancy (VLAN grouping)
- downstream awareness of failure
- transparent to and no interoperability with third-party equipment
- multiple backups
- hierarchical FVRP regions

Implementation

FVRP must be enabled globally and enabled on at least one VLAN on a participating switch. VLAN members (physical interfaces or Port Channels) must be tagged (802.1Q enabled) and the VLAN cannot have an IP address configured on it.

FVRP is supported on Gigabit Ethernet, 10 Gigabit Ethernet, and port channel interfaces.

The Default VLAN does not support FVRP.

An E-Series is called a FVRP-aware switch. With FVRP enabled, it participates in the Master election process and either becomes a Master or Standby switch. FVRP core and standby switches exchange FVRP Configuration messages every Hello timer interval to determine the mastership of one or more VLANs.

If a link fails, FVRP uses the redundant links to reroute traffic. If an access link goes down, the Master re-routes the traffic to different access links to ensure the traffic reaches its destination. After a failover, the Flush Address Message is sent twice to ensure that all MAC addresses for the VLAN are removed from the Access switch. This failover can occur in approximately two seconds, depending on the topology.

For "non-Force10" access switches, the link flap mechanism ensures that stale MAC addresses are removed during topology change events. If the non-Force10 switch does not detect the link flap, it may be necessary to manually shut/no-shut the link.

If an uplink goes down, the FVRP process adds the uplink’s priority value to the Master priority for all VLANs for which the uplink is carrying traffic. By modifying the VLAN priority, the standby switch can take over mastership of the VLAN after receiving configuration messages from the existing master.

During topology initialization, all links are blocked and each FVRP-aware switch transmits FVRP configuration messages. Using the criteria listed below, the switches begin the Master election process and select a Master switch for each FVRP VLAN in the topology.

FVRP Master Election

The following criteria determines which switch in a FVRP VLAN becomes the Master switch:

- FVRP access port availability (that is, the VLAN with the most active access interfaces)
- highest priority (lowest priority value, using the **fvrp priority** command)
- lowest MAC address (on the Management port)

Configuration Task List for FVRP

To configure FVRP, use commands in the PROTOCOL FVRP mode, the INTERFACE mode, and the VLAN mode.

The following list includes the configuration tasks for FVRP:

- enable FVRP on an interface on page 276 (required)
- enable FVRP on a VLAN on page 277 (required)
- enable FVRP globally on page 278 (required)
- changing FVRP parameters on page 279 (optional)

For a complete listing of all commands related to FVRP, refer to .

enable FVRP on an interface

FVRP is supported on Layer-2 interfaces only and the interfaces must be tagged members of a FVRP VLAN. The interfaces must also be enabled.

When you add a Layer-2 interface to a FVRP VLAN, it is considered a core link, by default. You must specify interfaces as an access link or an uplink depending on what the link connects. Uplinks must connect a FVRP core switch to a different network, while access links connect FVRP core switches to an access switch. You can also specify that the interface is connected to an FVRP-aware system.

Core links do not require additional configuration. They are detected by the software and assigned Core port status. Use the **show fvrp vlan** command to view the different ports of a FVRP VLAN.

To enable FVRP on a Layer-2 interface and define its role, use one of the following commands in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
fvrp access [**region** *region-name*]	INTERFACE	Specify the interfaces as a FVRP access link. To assign an FVRP region, the interface must already have the **fvrp access** command configured.
fvrp aware	INTERFACE	Identify that the interface is connected to an FVRP-aware system.
fvrp uplink	INTERFACE	Specify the interface as a FVRP uplink. This link must connect a FVRP core switch to another network.

To view the interface configuration, use the **show config** command in the INTERFACE mode or the **show running interface** command (Figure 146) in the EXEC privilege mode.

```
Force10#show running-config interface gigabitEthernet 1/0
!
interface GigabitEthernet 1/0
 no ip address
 switchport
 no shutdown
 fvrp access
 fvrp access region X
 fvrp aware
Force10#
```

Figure 146 show running-config interface Command Example

enable FVRP on a VLAN

You must enable FVRP on a VLAN and assign FVRP parameters to the members of the VLAN. You cannot assign an IP address to a FVRP VLAN.

To enable FVRP on a VLAN, use the following commands in the VLAN mode:

Command Syntax	Command Mode	Purpose
no fvrp disable	VLAN	Enable FVRP on a VLAN.
fvrp core	VLAN	For core switches only, specify that the VLAN participates in the Master election process and can be either a Master or Standby. If the E-Series is an edge switch in the FVRP network, do not enter this command.

Note: The VLAN stack feature is not supported in FVRP.

Once you enable FVRP on a VLAN, you must specify which of the interfaces in the VLAN participate in the FVRP protocol. Redundant links can be grouped and one of the interfaces in that group will become the master access link for that VLAN. The active interface with the highest priority (a configurable parameter) becomes the master access link for that FVRP VLAN.

To ensure redundant links between FVRP Master and Standby switches, configure a port channel between the switches. If you do not specify it as an access or uplink, the software will recognize it as a core link.

To assign an interface to a FVRP group and assign it a priority, use the following commands in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
fvrp *interface* **group** *number*	VLAN	Assign a VLAN member to an FVRP group. • *interface*: enter the interface-type and slot/port information. • *number* range: 0-256. The default is 0.
fvrp *interface* **priority** *priority*	VLAN	Assign a priority value to an interface in a FVRP-aware VLAN. • *interface:* enter the interface-type and slot/port information. • *priority* range: 1 - 256. The default is 128.

To view the VLAN configuration, use the **show config** command in the VLAN mode or the **show running vlan** command in the EXEC privilege mode.

Once FVRP is completely configured, use the **show fvrp vlan** *vlan-id* command in the EXEC privilege mode (Figure 147) to view the status of interfaces and parameters.

```
Force10#sh fvrp vlan 100
     FVRP Vlan 100 Information
          FVRP Vlan Enabled
          FVRP Vlan Mode: Core
          FVRP Vlan State: StandBy
          FVRP Vlan priority: 128
          FVRP Vlan Hello time: 1
     Access Port 121 (GigabitEthernet 5/0) group 1 priority 128 is Blocking
     Access Port 122 (GigabitEthernet 5/1) group 1 priority 128 is Blocking
     Uplink Port 123 (GigabitEthernet 5/2) priority 128 is Blocking
Force10#
```

Figure 147 show fvrp vlan Command Example

enable FVRP globally

By default, FVRP is not enabled.

FVRP must be enabled on all E-Series with VLANs participating in the FVRP network.

To enable FVRP, use these commands in the following sequence starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**protocol fvrp**	CONFIGURATION	Enter the global FVRP mode.
2	**no disable**	FVRP	Enable FVRP globally. After you enter **no disable**, the software initially blocks all access ports for 35 seconds to complete the Master election.

To view the FVRP global configuration, enter the **show config** command in the FVRP mode or the **show running-config fvrp** command in the EXEC privilege mode.

```
Force10#show run fvrp
!
protocol fvrp
 no disable

Force10#
```

Figure 148 show running-config fvrp Command Example

changing FVRP parameters

FTOS provides different configurable parameters to affect the Master VLAN or Master link election process or the timing of FVRP messages. The Master VLAN and Master link are elected based on different factors, but you can influence which VLAN or link is chosen by configuring the priority value. The higher the priority (that is, the lower the number assigned) the more likely it is that the VLAN or link will be the Master.

The software uses the following criteria to choose the Master link:

- active port status (the interface must be enabled);
- best priority (the lowest **fvrp** *interface* **priority** command value); and
- lowest port index (an system internal parameter).

The criteria used to determine the Master VLAN are discussed in the .

Other parameters you can change include preemption, FVRP hello and hold timers.

Use the following commands to change FVRP VLAN parameters in the VLAN mode:

Command Syntax	Command Mode	Purpose
fvrp priority *priority*	VLAN	Change the priority of the VLAN to affect the Master election process. • *priority* range: 1 - 256. The default is 128. When an uplink for the VLAN goes down, the priority value for that uplink is added to the VLAN's priority value and a Master election process may begin.
fvrp hello-time *seconds*	VLAN	Change the time interval between FVRP configuration messages. • *seconds* range: 1 - 256. The default is 1 second.
fvrp hold-time *seconds*	VLAN	Change the delay the preemption of the Master to ensure that the switch's MAC address tables are stabilized. • *seconds* range: 1 - 256. The default is 1 second.
fvrp *interface* **preempt**	VLAN	Enable interface to assume Master access link for the VLAN. • *interface:* enter the interface-type and slot/port information.

Configuration Example

E1200-1 and E1200-2 are configured as FVRP core switches (see Figure 149). VLAN 100 and VLAN 200 are configured on all E1200s, and all ports, including the access ports, are added to both VLANs as tagged.

Configure the priority for VLAN 100 on E1200-1 so that E1200-1 is Master for VLAN 100. Configure the priority for VLAN 200 on E1200-2 so that E1200-2 is Master for VLAN 200. E1200-1 is the FVRP standby for VLAN 200 and E1200-2 is a standby switch for VLAN 100.

This configuration is represented in Figure 149. The solid lines represent the active access links for FVRP VLAN 100. The dotted lines are access links blocked by FVRP. On E1200-1, the forwarding access links forward VLAN 100 traffic and block all VLAN 200 traffic. On E1200-2, the forwarding access links block VLAN 100 traffic and forward all VLAN 200 traffic.

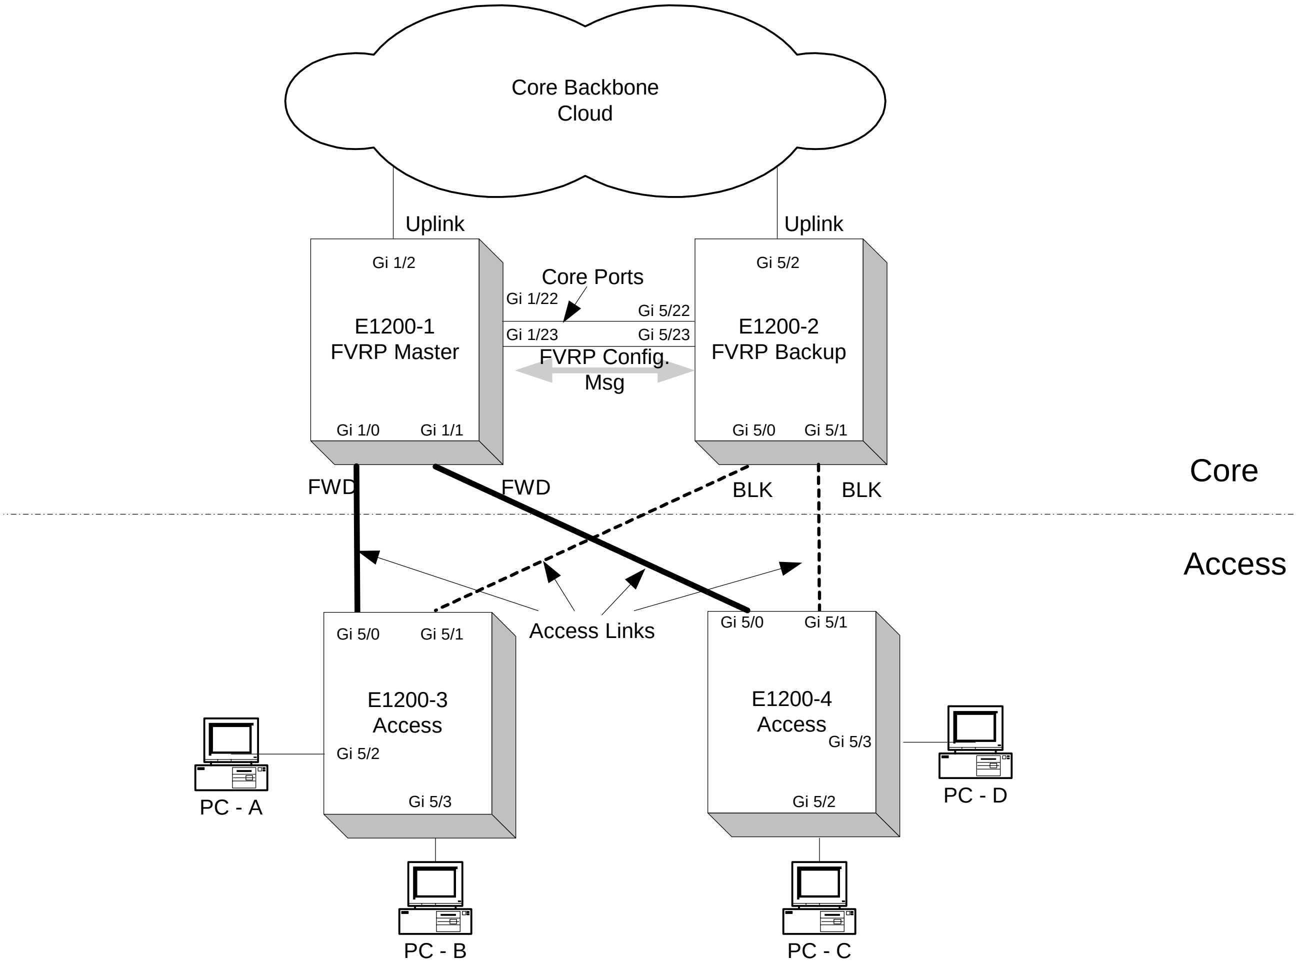

Figure 149 FVRP Network Diagram

E1200-1 Configuration

```
E1200-1#sh running-config fvrp
```

```
!
protocol fvrp
 fvrp control-vlan 10
 fvrp core region X
 no disable


E1200-1#sh running-config interface gi 1/2
!
interface GigabitEthernet 1/2
 no ip address
 switchport
 no shutdown
 fvrp uplink
E1200-1#sh running-config interface gi 1/0
!
interface GigabitEthernet 1/0
 no ip address
 switchport
 no shutdown
 fvrp access
 fvrp access region X
 fvrp aware
E1200-1#sh running-config interface gi 1/1
!
interface GigabitEthernet 1/1
 no ip address
 switchport
 no shutdown
 fvrp access
 fvrp access region X
 fvrp aware
E1200-1#sh running-config interface port-channel 10
!
interface port-channel 10
 no ip address
 switchport
 channel-member GigabitEthernet 1/22-23
 no shutdown

E1200-1#sh running-config interface gi 1/22
!
interface GigabitEthernet 1/22
 no shutdown
 no ip address

E1200-1#sh running-config interface gi 1/23
!
interface GigabitEthernet 1/23
 no shutdown
 no ip address

E1200-1#
E1200-1#sh running-config interface vlan 10
!
interface VLAN 10
 no ip address
 tagged GigabitEthernet 1/0-2
 tagged Port-channel 10
 no shutdown

E1200-1#sh running-config interface vlan 100
!
interface Vlan 100
 no ip address
 tagged GigabitEthernet 1/0-2
 no shutdown
 no fvrp disable
 fvrp priority 10
 fvrp GigabitEthernet 1/0 group 1
 fvrp GigabitEthernet 1/1 group 2
 fvrp core

E1200-1#sh running-config interface vlan 200
!
interface Vlan 200
 no ip address
 tagged GigabitEthernet 1/0-2
 no shutdown
 no fvrp disable
 fvrp GigabitEthernet 1/0 group 1
 fvrp GigabitEthernet 1/1 group 2
 fvrp core
E1200-1#
```

E1200-2 Configuration

```
E1200-2#sh running-config fvrp
!
protocol fvrp
 fvrp control-vlan 10
 fvrp core region X
 no disable

E1200-2#sh running-config interface gi 5/2
!
interface GigabitEthernet 5/2
 no ip address
 switchport
 no shutdown
 fvrp uplink
E1200-2#sh running-config interface gi 5/0
!
interface GigabitEthernet 5/0
 no ip address
 switchport
 no shutdown
 fvrp access
 fvrp access region X
 fvrp aware
E1200-2#sh running-config interface gi 5/1
!
interface GigabitEthernet 5/1
 no ip address
 switchport
 no shutdown
 fvrp access
 fvrp access region X
 fvrp aware
E1200-2#sh running-config interface gi 5/22
!
interface GigabitEthernet 5/22
 no ip address
 no shutdown

E1200-2#sh running-config interface gi 5/23
!
interface GigabitEthernet 5/23
```

```
 no ip address
 no shutdown

E1200-2#sh running-config interface port-channel 10
!
interface port-channel 10
 no ip address
 switchport
 channel-member GigabitEthernet 5/22-23
 no shutdown

E1200-2#sh running-config interface vlan 10
!
interface VLAN 10
 no ip address
 tagged GigabitEthernet 5/0-2
 tagged Port-channel 10

E1200-2#sh running-config interface vlan 100
!
interface Vlan 100
 no ip address
 tagged GigabitEthernet 5/0-2
 no fvrp disable
 fvrp GigabitEthernet 5/0 group 1
 fvrp GigabitEthernet 5/1 group 2
 fvrp core

E1200-2#sh running-config interface vlan 200
!
interface Vlan 200
 no ip address
 tagged GigabitEthernet 5/0-2
 no fvrp disable
 fvrp priority 50
 fvrp GigabitEthernet 5/0 group 1
 fvrp GigabitEthernet 5/1 group 2
 fvrp core
E1200-2#


Configuration for access switches E1200-3 and 4
E1200-3:



E1200-3#sh running-config fvrp
!
protocol fvrp
fvrp control-vlan 10
fvrp core region X
no disable

E1200-3#sh running-config interface gi 5/0
!
interface GigabitEthernet 5/0
no ip address
switchport
fvrp access
no shutdown

E1200-3#sh running-config interface gi 5/1
!
interface GigabitEthernet 5/1
no ip address
```

```
switchport
fvrp access
no shutdown

E1200-3#sh running-config interface gi 5/2
!
interface GigabitEthernet 5/2
no ip address
switchport
no shutdown

E1200-3#sh running-config interface gi 5/3
!
interface GigabitEthernet 5/3
no ip address
switchport
no shutdown

E1200-3#sh running-config interface vlan 10
!
interface VLAN 10
no ip address
tagged GigabitEthernet 5/0-1

E1200-3#sh running-config interface vlan 100
!
interface VLAN 100
no ip address
tagged GigabitEthernet 5/0-3
no fvrp disable

E1200-3#sh running-config interface vlan 200
!
interface VLAN 200
no ip address
tagged GigabitEthernet 5/0-3
no fvrp disable
```

E1200-4

```
E1200-4#sh running-config fvrp
!
protocol fvrp
fvrp control-vlan 10
fvrp core region X
no disable

E1200-4#sh running-config interface gi 5/0
!
interface GigabitEthernet 5/0
no ip address
switchport
fvrp access
no shutdown

E1200-4#sh running-config interface gi 5/1
!
interface GigabitEthernet 5/1
no ip address
switchport
fvrp access
no shutdown
```

```
E1200-4#sh running-config interface gi 5/2
!
interface GigabitEthernet 5/2
no ip address
switchport
no shutdown

E1200-4#sh running-config interface gi 5/3
!
interface GigabitEthernet 5/3
no ip address
switchport
no shutdown

E1200-4#sh running-config interface vlan 10
!
interface VLAN 10
no ip address
tagged GigabitEthernet 5/0-1

E1200-4#sh running-config interface vlan 100
!
interface VLAN 100
no ip address
tagged GigabitEthernet 5/0-3
no fvrp disable
E1200-4#sh running-config interface vlan 200
!
interface VLAN 200
no ip address
tagged GigabitEthernet 5/0-3
 no fvrp disable
```

Viewing FVRP Configuration

To view the changes in FVRP operations for the VLAN, use the **show fvrp vlan** command (Figure 147) in the EXEC privilege mode.

Showing Configuration for FVRP Core Switches E1200-1 and E1200-2

```
E1200-1#sh fvrp vlan 100
     FVRP Vlan 100 Information
          FVRP Vlan Enabled
          FVRP Vlan Mode: Core
          FVRP Vlan State: Master
          FVRP Vlan priority: 10
          FVRP Vlan Hello time: 1
     Access Port 25 (GigabitEthernet 1/0) group 1 priority 128 is Forwarding
     Access Port 26 (GigabitEthernet 1/1) group 1 priority 128 is Forwarding
     Uplink Port 27 (GigabitEthernet 1/2) priority 128 is Forwarding
E1200-1#sh fvrp vlan 200
     FVRP Vlan 200 Information
          FVRP Vlan Enabled
          FVRP Vlan Mode: Core
          FVRP Vlan State: Standby
          FVRP Vlan priority: 128
          FVRP Vlan Hello time: 1
     Access Port 25 (GigabitEthernet 1/0) group 1 priority 128 is Blocking
     Access Port 26 (GigabitEthernet 1/1) group 1 priority 128 is Blocking
     Uplink Port 27 (GigabitEthernet 1/2) priority 128 is Blocking


E1200-2#sh fvrp vlan 100
     FVRP Vlan 100 Information
          FVRP Vlan Enabled
          FVRP Vlan Mode: Core
          FVRP Vlan State: StandBy
          FVRP Vlan priority: 128
          FVRP Vlan Hello time: 1
     Access Port 121 (GigabitEthernet 5/0) group 1 priority 128 is Blocking
     Access Port 122 (GigabitEthernet 5/1) group 1 priority 128 is Blocking
     Uplink Port 123 (GigabitEthernet 5/2) priority 128 is Blocking
E1200-2#sh fvrp vlan 200
     FVRP Vlan 200 Information
          FVRP Vlan Enabled
          FVRP Vlan Mode: Core
          FVRP Vlan State: Master
          FVRP Vlan priority: 50
          FVRP Vlan Hello time: 1
     Access Port 121 (GigabitEthernet 5/0) group 1 priority 128 is Forwarding
     Access Port 122 (GigabitEthernet 5/1) group 1 priority 128 is Forwarding
     Uplink Port 123 (GigabitEthernet 5/2) priority 128 is Forwarding
```

Figure 150 show fvrp vlan Command Examples

Chapter 16 IP Addressing

The E-Series software supports various IP addressing features. This chapter explains the basics of Domain Name Service (DNS), Address Resolution Protocol (ARP), and routing principles and their implementation in FTOS.

The E-Series software supports various IP addressing features:

- IP Addresses on page 287
- Directed Broadcast on page 292
- DHCP on page 292
- Resolution of Host Names on page 293
- ARP on page 296
- ICMP on page 298

Table 20 lists the defaults for the IP addressing features described in this chapter.

Table 20 IP Defaults

IP Feature	Default
DNS	Disabled
Directed Broadcast	Disabled
Proxy ARP	Enabled
ICMP Unreachable	Disabled
ICMP Redirect	Disabled

IP Addresses

FTOS supports IP version 4, as described in RFC 791. The software also supports classful routing and Variable Length Subnet Masks (VLSM). With VLSM one network can be can configured with different masks. Supernetting, which increases the number of subnets, is also supported. Subnetting is when a mask is added to the IP address to separate the network and host portions of the IP address.

At its most basic level, an IP address is 32-bits composed of network and host portions and represented in dotted decimal format. For example,

00001010110101100101011110000011

is represented as 10.214.87.131

For more information on IP Address, refer to RFC 791, *Internet Protoco*l.

Implementation Information

In FTOS, you can configure any IP address as a static route except IP addresses already assigned to interfaces.

Configuration Task List for IP Addresses

The following list includes the configuration tasks for IP addresses.

- assign IP addresses to an interface on page 288 (mandatory)
- configure static routes for the E-Series on page 290 (optional)
- configure static routes for the management interface on page 291 (optional)

For a complete listing of all commands related to IP addressing, refer to .

assign IP addresses to an interface

Assign primary and secondary IP addresses to physical or logical (for example, VLAN or port channel) interfaces to enable IP communication between the E-Series and hosts connected to that interface. In FTOS, you can assign one primary address and up to eight secondary IP addresses to each interface.

To assign an IP address to an interface, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**interface** *interface*	CONFIGURATION	Enter the keyword **interface** followed by the type of interface and slot/port information: • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a Loopback interface, enter the keyword **loopback** followed by a number from 0 to 16383. • For the Management interface on the RPM, enter the keyword **ManagementEthernet** followed by the slot/port information. The slot range is 0-1 and the port range is 0. • For a port channel interface, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN interface, enter the keyword **vlan** followed by a number from 1 to 4094.
2	**no shutdown**	INTERFACE	Enable the interface.
3	**ip address** *ip-address mask* [**secondary**]	INTERFACE	Configure a primary IP address and mask on the interface. • *ip-address mask:* IP address must be in dotted decimal format (A.B.C.D) and the mask must be in slash prefix-length format (/24). Add the keyword **secondary** if the IP address is the interface's backup IP address. You can configure up to eight secondary IP addresses.

To view the configuration, use the **show config** command (Figure 151) in the INTERFACE mode or **show ip interface** in the EXEC privilege mode (Figure 152).

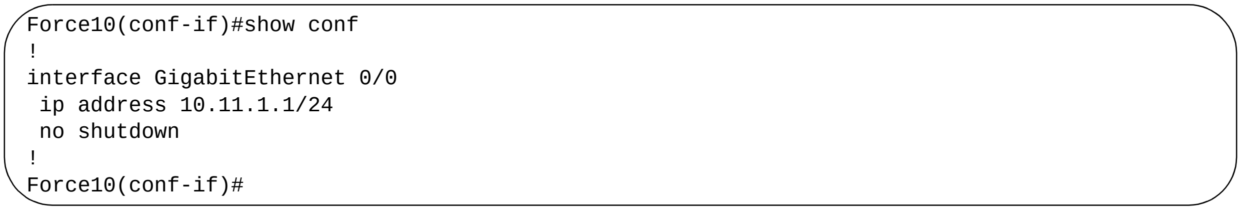

```
Force10(conf-if)#show conf
!
interface GigabitEthernet 0/0
 ip address 10.11.1.1/24
 no shutdown
!
Force10(conf-if)#
```

Figure 151 show config Command Example in the INTERFACE Mode

```
Force10#show ip int gi 0/8
GigabitEthernet 0/8 is up, line protocol is up
Internet address is 10.69.8.1/24
Broadcast address is 10.69.8.255
Address determined by config file
MTU is 1554 bytes
Inbound  access list is not set
Proxy ARP is enabled
Split Horizon is enabled
Poison Reverse is disabled
ICMP redirects are not sent
ICMP unreachables are not sent


Force10#
```

Figure 152 show ip interface Command Example

configure static routes for the E-Series

A static route is an IP address that is manually configured and not learned by a routing protocol, such as OSPF. Often static routes are used as backup routes in case other dynamically learned routes are unreachable.

To configure a static route, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ip route *ip-address mask* {*ip-address* | *interface* [*ip-address*]} [*distance*] [**permanent**] [**tag** *tag-value*]	CONFIGURATION	Configure a static IP address. Use the following required and optional parameters: • *ip-address*: Enter an address in dotted decimal format (A.B.C.D). • *mask*: Enter a mask in slash prefix-length format (/X). • *interface*: Enter an interface type followed by slot/port information. • *distance* range: 1 to 255 (optional). • **permanent:** Keep the static route in the routing table (if *interface* option is used) even if the interface with the route is disabled. (optional) • **tag** *tag-value* range: 1 to 4294967295. (optional)

You can enter as many static IP addresses as necessary.

To view the configured routes, use the **show ip route static** command.

```
Force10#show ip route static
       Destination        Gateway                    Dist/Metric Last Change
       -----------        -------                    ----------- -----------
  S    2.1.2.0/24         Direct, Nu 0                        0/0    00:02:30
  S    6.1.2.0/24         via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.2/32         via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.3/32         via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.4/32         via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.5/32         via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.6/32         via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.7/32         via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.8/32         via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.9/32         via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.10/32        via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.11/32        via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.12/32        via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.13/32        via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.14/32        via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.15/32        via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.16/32        via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    6.1.2.17/32        via 6.1.20.2, Te 5/0                1/0    00:02:30
  S    11.1.1.0/24        Direct, Nu 0                        0/0    00:02:30
                          Direct, Lo 0
--More--
```

Figure 153 show ip route static Command Example (partial)

The software installs a next hop that is on the directly connected subnet of current IP address on the interface (for example, if interface gig 0/0 is on 172.31.5.0 subnet, FTOS installs the static route).

The software also installs a next hop that is not on the directly connected subnet but which recursively resolves to a next hop on the interface's configured subnet. For example, if gig 0/0 has ip address on subnet 2.2.2.0 and if 172.31.5.43 recursively resolves to 2.2.2.0, FTOS installs the static route.

- When interface goes down, FTOS withdraws the route.
- When interface comes up, FTOS re-installs the route.
- When recursive resolution is “broken,” FTOS withdraws the route.
- When recursive resolution is satisfied, FTOS re-installs the route.

configure static routes for the management interface

When an IP address used by a protocol and a static management route exists for the same prefix, the protocol route takes precedence over the static management route.

To configure a static route for the management port, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
management route *ip-address mask* {*forwarding-router-address* | **ManagementEthernet** *slot/port*}	CONFIGURATION	Assign a static route to point to the Management interface or forwarding router.

To view the configured static routes for the Management port, use the **show ip management-route** command in the EXEC privilege mode.

```
Force10>show ip management-route

Destination        Gateway                        State
-----------        -------                        -----
1.1.1.0/24         172.31.1.250                   Active
172.16.1.0/24      172.31.1.250                   Active
172.31.1.0/24      ManagementEthernet 1/0         Connected

Force10>
```

Figure 154 show ip management-route Command Example

Directed Broadcast

By default, FTOS drops directed broadcast packets destined for an interface. This default setting provides some protection against Denial of Service (DOS) attacks.

To enable FTOS to receive directed broadcasts, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip directed-broadcast	INTERFACE	Enable directed broadcast.

To view the configuration, use the **show config** command in the INTERFACE mode.

DHCP

For protocols such as Dynamic Host Configuration Protocol (DHCP), relay devices respond to UDP broadcasts with information such as boot-up information. You can configure the IP address of a relay device (or the helper address) on an interface. Add multiple DHCP servers by entering the **ip helper-address** command multiple times. If multiple servers are defined, an incoming request is sent simultaneously to all configured servers and the reply is forwarded to the DHCP client.

FTOS uses standard DHCP ports, that is UDP ports 67 (server) and 68 (client) for DHCP relay services. It listens on port 67 and if it receives a broadcast, the software converts it to unicast, and forwards to it to the DHCP-server with source port=68 and destination port=67.

The server replies with source port=67, destination port=67 and FTOS forwards to the client with source port=67, destination port=68.

To configure a helper address, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip helper-address *ip-address*	INTERFACE	Configure the IP address of a relay device.

To view the configuration, use the **show ip interface** command (Figure 155) in EXEC privilege mode.

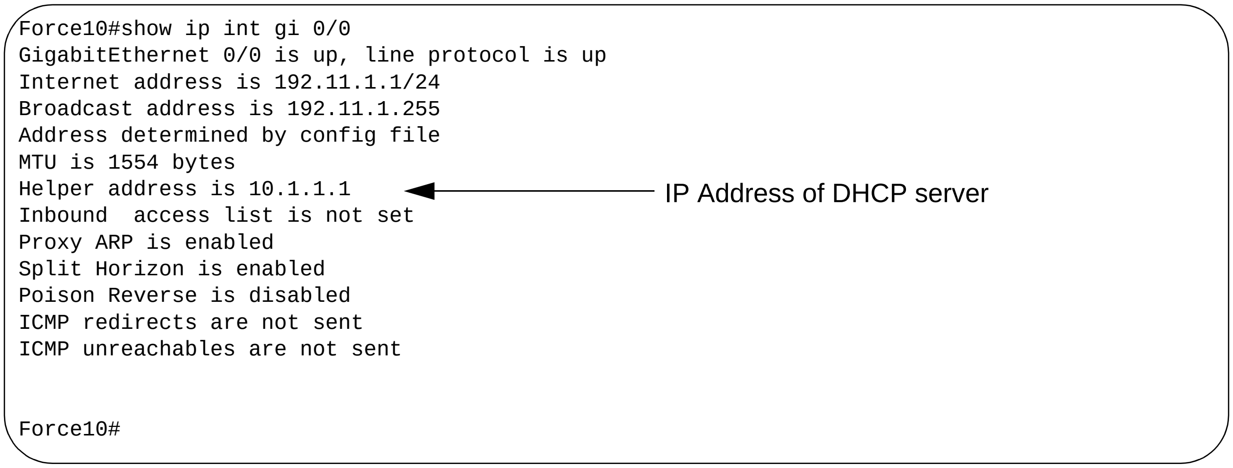

```
Force10#show ip int gi 0/0
GigabitEthernet 0/0 is up, line protocol is up
Internet address is 192.11.1.1/24
Broadcast address is 192.11.1.255
Address determined by config file
MTU is 1554 bytes
Helper address is 10.1.1.1
Inbound  access list is not set
Proxy ARP is enabled
Split Horizon is enabled
Poison Reverse is disabled
ICMP redirects are not sent
ICMP unreachables are not sent


Force10#
```

Figure 155 show ip interface Command Example

Resolution of Host Names

Domain Name Service (DNS) maps host names to IP addresses. This feature simplifies such commands as Telnet and FTP by allowing you to enter a name instead of an IP address.

Dynamic resolution of host names is disabled by default. Unless the feature is enabled, the system resolves only host names entered into the host table with the **ip host** command.

- enable dynamic resolution of host names on page 293
- specify local system domain and a list of domains on page 294
- DNS with traceroute on page 295

enable dynamic resolution of host names

By default, dynamic resolution of host names (DNS) is disabled.

To enable DNS, use the following commands in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ip domain-lookup	CONFIGURATION	Enable dynamic resolution of host names.
ip name-server *ip-address* [*ip-address2* ... *ip-address6*]	CONFIGURATION	Specify up to 6 name servers. The order you entered the servers determines the order of their use.

To view current bindings, use the **show hosts** command.

```
Force10>show host
Default domain is force10networks.com
Name/address lookup uses domain service
Name servers are not set
Host                     Flags      TTL    Type   Address
--------                 -----      ----   ----   -------
ks                       (perm, OK) -       IP    2.2.2.2
patch1                   (perm, OK) -       IP    192.68.69.2
tomm-3                   (perm, OK) -       IP    192.68.99.2
gxr                      (perm, OK) -       IP    192.71.18.2
f00-3                    (perm, OK) -       IP    192.71.23.1
Force10>
```

Figure 156 show hosts Command Example

To view the current configuration, use the **show running-config resolve** command.

specify local system domain and a list of domains

If you enter a partial domain, FTOS can search different domains to finish or fully qualify that partial domain. A fully qualified domain name (FQDN) is any name that is terminated with a period/dot. FTOS searches the host table first to resolve the partial domain. The host table contains both statically configured and dynamically learnt host and IP addresses. If the software cannot resolve the domain, it tries the domain name assigned to the local system. If that does not resolve the partial domain, the software searches the list of domains configured

To configure a domain name, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ip domain-name *name*	CONFIGURATION	Configure one domain name for the E-Series

To configure a list of domain names, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ip domain-list *name*	CONFIGURATION	Configure names to complete unqualified host names. Configure this command up to 6 times to specify a list of possible domain names. The software searches the domain names in the order they were configured until a match is found or the list is exhausted.

DNS with traceroute

To configure your switch to perform DNS with traceroute, follow the steps below in the CONFIGURATION mode.

Command Syntax	Command Mode	Purpose
ip domain-lookup	CONFIGURATION	Enable dynamic resolution of host names.
ip name-server *ip-address* [*ip-address2* ... *ip-address6*]	CONFIGURATION	Specify up to 6 name servers. The order you entered the servers determines the order of their use.
traceroute [*host* | *ip-address*]	CONFIGURATION	When you enter the traceroute command without specifying an IP address (Extended Traceroute), you are prompted for a target and source IP address, timeout in seconds (default is 5), a probe count (default is 3), minimum TTL (default is 1), maximum TTL (default is 30), and port number (default is 33434). To keep the default setting for those parameters, press the ENTER key.

Figure 157 is an example output of DNS using the traceroute command.

Figure 157 Traceroute command example

```
Force10#traceroute www.force10networks.com

Translating "www.force10networks.com"...domain server (10.11.0.1) [OK]
Type Ctrl-C to abort.

------------------------------------------------------------------------------------------
Tracing the route to www.force10networks.com (10.11.84.18), 30 hops max, 40 byte packets
------------------------------------------------------------------------------------------

 TTL Hostname             Probe1      Probe2      Probe3
  1  10.11.199.190         001.000 ms  001.000 ms  002.000 ms
  2  gwegress-sjc-02.force10networks.com (10.11.30.126) 005.000 ms  001.000 ms  001.000 ms
  3  fw-sjc-01.force10networks.com (10.11.127.254) 000.000 ms  000.000 ms  000.000 ms
  4  www.force10networks.com (10.11.84.18) 000.000 ms  000.000 ms  000.000 ms
Force10#
```

ARP

FTOS uses two forms of address resolution: ARP and Proxy ARP.

Address Resolution Protocol (ARP) runs over Ethernet and enables endstations to learn the MAC addresses of neighbors on an IP network. Over time, FTOS creates a forwarding table mapping the MAC addresses to their corresponding IP address. This table is called the ARP Cache and dynamically learned addresses are removed after a defined period of time.

For more information on ARP, see RFC 826, *An Ethernet Address Resolution Protocol*.

In FTOS, Proxy ARP enables hosts with knowledge of the network to accept and forward packets from hosts that contain no knowledge of the network. Proxy ARP makes it possible for hosts to be ignorant of the network, including subnetting.

For more information on Proxy ARP, refer to RFC 925, *Multi-LAN Address Resolution*, and RFC 1027, *Using ARP to Implement Transparent Subnet Gateways*.

Configuration Task List for ARP

The following list includes configuration tasks for ARP:

- configure static ARP entries on page 296 (optional)
- enable Proxy ARP on page 297 (optional)
- clear ARP cache on page 297 (optional)

For a complete listing of all ARP-related commands, refer to .

configure static ARP entries

ARP dynamically maps the MAC and IP addresses, and while most network host support dynamic mapping, you can configure an ARP entry (called a static ARP) for the ARP cache.

To configure a static ARP entry, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
arp *ip-address mac-address interface*	CONFIGURATION	Configure an IP address and MAC address mapping for an interface. • *ip-address:* IP address in dotted decimal format (A.B.C.D). • *mac-address:* MAC address in nnnn.nnnn.nnnn format • *interface:* enter the interface type slot/port information.

These entries do not age and can only be removed manually. To remove a static ARP entry, use the **no arp** ***ip-address*** command syntax.

To view the static entries in the ARP cache, use the **show arp static** command (Figure 158) in the EXEC privilege mode.

```
Force10#show arp

Protocol    Address         Age(min)  Hardware Address    Interface  VLAN    CPU
--------------------------------------------------------------------------------
Internet    10.1.2.4             17   08:00:20:b7:bd:32   Ma 1/0      -      CP
Force10#
```

Figure 158 show arp static Command Example

enable Proxy ARP

By default, Proxy ARP is enabled. To disable Proxy ARP, use **no proxy-arp** command in the interface mode.

To re-enable Proxy ARP, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip proxy-arp	INTERFACE	Re-enable Proxy ARP.

To view if Proxy ARP is enabled on the interface, use the **show config** command in the INTERFACE mode. If it is not listed in the show config command output, it is enabled. Only nondefault information is displayed in the **show config** command output.

clear ARP cache

To clear the ARP cache of dynamically learnt ARP information, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Purpose
clear arp-cache [*interface* | *ip ip-address*] [**no-refresh**]	EXEC privilege	Clear the ARP caches for all interfaces or for a specific interface by entering the following information: • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a port channel interface, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN interface, enter the keyword **vlan** followed by a number between 1 and 4094. **ip** *ip-address* (OPTIONAL) Enter the keyword ip followed by the IP address of the ARP entry you wish to clear. **no-refresh** (OPTIONAL) Enter the keyword **no-refresh** to delete the ARP entry from CAM. Or use this option with *interface* or **ip** *ip-address* to specify which dynamic ARP entires you want to delete. **Note:** Transit traffic may not be forwarded during the period when deleted ARP entries are resolved again and re-installed in CAM. Use this option with extreme caution.

ICMP

For diagnostics, Internet Control Message Protocol (ICMP) provide routing information to end stations by choosing the best route (ICMP redirect messages) or determining if a router is reachable (ICMP Echo or Echo Reply). ICMP Error messages inform the router of problems in a particular packet. These messages are sent only on unicast traffic

Configuration Task List for ICMP

Use the following steps to configure ICMP:

- enable ICMP unreachable messages on page 298
- enable ICMP redirects on page 299

See the for a complete listing of all commands related to ICMP.

enable ICMP unreachable messages

By default ICMP unreachable messages are disabled. When enabled ICMP unreachable messages are created and sent out all interfaces. To disable ICMP unreachable messages, use the **no ip unreachable** command syntax.

To reenable the creation of ICMP unreachable messages on the interface, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip unreachable	INTERFACE	Set FTOS to create and send ICMP unreachable messages on the interface.

To view if ICMP unreachable messages are sent on the interface, use the **show config** command in the INTERFACE mode. If it is not listed in the **show config** command output, it is enabled. Only nondefault information is displayed in the **show config** command output.

enable ICMP redirects

C-Series	NO
E-Series	✓

Platform Specific Feature: enable ICMP redirects is supported on E-Series only.

By default, ICMP redirect messages is disabled. When enabled, ICMP redirect messages are created and sent out all interfaces. To disable ICMP redirect messages, use the **no ip redirect** command syntax.

To reenable the creation of ICMP redirect messages on the interface, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip redirect	INTERFACE	Set FTOS to create and send ICMP redirect messages on the interface.

To view if ICMP redirect messages are sent on the interface, use the **show config** command in the INTERFACE mode. If it is not listed in the **show config** command output, it is enabled. Only nondefault information is displayed in the **show config** command output.

Chapter 17 IP Access Control Lists, Prefix Lists, and Route-maps

IP access control lists (ACLs), IP prefix lists, and route maps enable you to filter traffic and manipulate routes into and out of the E-Series.

This chapter covers the following topics:

IP Access Control Lists

An ACL is a series of sequential filters that contain a matching criterion (examine IP, TCP, or UDP packets) and an action (permit or deny). The filters are processed in sequence so that if a packet does not match the criterion in the first filter, the second filter (if configured) is applied. When the packet matches a filter, the E-Series drops or forwards the packet based on the filter’s designated action. If the packet does not match any of the filters in the ACL, the packet is dropped (that is, implicit deny).

In the E-Series, you can create two different types of IP ACLs: standard or extended. A standard ACL filters packets based on the source IP packet. An extended ACL filters traffic based on the following criteria (for more information on ACL supported options see *FTOS Command Line Interface Reference*):

- IP protocol number
- Source IP address
- Destination IP address
- Source TCP port number

- Destination TCP port number
- Source UDP port number
- Destination UDP port number

For extended ACL TCP and UDP filters, you can match criteria on specific or ranges of TCP or UDP ports. For extended ACL TCP filters, you can also match criteria on established TCP sessions.

When creating an access list, the sequence of the filters is important. You have a choice of assigning sequence numbers to the filters as you enter them, or FTOS will assign numbers in the order the filters are created. The sequence numbers, whether configured or assigned by FTOS, are listed in the **show config** and **show ip accounting access-list** command display output.

Ingress and egress Hot Lock ACLs allow you to appending or deleting new rules into an existing ACL (already written into CAM) without disruption to traffic flow. Existing entries in CAM simply are shuffled to accommodate new entries. Hot Lock ACLs are enabled by default and support both standard and extended ACLs.

C-Series	NO
E-Series	✓

Platform Specific Feature: Egress Hot Lock is supported on E-Series only.

Implementation Information

In the E-Series, you can assign one IP ACL per interface. If an ACL is not assigned to any interface, it is not used by the software in any other capacity.

The number of entries allowed per ACL is hardware-dependent. Refer to your line card documentation for detailed specification on entries allowed per ACL.

For the following features if counters are enabled on rules that are already configured, when a new rule is either inserted or prepended, all the existing counters are reset:

- L2 Ingress Access list
- L3 Egress Access list
- L2 Egress Access list

C-Series	NO
E-Series	✓

Platform Specific Feature: L2 and L3 Egress Access Lists are supported on E-Series only.

If a rule is simply appended then the existing counters are not affected.

Note: IP ACLs are supported over VLANs in Version 6.2.1.1 and higher.

Configuration Task List for IP ACLs

To configure an ACL, use commands in the IP ACCESS LIST mode and the INTERFACE mode. The following list includes the configuration tasks for IP ACLs:

- configure a standard IP ACL on page 303 (mandatory)
- configure an extended IP ACL on page 305 (mandatory)

For a complete listing of all commands related to IP ACLs, refer to the *FTOS Command Line Interface Reference* document.

configure a standard IP ACL

A standard IP ACL uses the source IP address as its match criterion.

To configure a standard IP access list, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip access-list standard** *access-listname*	CONFIGURATION (conf-std-nacl)	Enter IP ACCESS LIST mode by naming a standard IP access list.
2	**seq** *sequence-number* {**deny** | **permit**} {*source* [*mask*] | **any** | **host** *ip-address*} [**count** [**byte**] | **log**] [**order**] [**monitor**]	IP ACCESS LIST	Configure a drop or forward filter. The parameters are: • **log** and **monitor** options are supported on E-Series only.

Note: When assigning sequence numbers to filters, keep in mind that you might need to insert a new filter. To prevent reconfiguring multiple filters, assign sequence numbers in multiples of five or another number.

To view ACLs configured on the E-Series, use the **show ip accounting access-list** command (Figure 159) in the EXEC privilege mode.

```
Force10#show ip accounting access ToOspf interface gig 1/6
Standard IP access list ToOspf
 seq 5 deny any
 seq 10 deny 10.2.0.0 /16
 seq 15 deny 10.3.0.0 /16
 seq 20 deny 10.4.0.0 /16
 seq 25 deny 10.5.0.0 /16
 seq 30 deny 10.6.0.0 /16
 seq 35 deny 10.7.0.0 /16
 seq 40 deny 10.8.0.0 /16
 seq 45 deny 10.9.0.0 /16
 seq 50 deny 10.10.0.0 /16
Force10#
```

Figure 159 show ip accounting access-list Command Example

Figure 160 illustrates how the **seq** command orders the filters according to the sequence number assigned. In the example, filter 25 was configured before filter 15, but the **show config** command displays the filters in the correct order.

```
Force10(config-std-nacl)#seq 25 deny ip host 10.5.0.0 any log
Force10(config-std-nacl)#seq 15 permit tcp 10.3.0.0 /16 any
Force10(config-std-nacl)#show config
!
ip access-list standard dilling
 seq 15 permit tcp 10.3.0.0/16 any
 seq 25 deny ip host 10.5.0.0 any log
Force10(config-std-nacl)#
```

Figure 160 seq Command Example

To delete a filter, use the **no seq** *sequence-number* command in the IP ACCESS LIST mode.

If you are creating a standard ACL with only one or two filters, you can let FTOS assign a sequence number based on the order in which the filters are configured. The software assigns filters in multiples of 5.

To configure a filter without a specified sequence number, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip access-list standard** *access-list-name*	CONFIGURATION	Create a standard IP ACL and assign it a unique name.
2	{**deny** | **permit**} {*source* [*mask*] | **any** | **host** *ip-address*} [**count** [**byte**] | **log**] [**order**] [**monitor**]	IP ACCESS LIST	Configure a drop or forward IP ACL filter. • **log** and **monitor** options are supported on E-Series only.

Figure 161 illustrates a standard IP ACL in which the sequence numbers were assigned by the E-Series software. The filters were assigned sequence numbers based on the order in which they were configured (for example, the first filter was given the lowest sequence number). The **show config** command in the IP ACCESS LIST mode displays the two filters with the sequence numbers 5 and 10.

```
Force10(config-route-map)#ip access standard kigali
Force10(config-std-nacl)#permit 10.1.0.0/16
Force10(config-std-nacl)#show config
!
ip access-list standard kigali
 seq 5 permit 10.1.0.0/16
Force10(config-std-nacl)#
```

Figure 161 Standard IP ACL Example

To view all configured IP ACLs, use the **show ip accounting access-list** command (Figure 162) in the EXEC privilege mode.

```
Force10#show ip accounting access example interface gig 4/12
Extended IP access list example
seq 10 deny tcp any any eq 111
 seq 15 deny udp any any eq 111
 seq 20 deny udp any any eq 2049
 seq 25 deny udp any any eq 31337
 seq 30 deny tcp any any range 12345 12346
 seq 35 permit udp host 10.21.126.225 10.4.5.0 /28
 seq 40 permit udp host 10.21.126.226 10.4.5.0 /28
 seq 45 permit udp 10.8.0.0 /16 10.50.188.118 /31 range 1812 1813
 seq 50 permit tcp 10.8.0.0 /16 10.50.188.118 /31 eq 49
 seq 55 permit udp 10.15.1.0 /24 10.50.188.118 /31 range 1812 1813
```

Figure 162 show ip accounting access-list Command Example

To delete a filter, enter the **show config** command in the IP ACCESS LIST mode and locate the sequence number of the filter you want to delete. Then use the **no seq** *sequence-number* command in the IP ACCESS LIST mode.

configure an extended IP ACL

Extended IP ACLs filter on source and destination IP addresses, IP host addresses, TCP addresses, TCP host addresses, UDP addresses, and UDP host addresses.

Since traffic passes through the filter in the order of the filter's sequence, you can configure the extended IP ACL by first entering the IP ACCESS LIST mode and then assigning a sequence number to the filter.

To create a filter for packets with a specified sequence number, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip access-list extended** *access-list-name*	CONFIGURATION (config-ext-nacl)	Enter the IP ACCESS LIST mode by creating an extended IP ACL.

Step	Command Syntax	Command Mode	Purpose
2	**seq** *sequence-number* {**deny** | **permit**} {*ip-protocol-number* | **icmp | ip | tcp | udp**} {*source mask* | **any** | **host** *ip-address*} {*destination mask* | **any** | **host** *ip-address*} [*operator port* [*port*]] [**count** [**byte**] | **log**] [**order**] [**monitor**]	IP ACCESS LIST	Configure a drop or forward filter. • **log** and **monitor** options are supported on E-Series only.

To create a filter for TCP packets with a specified sequence number, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip access-list extended** *access-list-name*	CONFIGURATION	Create an extended IP ACL and assign it a unique name.
2	**seq** *sequence-number* {**deny** | **permit**} **tcp** {*source mask* | **any** | **host** *ip-address*}} [**count** [**byte**] | **log**] [**order**] [**monitor**]	IP ACCESS LIST	Configure an extended IP ACL filter for TCP packets. • **log** and **monitor** options are supported on E-Series only.

To create a filter for UDP packets with a specified sequence number, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip access-list extended** *access-list-name*	CONFIGURATION	Create a extended IP ACL and assign it a unique name.
2	**seq** *sequence-number* {**deny** | **permit**} {*ip-protocol-number* **udp**} {*source mask* | **any** | **host** *ip-address*} {*destination mask* | **any** | **host** *ip-address*} [*operator port* [*port*]] [**count** [**byte**] | **log**] [**order**] [**monitor**]	IP ACCESS LIST	Configure an extended IP ACL filter for UDP packets. • **log** and **monitor** options are supported on E-Series only.

When you create the filters with a specific sequence number, you can create the filters in any order and the filters are placed in the correct order.

Note: When assigning sequence numbers to filters, keep in mind that you might need to insert a new filter. To prevent reconfiguring multiple filters, assign sequence numbers in multiples of five or another number.

Figure 163 illustrates how the **seq** command orders the filters according to the sequence number assigned. In the example, filter 15 was configured before filter 5, but the **show config** command displays the filters in the correct order.

```
Force10(config-ext-nacl)#seq 15 deny ip host 112.45.0.0 any log
Force10(config-ext-nacl)#seq 5 permit tcp 12.1.3.45 0.0.255.255 any
Force10(config-ext-nacl)#show confi
!
ip access-list extended dilling
 seq 5 permit tcp 12.1.0.0 0.0.255.255 any
 seq 15 deny ip host 112.45.0.0 any log
Force10(config-ext-nacl)#
```

Figure 163 Extended IP ACL Using seq Command Example

If you are creating an extended ACL with only one or two filters, you can let FTOS assign a sequence number based on the order in which the filters are configured. FTOS assigns filters in multiples of 5.

To configure a filter for an extended IP ACL without a specified sequence number, use any or all of the following commands in the IP ACCESS LIST mode:

Command Syntax	Command Mode	Purpose
{**deny** | **permit**} {*source mask* | **any** | **host** *ip-address*} [**count** [**byte**] | **log**] [**order**] [**monitor**]	IP ACCESS LIST	Configure a deny or permit filter to examine IP packets. • **dlog** and **monitor** options are supported on E-Series only.
{**deny** | **permit**} **tcp** {*source mask*] | **any** | **host** *ip-address*}} [**count** [**byte**] | **log**] [**order**] [**monitor**]	IP ACCESS LIST	Configure a deny or permit filter to examine TCP packets. • **log** and **monitor** options are supported on E-Series only.
{**deny** | **permit**}**udp** {*source mask* | **any** | **host** *ip-address*}} [**count** [**byte**] | **log**] [**order**] [**monitor**]	IP ACCESS LIST	Configure a deny or permit filter to examine UDP packets. • **log** and **monitor** options are supported on E-Series only.

Figure 164 illustrates an extended IP ACL in which the sequence numbers were assigned by the E-Series software. The filters were assigned sequence numbers based on the order in which they were configured (for example, the first filter was given the lowest sequence number). The **show config** command in the IP ACCESS LIST mode displays the two filters with the sequence numbers 5 and 10.

```
Force10(config-ext-nacl)#deny tcp host 123.55.34.0 any
Force10(config-ext-nacl)#permit udp 154.44.123.34 0.0.255.255 host 34.6.0.0
Force10(config-ext-nacl)#show config
!
ip access-list extended nimule
 seq 5 deny tcp host 123.55.34.0 any
 seq 10 permit udp 154.44.0.0 0.0.255.255 host 34.6.0.0
Force10(config-ext-nacl)#
```

Figure 164 Extended IP ACL Example

To view all configured IP ACLs and the number of packets processed through the ACL, use the **show ip accounting access-list** command (Figure 162) in the EXEC privilege mode.

Established Flag

The **est** (established) flag is deprecated for Terascale series line cards.The flag is only available on legacy Etherscale linecards. Employ the **ack** and **rst** flags in their stead to achieve the same functionality.

To obtain the functionality of **est,** use the following ACLs:

- permit tcp any any rst
- permit tcp any any ack

Configuring Layer 2 and Layer 3 ACLs on an Interface

Both Layer 2 and Layer 3 ACLs may be configured on an interface in Layer 2 mode. If both L2 and L3 ACLs are applied to an interface, the following rules apply:

- The packets routed by Force10 are governed by the L3 ACL only, since they are not filtered against an L2 ACL.
- The packets switched by Force10 are first filtered by an L3 ACL, then by an L2 ACL.
- When packets are switched by Force10, the egress L3 ACL does not filter the packet.

For the following features if counters are enabled on rules that have already been configured, and when a new rule is either inserted or prepended, all the existing counters will be reset:

- L2 Ingress Access list
- L3 Egress Access list

- L2 Egress Access list

C-Series	NO
E-Series	✓

Platform Specific Feature: L2 and L3 Egress Access Lists are supported on E-Series only.

If a rule is simply appended, existing counters are not affected.

Please see the table on expected behavior:

Table 21 L2 and L3 ACL Filtering on Switched Packets

L2 ACL Behavior	L3 ACL Behavior	Decision on Targeted Traffic
Deny	Deny	Denied by L3 ACL
Deny	Permit	Permitted by L3 ACL
Permit	Deny	Denied by L2 ACL
Permit	Permit	Permitted by L2 ACL

Note: If an interface is configured as a “vlan-stack access” port, the packets are filtered by an L2 ACL only. The L3 ACL applied to such a port does not affect traffic. That is, existing rules for other features (such as trace-list, PBR, and QoS) are applied accordingly to the permitted traffic.

For information on Layer 2 or MAC ACLs, refer to MAC Addressing and MAC Access Lists on page 169.

Assign an IP ACL to an Interface

To pass traffic through a configured IP ACL, you must assign that ACL to a physical or port channel interface. The IP ACL is applied to all traffic entering a physical or port channel interface and the traffic is either forwarded or dropped depending on the criteria and actions specified in the ACL.

The same ACL may be applied to different interfaces and that changes its functionality. For example, you can take ACL "ABCD", and apply it using the **in** keyword and it becomes an ingress access list. If you apply the same ACL using the **out** keyword, it becomes an egress access list. If you apply the same ACL to the loopback interface, it becomes a loopback access list.

This chapter covers the following topics:

- Configuring Ingress ACLs on page 311
- Configuring Egress ACLs on page 312
- Configuring ACLs to Loopback on page 313

For more information on Layer-3 interfaces, refer to Chapter 9, Interfaces, on page 197.

To apply an IP ACL (standard or extended) to a physical or port channel interface, use these commands in the following sequence in the INTERFACE mode:

Step	Command Syntax	Command Mode	Purpose
1	**interface** *interface slot/port*	CONFIGURATION	Enter the interface number.
2	**ip address** *ip-address*	INTERFACE	Configure an IP address for the interface, placing it in Layer-3 mode.
3	**ip access-group** *access-list-name* {**in** | **out**} [**implicit-permit**] [**vlan** *vlan-range*]	INTERFACE	Apply an IP ACL to traffic entering or exiting an interface. • **out:** configure the ACL to filter outgoing traffic. This keyword is supported only on E-Series. **Note:** The number of entries allowed per ACL is hardware-dependent. Refer to your line card documentation for detailed specification on entries allowed per ACL.
4	**ip access-list [standard | extended]** *name*	INTERFACE	Apply rules to the new ACL.

To view which IP ACL is applied to an interface, use the **show config** command (Figure 165) in the INTERFACE mode or the **show running-config** command in the EXEC mode.

```
Force10(conf-if)#show conf
!
interface GigabitEthernet 0/0
 ip address 10.2.1.100 255.255.255.0
 ip access-group nimule in
 no shutdown
Force10(conf-if)#
```

Figure 165 show config Command in the INTERFACE Mode

Use only Standard ACLs in the **access-class** command to filter traffic on Telnet sessions.

Counting ACL Hits

You can view the number of packets matching the ACL by using the **count** option when creating ACL entries. E-Series supports packet and byte counts simultaneously. C-Series supports only one at any given time.

To view the number of packets match an ACL that is applied to an interface:

Step	Task
1	Create an ACL that uses rules with the count option. See Configuration Task List for IP ACLs on page 303

Step	Task
2	Apply the ACL as an inbound or outbound ACL on an interface. See Assign an IP ACL to an Interface on page 309
3	View the number of packets matching the ACL using the **show ip accounting access-list** from EXEC Privilege mode.

Configuring Ingress ACLs

Ingress ACLs are applied to interfaces and to traffic entering the system.These system-wide ACLs eliminate the need to apply ACLs onto each interface and achieves the same results. By localizing target traffic, it is a simpler implementation.

To create an ingress ACLs, use the **ip access-group** command (Figure 159) in the EXEC privilege mode. This example also shows applying the ACL, applying rules to the newly created access group, and viewing the access list:

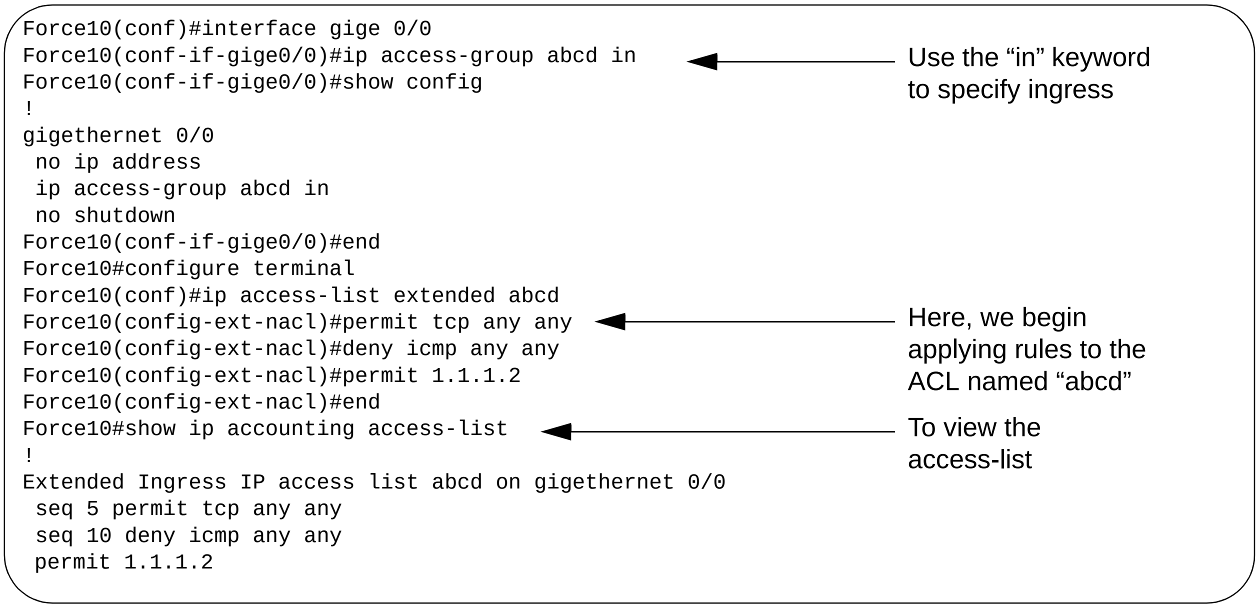

```
Force10(conf)#interface gige 0/0
Force10(conf-if-gige0/0)#ip access-group abcd in
Force10(conf-if-gige0/0)#show config
!
gigethernet 0/0
 no ip address
 ip access-group abcd in
 no shutdown
Force10(conf-if-gige0/0)#end
Force10#configure terminal
Force10(conf)#ip access-list extended abcd
Force10(config-ext-nacl)#permit tcp any any
Force10(config-ext-nacl)#deny icmp any any
Force10(config-ext-nacl)#permit 1.1.1.2
Force10(config-ext-nacl)#end
Force10#show ip accounting access-list
!
Extended Ingress IP access list abcd on gigethernet 0/0
 seq 5 permit tcp any any
 seq 10 deny icmp any any
 permit 1.1.1.2
```

Use the “in” keyword to specify ingress

Here, we begin applying rules to the ACL named “abcd”

To view the access-list

Figure 166 Creating an Ingress ACL Example

Configuring Egress ACLs

C-Series	NO
E-Series	✓

Platform Specific Feature: Configuring Egress ACLs is supported on E-Series only.

Egress ACLs are applied to line cards and affect the traffic leaving the system. Configuring egress ACLs onto physical interfaces protects the system infrustructure from attack—malicious and incidental—by explictly allowing only authorized traffic.These system-wide ACLs eliminate the need to apply ACLs onto each interface and achieves the same results. By localizing target traffic, it is a simpler implementation.

Packets originated from the system, are not filtered by egress ACLs. This means if you initiate a ping session from the system, and apply an egress ACL to block this type of traffic on the interface, the ACL does not affect that ping traffic.

An egress ACL is used when users would like to restrict egress traffic. For example, when a DOS attack traffic is isolated to one particular interface, the user can apply an egress ACL to block that particular flow from exiting the box, thereby protecting downstream devices.

To create an egress ACLs, use the **ip access-group** command (Figure 159) in the EXEC privilege mode. This example also shows viewing the configuration, applying rules to the newly created access group, and viewing the access list:

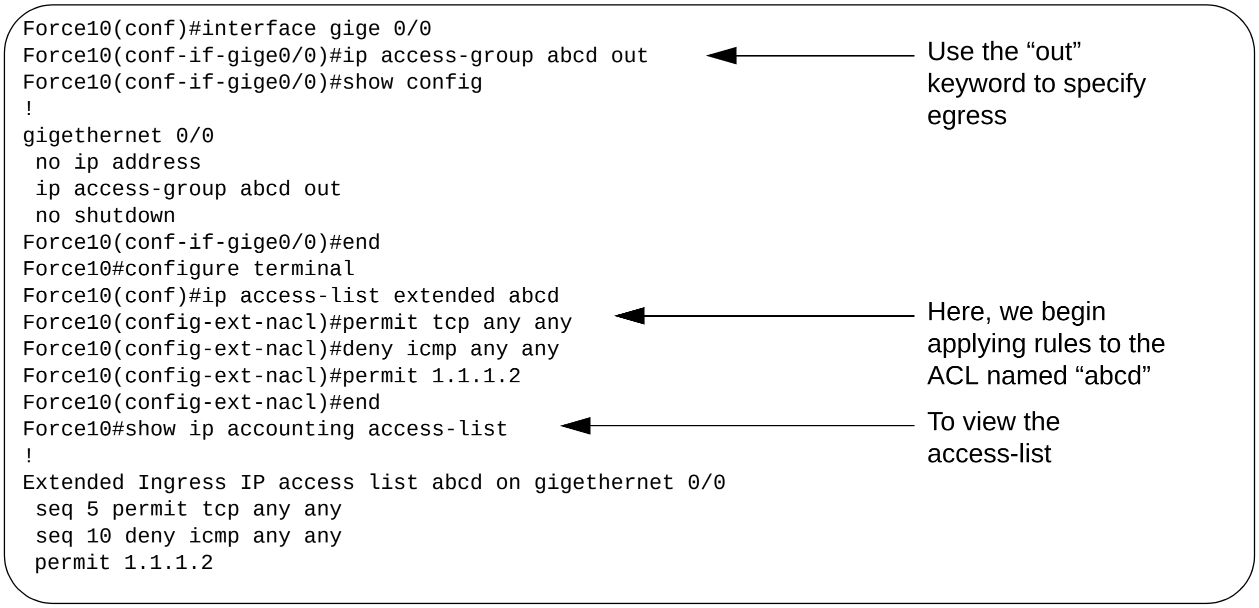

```
Force10(conf)#interface gige 0/0
Force10(conf-if-gige0/0)#ip access-group abcd out
Force10(conf-if-gige0/0)#show config
!
gigethernet 0/0
 no ip address
 ip access-group abcd out
 no shutdown
Force10(conf-if-gige0/0)#end
Force10#configure terminal
Force10(conf)#ip access-list extended abcd
Force10(config-ext-nacl)#permit tcp any any
Force10(config-ext-nacl)#deny icmp any any
Force10(config-ext-nacl)#permit 1.1.1.2
Force10(config-ext-nacl)#end
Force10#show ip accounting access-list
!
Extended Ingress IP access list abcd on gigethernet 0/0
 seq 5 permit tcp any any
 seq 10 deny icmp any any
 permit 1.1.1.2
```

Use the “out” keyword to specify egress

Here, we begin applying rules to the ACL named “abcd”

To view the access-list

Figure 167 Creating an Egress ACL Example

Configuring ACLs to Loopback

C-Series	NO
E-Series	✓

Platform Specific Feature: Configuring ACLs to Loopback is supported on E-Series only.

Configuring ACLs onto the CPU for loopback protects the system infrustructure from attack—malicious and incidental—by explictly allowing only authorized traffic.

The ACLs on loopback are applied only to the CPU on the RPM—this eliminates the need to apply specific ACLs onto all ingress interfaces and achieves the same results. By localizing target traffic, it is a simpler implementation.

The ACLs target and handle Layer 3 traffic destined to terminate on the system including routing protocols, remote access, SNMP, ICMP, and etc. Effective filtering of L3 traffic from L3 routers reduces the risk of attack.

Note: Loopback ACLs are supported only on ingress traffic.

See also Loopback Interfaces on page 211.

Applying an ACL to Loopback

To apply an ACL (standard or extended) for loopback, use these commands in the following sequence in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**interface loopback 0**	CONFIGURATION	Only loopback 0 is supported for loopback ACL.
2	**ip access-group** *name* **in**	CONFIGURATION	Apply an ACL to traffic entering loopback. • **in:** configure the ACL to filter incoming traffic **Note:** ACLs for loopback can only be applied to incoming traffic.
3	ip access-list [standard | extended] *name*		Apply rules to the new ACL.

To apply ACLs on loopback, use the **ip access-group** command (Figure 159) in the EXEC privilege mode. This example also shows viewing the configuration, applying rules to the newly created access group, and viewing the access list:

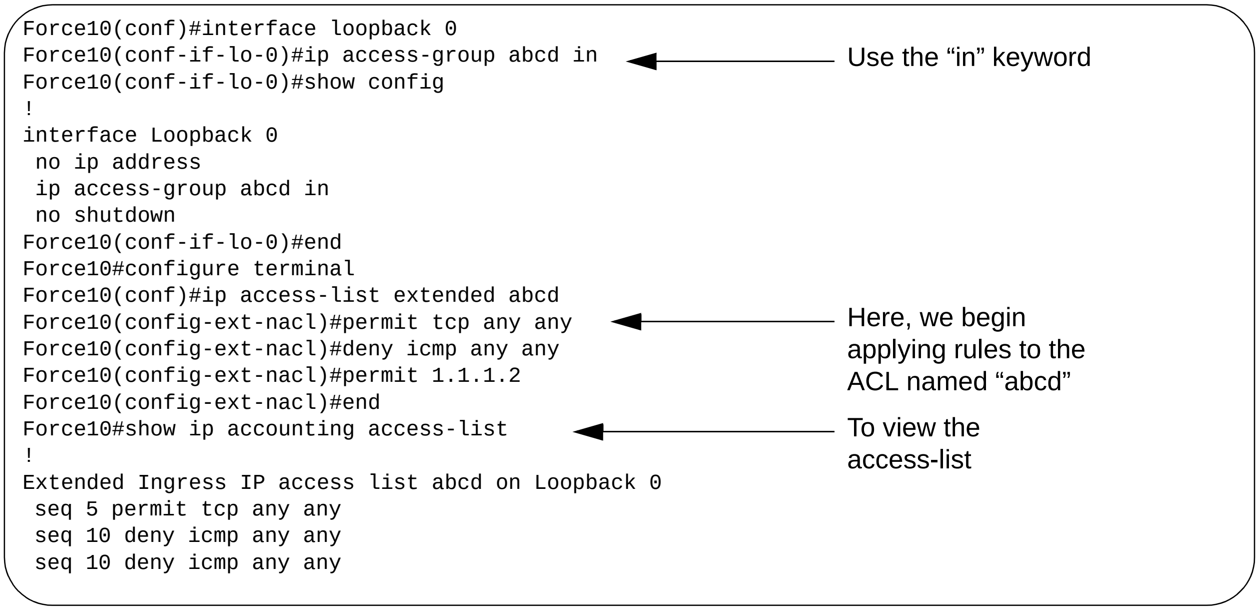

```
Force10(conf)#interface loopback 0
Force10(conf-if-lo-0)#ip access-group abcd in
Force10(conf-if-lo-0)#show config
!
interface Loopback 0
 no ip address
 ip access-group abcd in
 no shutdown
Force10(conf-if-lo-0)#end
Force10#configure terminal
Force10(conf)#ip access-list extended abcd
Force10(config-ext-nacl)#permit tcp any any
Force10(config-ext-nacl)#deny icmp any any
Force10(config-ext-nacl)#permit 1.1.1.2
Force10(config-ext-nacl)#end
Force10#show ip accounting access-list
!
Extended Ingress IP access list abcd on Loopback 0
 seq 5 permit tcp any any
 seq 10 deny icmp any any
 seq 10 deny icmp any any
```

Use the “in” keyword

Here, we begin applying rules to the ACL named “abcd”

To view the access-list

Figure 168 Applying an ACL to Loopback Example

IP Prefix Lists

IP prefix lists control routing policy. An IP prefix list is a series of sequential filters that contain a matching criterion (examine IP route prefix) and an action (permit or deny) to process routes. The filters are processed in sequence so that if a route prefix does not match the criterion in the first filter, the second filter (if configured) is applied. When the route prefix matches a filter, FTOS drops or forwards the packet based on the filter’s designated action. If the route prefix does not match any of the filters in the prefix list, the route is dropped (that is, implicit deny).

A route prefix is an IP address pattern that matches on bits within the IP address. The format of a route prefix is A.B.C.D/X where A.B.C.D is a dotted-decimal address and /X is the number of bits that should be matched of the dotted decimal address. For example, in 112.24.0.0/16, the first 16 bits of the address 112.24.0.0 match all addresses between 112.24.0.0 to 112.24.255.255.

Below are some examples that permit or deny filters for specific routes using the **le** and **ge** parameters, where x.x.x.x/x represents a route prefix:

- To deny only /8 prefixes, enter `deny x.x.x.x/x ge 8 le 8`
- To permit routes with the mask greater than /8 but less than /12, enter `permit x.x.x.x/x ge 8 le 12`
- To deny routes with a mask less than /24, enter `deny x.x.x.x/x le 24`
- To permit routes with a mask greater than /20, enter `permit x.x.x.x/x ge 20`

The following rules apply to prefix lists:

- A prefix list without any permit or deny filters allows all routes.
- An “implicit deny” is assumed (that is, the route is dropped) for all route prefixes that do not match a permit or deny filter in a configured prefix list.
- Once a route matches a filter, the filter’s action is applied. No additional filters are applied to the route.

Implementation Information

In FTOS, prefix lists are used in processing routes for routing protocols (for example, RIP, OSPF, and BGP).

Configuration Task List for Prefix Lists

To configure a prefix list, you must use commands in the PREFIX LIST, the ROUTER RIP, ROUTER OSPF, and ROUTER BGP modes. Basically, you create the prefix list in the PREFIX LIST mode, and assign that list to commands in the ROUTER RIP, ROUTER OSPF and ROUTER BGP modes.

The following list includes the configuration tasks for prefix lists:

- configure a prefix list on page 315 (mandatory)
- use a prefix list for route redistribution on page 318 (mandatory)

For a complete listing of all commands related to prefix lists, refer to .

configure a prefix list

To configure a prefix list, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip prefix-list** *prefix-name*	CONFIGURATION	Create a prefix list and assign it a unique name. You are in the PREFIX LIST mode.
2	**seq** *sequence-number* {**deny** | **permit**} *ip-prefix* [**ge** *min-prefix-length*] [**le** *max-prefix-length*]	PREFIX LIST	Create a prefix list with a sequence number and a deny or permit action. The optional parameters are: • **ge** *min-prefix-length:* is the minimum prefix length to be matched (0 to 32). • **le** *max-prefix-length:* is the maximum prefix length to be matched (0 to 32).

If you want to forward all routes that do not match the prefix list criteria, you must configure a prefix list filter to permit all routes (**permit 0.0.0.0/0 le 32**). The “permit all” filter should be the last filter in your prefix list. To permit the default route only, enter **permit 0.0.0.0/0**.

Figure 169 illustrates how the **seq** command orders the filters according to the sequence number assigned. In the example, filter 20 was configured before filter 15 and 12, but the **show config** command displays the filters in the correct order.

```
Force10(conf-nprefixl)#seq 20 permit 0.0.0.0/0 le 32
Force10(conf-nprefixl)#seq 12 deny 134.23.0.0 /16
Force10(conf-nprefixl)#seq 15 deny 120.23.14.0 /8 le 16
Force10(conf-nprefixl)#show config
!
ip prefix-list juba
 seq 12 deny 134.23.0.0/16
 seq 15 deny 120.0.0.0/8 le 16
 seq 20 permit 0.0.0.0/0 le 32
Force10(conf-nprefixl)#
```

Figure 169 seq Command Example

Note the last line in the prefix list Juba contains a “permit all” statement. By including this line in a prefix list, you specify that all routes not matching any criteria in the prefix list are forwarded.

To delete a filter, use the **no seq** *sequence-number* command in the PREFIX LIST mode.

If you are creating a standard prefix list with only one or two filters, you can let the E-Series software assign a sequence number based on the order in which the filters are configured. The E-Series software assigns filters in multiples of five.

To configure a filter without a specified sequence number, use these commands in the following sequence starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip prefix-list** *prefix-name*	CONFIGURATION	Create a prefix list and assign it a unique name.
2	{**deny** | **permit**} *ip-prefix* [**ge** *min-prefix-length*] [**le** *max-prefix-length*]	PREFIX LIST	Create a prefix list filter with a deny or permit action. The optional parameters are: • **ge** *min-prefix-length:* is the minimum prefix length to be matched (0 to 32). • **le** *max-prefix-length:* is the maximum prefix length to be matched (0 to 32).

Figure 170 illustrates a prefix list in which the sequence numbers were assigned by the E-Series software. The filters were assigned sequence numbers based on the order in which they were configured (for example, the first filter was given the lowest sequence number). The **show config** command in the PREFIX LIST mode displays the two filters with the sequence numbers 5 and 10.

```
Force10(conf-nprefixl)#permit 123.23.0.0 /16
Force10(conf-nprefixl)#deny 133.24.56.0 /8
Force10(conf-nprefixl)#show conf
!
ip prefix-list awe
 seq 5 permit 123.23.0.0/16
 seq 10 deny 133.0.0.0/8
Force10(conf-nprefixl)#
```

Figure 170 Prefix List Example

To delete a filter, enter the **show config** command in the PREFIX LIST mode and locate the sequence number of the filter you want to delete; then use the **no seq** *sequence-number* command in the PREFIX LIST mode.

To view all configured prefix lists, use either of the following commands in the EXEC mode:

Command Syntax	Command Mode	Purpose
show ip prefix-list detail [*prefix-name*]	EXEC privilege	Show detailed information about configured Prefix lists.
show ip prefix-list summary [*prefix-name*]	EXEC privilege	Show a table of summarized information about configured Prefix lists.

```
Force10>show ip prefix detail
Prefix-list with the last deletion/insertion: filter_ospf
ip prefix-list filter_in:
count: 3, range entries: 3, sequences: 5 - 10
   seq 5 deny 1.102.0.0/16 le 32 (hit count: 0)
   seq 6 deny 2.1.0.0/16 ge 23 (hit count: 0)
   seq 10 permit 0.0.0.0/0 le 32 (hit count: 0)
ip prefix-list filter_ospf:
count: 4, range entries: 1, sequences: 5 - 10
   seq 5 deny 100.100.1.0/24 (hit count: 0)
   seq 6 deny 200.200.1.0/24 (hit count: 0)
   seq 7 deny 200.200.2.0/24 (hit count: 0)
   seq 10 permit 0.0.0.0/0 le 32 (hit count: 0)
Force10>
```

Figure 171 show ip prefix-list detail Command Example

```
Force10>show ip prefix summary
Prefix-list with the last deletion/insertion: filter_ospf
ip prefix-list filter_in:
count: 3, range entries: 3, sequences: 5 - 10
ip prefix-list filter_ospf:
count: 4, range entries: 1, sequences: 5 - 10
Force10>
```

Figure 172 show ip prefix-list summary Command Example

use a prefix list for route redistribution

To pass traffic through a configured prefix list, you must use the prefix list in a route redistribution command. The prefix list is applied to all traffic redistributed into the routing process and the traffic is either forwarded or dropped depending on the criteria and actions specified in the prefix list.

To apply a filter to routes in RIP (RIP is supported on E-Series only.), use either of the following commands in the ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
distribute-list *prefix-list-name* **in** [*interface*]	ROUTER RIP	Apply a configured prefix list to incoming routes. You can specify an interface. If you enter the name of a nonexistent prefix list, all routes are forwarded.
distribute-list *prefix-list-name* **out** [*interface* | **connected** | **static** | **ospf**]	ROUTER RIP	Apply a configured prefix list to outgoing routes. You can specify an interface or type of route. If you enter the name of a non-existent prefix list, all routes are forwarded.

To view the configuration, use the **show config** command in the ROUTER RIP mode (Figure 173) or the **show running-config rip** command in the EXEC mode.

```
Force10(conf-router_rip)#show config
!
router rip
 distribute-list prefix juba out
 network 10.0.0.0
Force10(conf-router_rip)#router ospf 34
```

Figure 173 show config Command in the ROUTER RIP Mode

To apply a filter to routes in OSPF, use either of the following commands in the ROUTER OSPF mode:

Command Syntax	Command Mode	Purpose
distribute-list *prefix-list-name* **in** [*interface*]	ROUTER OSPF	Apply a configured prefix list to incoming routes. You can specify an interface. If you enter the name of a non-existent prefix list, all routes are forwarded.
distribute-list *prefix-list-name* **out** [**connected** | **rip** | **static**]	ROUTER OSPF	Apply a configured prefix list to incoming routes. You can specify which type of routes are affected. If you enter the name of a non-existent prefix list, all routes are forwarded.

To view the configuration, use the **show config** command in the ROUTER OSPF mode (Figure 177) or the **show running-config ospf** command in the EXEC mode.

```
Force10(conf-router_ospf)#show config
!
router ospf 34
 network 10.2.1.1 255.255.255.255 area 0.0.0.1
 distribute-list prefix awe in
Force10(conf-router_ospf)#
```

Figure 174 show config Command Example in ROUTER OSPF Mode

ACL Resequencing

ACL Resequencing allows you to re-number the rules and remarks in an access or prefix list. The placement of rules within the list is critical because packets are matched against rules in sequential order. Use resquencing whenever there is no longer an opportunity to order new rules as desired using current numbering scheme.

For example, Table 22 contains some rules that are numbered in increments of 1. No new rules can be placed between these, so apply resequencing to create numbering space, as shown in Table 23. In the same example, apply resequencing if more than two rules must be placed between rules 7 and 10.

IPv4 and IPv6 ACLs and prefixes and MAC ACLs can be resequenced.

Note: ACL Resequencing does not affect the rules or remarks or the order in which they are applied. It merely renumbers them so that new rules can be placed within the list as desired.

Table 22 ACL Resequencing Example (Insert New Rules)

seq 5 permit any host 1.1.1.1
seq 6 permit any host 1.1.1.2
seq 7 permit any host 1.1.1.3
seq 10 permit any host 1.1.1.4

Table 23 ACL Resequencing Example (Resequenced)

seq 5 permit any host 1.1.1.1
seq 10 permit any host 1.1.1.2
seq 15 permit any host 1.1.1.3
seq 20 permit any host 1.1.1.4

Resequencing an ACL or Prefix List

Resequencing is available for IPv4 and IPv6 ACLs and prefix lists and MAC ACLs. To resequence an ACL or prefix list use the appropriate command in Table 24. You must specify the list name, starting number, and increment when using these commands.

C-Series	NO
E-Series	✓

Platform Specific Feature: IPv6 is supported on E-Series only.

Table 24 Resequencing ACLs and Prefix Lists

List	Command	Command Mode
IPv4, IPv6, or MAC ACL	**resequence access-list** {**ipv4** | **ipv6** | **mac**} {*access-list-name StartingSeqNum Step-to-Increment*}	Exec
IPv4 or IPv6 prefix-list	**resequence prefix-list** {**ipv4** | **ipv6**} {*prefix-list-name StartingSeqNum Step-to-Increment*}	Exec

Figure 175 shows the resequencing of an IPv4 access-list beginning with the number 2 and incrementing by 2.

Figure 175 Resequencing ACLs

```
Force10(config-ext-nacl)# show config
!
ip access-list extended test
remark 4 XYZ
remark 5 this remark corresponds to permit any host 1.1.1.1
seq 5 permit ip any host 1.1.1.1
remark 9 ABC
remark 10 this remark corresponds to permit ip any host 1.1.1.2
seq 10 permit ip any host 1.1.1.2
seq 15 permit ip any host 1.1.1.3
seq 20 permit ip any host 1.1.1.4
Force10# end
Force10# resequence access-list ipv4 test 2 2
Force10# show running-config acl
!
ip access-list extended test
remark 2 XYZ
remark 4 this remark corresponds to permit any host 1.1.1.1
seq 4 permit ip any host 1.1.1.1
remark 6 this remark has no corresponding rule
remark 8 this remark corresponds to permit ip any host 1.1.1.2
seq 8 permit ip any host 1.1.1.2
seq 10 permit ip any host 1.1.1.3
seq 12 permit ip any host 1.1.1.4
```

Remarks and rules that originally have the same sequence number have the same sequence number after the **resequence** command is applied. Remarks that do not have a corresponding rule will be incremented as as a rule. These two mechanisms allow remarks to retain their original position in the list.

For example, in Figure 176, remark 10 corresponds to rule 10 and as such they have the same number before and after the command is entered. Remark 4 is incremented as a rule, and all rules have retained their original positions.

Figure 176 Resequencing Remarks

```
Force10(config-ext-nacl)# show config
!
ip access-list extended test
remark 4 XYZ
remark 5 this remark corresponds to permit any host 1.1.1.1
seq 5 permit ip any host 1.1.1.1
remark 9 ABC
remark 10 this remark corresponds to permit ip any host 1.1.1.2
seq 10 permit ip any host 1.1.1.2
seq 15 permit ip any host 1.1.1.3
seq 20 permit ip any host 1.1.1.4
Force10# end
Force10# resequence access-list ipv4 test 2 2
Force10# show running-config acl
!
ip access-list extended test
remark 2 XYZ
remark 4 this remark corresponds to permit any host 1.1.1.1
seq 4 permit ip any host 1.1.1.1
remark 6 this remark has no corresponding rule
remark 8 this remark corresponds to permit ip any host 1.1.1.2
seq 8 permit ip any host 1.1.1.2
seq 10 permit ip any host 1.1.1.3
seq 12 permit ip any host 1.1.1.4
```

Route Maps

Like ACLs and prefix lists, route maps are composed of a series of commands that contain a matching criterion and an action, yet route maps can change the packets meeting the criterion. ACLs and prefix lists can only drop or forward the packet or traffic. Route maps process routes for route redistribution. For example, a route map can be called to filter only specific routes and to add a metric.

Route maps also have an “implicit deny.” Unlike ACLs and prefix lists, however, where the packet or traffic is dropped, in route maps, if a route does not match any of the route map conditions, the route is not redistributed.

Implementation Information

The FTOS implementation of route maps allows route maps with no match command or no set command. When there is no match command, all traffic matches the route map and the set command applies.

Important Points to Remember

- For route-maps with more than one match clause:

    - Two or more match clauses within the same route-map sequence have the *same* match commands (though the values are different), matching a packet against these clauses is a logical OR operation.
    - Two or more match clauses within the same route-map sequence have *different* match commands, matching a packet against these clauses is a logical AND operation.
- If no match is found in a route-map sequence, the process moves to the next route-map sequence until a match is found, or there are no more sequences.
- When a match is found, the packet is forwarded; no more route-map sequences are processed.
    - If a continue clause is included in the route-map sequence, the next or a specified route-map sequence is processed after a match is found.

Configuration Task List for Route Maps

You configure route maps in the ROUTE-MAP mode and apply them in various commands in the ROUTER RIP and ROUTER OSPF modes.

The following list includes the configuration tasks for route maps:

- create a route map on page 324 (mandatory)
- configure route map filters on page 326 (optional)
- configure a route map for route redistribution on page 329 (optional)
- configure a route map for route tagging on page 329 (optional)

create a route map

Route maps, ACLs, and prefix lists are similar in composition because all three contain filters, but route map filters are do not contain the permit and deny actions found in ACLs and prefix lists. Route map filters match certain routes and set or specify values.

To create a route map and enter the ROUTE-MAP mode, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
route-map *map-name* [**permit** | **deny**] [*sequence-number*]	CONFIGURATION	Create a route map and assign it a unique name. The optional **permit** and **deny** keywords are the action of the route map. The default is **permit**. The optional parameter **seq** allows you to assign a sequence number to the route map instance.

The default action is permit and the default sequence number starts at 10. When the keyword **deny** is used in configuring a route map, routes that meet the match filters are not redistributed.

To view the configuration, use the **show config** command in the ROUTE-MAP mode (Figure 177).

```
Force10(config-route-map)#show config
!
route-map dilling permit 10
Force10(config-route-map)#
```

Figure 177 show config Command Example in the ROUTE-MAP Mode

You can create multiple instances of this route map by using the sequence number option to place the route maps in the correct order. FTOS processes the route maps with the lowest sequence number first. When a configured route map is applied to a command, like **redistribute**, traffic passes through all instances of that route map until a match is found.

Figure 178 shows an example with two instances of a route map.

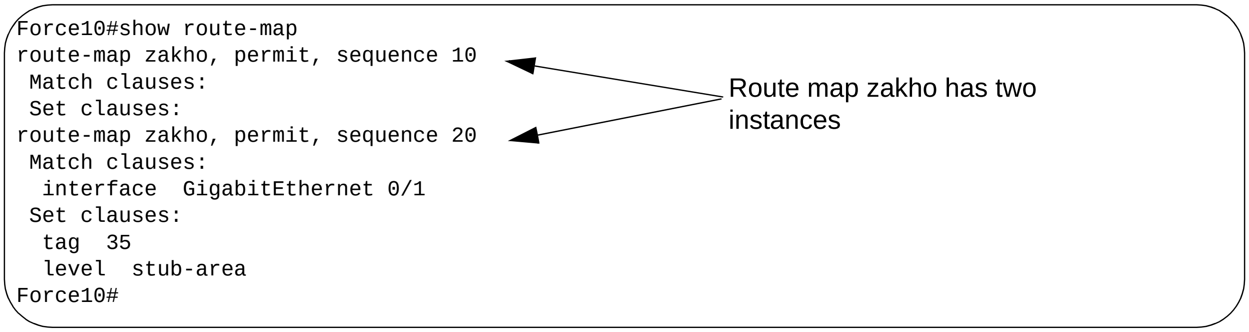

```
Force10#show route-map
route-map zakho, permit, sequence 10
 Match clauses:
 Set clauses:
route-map zakho, permit, sequence 20
 Match clauses:
  interface  GigabitEthernet 0/1
 Set clauses:
  tag  35
  level  stub-area
Force10#
```

Route map zakho has two instances

Figure 178 show route-map Command Example with Multiple Instances of a Route Map

To delete all instances of that route map, use the **no route-map** ***map-name*** command. To delete just one instance, add the sequence number to the command syntax (Figure 179).

```
Force10(conf)#no route-map zakho 10
Force10(conf)#end
Force10#show route-map
route-map zakho, permit, sequence 20
 Match clauses:
  interface  GigabitEthernet 0/1
 Set clauses:
  tag  35
  level  stub-area
Force10#
```

Figure 179 Example of Deleting One Instance of a Route Map

Figure 180 shows an example of a route map with multiple instances. The **show config** command displays only the configuration of the current route map instance. To view all instances of a specific route map, use the **show route-map** command.

```
Force10#show route-map dilling
route-map dilling, permit, sequence 10
 Match clauses:
 Set clauses:
route-map dilling, permit, sequence 15
 Match clauses:
  interface  Loopback 23
 Set clauses:
  tag  3444
Force10#
```

Figure 180 show route-map Command Example

To delete a route map, use the **no route-map** *map-name* command in the CONFIGURATION mode.

configure route map filters

Within the ROUTE-MAP mode, there are **match** and **set** commands. Basically, **match** commands search for a certain criterion in the routes and the **set** commands change the characteristics of those routes, either adding something or specifying a level.

When there are multiple match commands of the same parameter under one instance of route-map, then FTOS does a match between either of those match commands. If there are multiple match commands of different parameter, then FTOS does a match ONLY if there is a match among ALL match commands. The following example explains better:

example 1:

```
Force10(conf)#route-map force permit 10
Force10(config-route-map)#match tag 1000
Force10(config-route-map)#match tag 2000
Force10(config-route-map)#match tag 3000
```

In the above route-map, if a route has any of the tag value specified in the match commands, then there is a match.

example 2:

```
Force10(conf)#route-map force permit 10
Force10(config-route-map)#match tag 1000
Force10(config-route-map)#match metric 2000
```

In the above route-map, *only* if a route has *both* the characteristics mentioned in the route-map, it is matched. Explaining further, the route *must* have a tag value of 1000 *and* a metric value of 2000. Only then is there a match.

Also, if there are different instances of the same route-map, then it's sufficient if a permit match happens in *any* instance of that route-map. As an example:

```
Force10(conf)#route-map force permit 10
Force10(config-route-map)#match tag 1000

Force10(conf)#route-map force deny 20
Force10(config-route-map)#match tag 1000

Force10(conf)#route-map force deny 30
Force10(config-route-map)#match tag 1000
```

In the above route-map, instance 10 permits the route having a tag value of 1000 and instances 20 & 30 denies the route having a tag value of 1000. In the above scenario, FTOS scans all the instances of the route-map for any permit statement. If there is a match anywhere, the route is permitted, though other instances of the route-map denies it.

To configure match criterion for a route map, use any or all of the following commands in the ROUTE-MAP mode:

Command Syntax	Command Mode	Purpose
match as-path *as-path-name*	ROUTE-MAP	Match BGP routes with the same AS-PATH numbers.
match community *community-list-name* [**exact**]	ROUTE-MAP	Match BGP routes with COMMUNITY list attributes in their path.
match interface *interface*	ROUTE-MAP	Match routes whose next hop is a specific interface. The parameters are: • For a Fast Ethernet interface, enter the keyword **FastEthernet** followed by the slot/port information. • For a 1-Gigabit Ethernet interface, enter the keyword **gigabitEthernet** followed by the slot/port information. • For a loopback interface, enter the keyword **loopback** followed by a number between zero (0) and 16383. • For a port channel interface, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **tengigabitEthernet** followed by the slot/port information. • For a VLAN, enter the keyword vlan followed by a number from 1 to 4094.
match ip address *prefix-list-name*	ROUTE-MAP	Match destination routes specified in a prefix list.

Command Syntax	Command Mode	Purpose
match ip next-hop {*access-list-name* | **prefix-list** *prefix-list-name*}	ROUTE-MAP	Match next-hop routes specified in a prefix list.
match ip route-source {*access-list-name* | **prefix-list** *prefix-list-name*}	ROUTE-MAP	Match source routes specified in a prefix list.
match metric *metric-value*	ROUTE-MAP	Match routes with a specific value.
match origin {**egp** | **igp** | **incomplete**}	ROUTE-MAP	Match BGP routes based on the ORIGIN attribute.
match route-type {**external** [**type-1** | **type-2**] | **internal** | **level-1** | **level-2** | **local** }	ROUTE-MAP	Match routes specified as internal or external to OSPF, ISIS level-1, ISIS level-2, or locally generated.
match tag *tag-value*	ROUTE-MAP	Match routes with a specific tag.

To configure a set condition, use any or all of the following commands in the ROUTE-MAP mode:

Command Syntax	Command Mode	Purpose
set as-path prepend *as-number* [... *as-number*]	ROUTE-MAP	Add an AS-PATH number to the beginning of the AS-PATH
set automatic-tag	ROUTE-MAP	Generate a tag to be added to redistributed routes.
set level {**backbone** | **level-1** | **level-1-2** | **level-2** | **stub-area** }	ROUTE-MAP	Specify an OSPF area or ISIS level for redistributed routes.
set local-preference *value*	ROUTE-MAP	Specify a value for the BGP route’s LOCAL_PREF attribute.
set metric {**+** | **-** | *metric-value*}	ROUTE-MAP	Specify a value for redistributed routes.
set metric-type {**external** | **internal** | **type-1** | **type-2**}	ROUTE-MAP	Specify an OSPF or ISIS type for redistributed routes.
set next-hop *ip-address*	ROUTE-MAP	Assign an IP address as the route’s next hop.
set origin {**egp** | **igp** | **incomplete**}	ROUTE-MAP	Assign a BGP ORIGIN attribute.
set tag *tag-value*	ROUTE-MAP	Specify a tag for the redistributed routes.
set weight *value*	ROUTE-MAP	Specify a value as the route’s BGP weight.

Use these commands to create route map instances. There is no limit to the number of set and match commands per route map, but the convention is to keep the number of match and set filters in a route map low. **Set** commands do not require a corresponding **match** command.

configure a route map for route redistribution

Route maps on their own cannot affect traffic and must be included in different commands to affect routing traffic. To apply a route map to traffic on the E-Series, you must call or include that route map in a command such as the **redistribute** or **default-information originate** commands in OSPF, ISIS, and BGP.

Route redistribution occurs when FTOS learns the advertising routes from static or directly connected routes or another routing protocol. Different protocols assign different values to redistributed routes to identify either the routes and their origins. The metric value is the most common attribute that is changed to properly redistribute other routes into a routing protocol. Other attributes that can be changed include the metric type (for example, external and internal route types in OSPF) and route tag. Use the **redistribute** command in OSPF, RIP, ISIS, and BGP to set some of these attributes for routes that are redistributed into those protocols.

Route maps add to that redistribution capability by allowing you to match specific routes and set or change more attributes when redistributing those routes.

In Figure 181, the **redistribute** command calls the route map `staticospf` to redistribute only certain static routes into OSPF. According to the route map `staticospf`, only routes that have a next hop of Gigabitethernet interface 0/0 and that have a metric of 255 will be redistributed into the OSPF backbone area.

Note: When re-distributing routes using route-maps, the user must take care to create the route-map defined in the **redistribute** command under the routing protocol. If no route-map is created, then NO routes are redistributed.

```
router ospf 34
 default-information originate metric-type 1
 redistribute static metric 20 metric-type 2 tag 0 route-map staticospf
!
route-map staticospf permit 10
 match interface  GigabitEthernet 0/0
 match metric  255
 set level  backbone
```

Figure 181 Route Redistribution into OSPF Example

configure a route map for route tagging

One method for identifying routes from different routing protocols is to assign a tag to routes from that protocol. As the route enters a different routing domain, it is tagged and that tag is passed along with the route as it passes through different routing protocols. This tag can then be used when the route leaves a routing domain to redistribute those routes again.

In Figure 182, the **redistribute ospf** command with a route map is used in the ROUTER RIP mode to apply a tag of 34 to all internal OSPF routes that are redistributed into RIP.

```
!
router rip
 redistribute ospf 34 metric 1 route-map torip
!
route-map torip permit 10
 match route-type  internal
 set tag  34
!
```

Figure 182 Tagging OSPF Routes Entering a RIP Routing Domain

continue clause

Normally, when a match is found, set clauses are executed, and the packet is then forwarded; no more route-map modules are processed. If the **continue** command is configured at the end of a module, the next module (or a specified module) is processed even after a match is found. Figure 183 shows a continue clause at the end of a route-map module. In this example, if a match is found in the route-map “test” module 10, module 30 will be processed.

Note: If the continue clause is configured without specifying a module, the next sequential module is processed.

Figure 183 continue Command Example

```
!
route-map test permit 10
match commu comm-list1
set community 1:1 1:2 1:3
set as-path prepend 1 2 3 4 5
continue 30!
```

Chapter 18 High Availability

The Force10 Networks E-Series line cards, SFMs, and RPMs can be hotswapped.

This chapter covers the following topics:

- Online Insertion and Removal (OIR) on page 331
- RPM Redundancy on page 335
- High Availabilty on C-Series on page 342

Online Insertion and Removal (OIR)

In the E-Series, you add, replace, or remove a line card, the redundant SFM or the Standby RPM without interrupting the system. While the system is online, you can replace a blank filler panel with a line card or hot swap a line card of the same type.

This section covers the following topics:

- Line Cards on page 331
- SFMs on page 334
- Standby RPM on page 335

Line Cards

When you insert a line card into an empty slot in a system that is in operation, the software detects the line card type, and then writes the line card information into the running-config. The running-config file keeps the information, even if the line card is removed from the slot.

To better control traffic during a line card hotswap, shut down all interfaces on the line card, using the **shutdown** command.

You can pre-configure a line card slot and interfaces by using the **linecard** command. For an empty slot, enter the **linecard** command with the specific card-type. Once you have entered that information, you may configure the interfaces that would normally be found on that line card. Figure 184 displays a portion of the **show running-config** command output for a line card and interface that are not installed in the chassis. The first line of the screenshot informs you that the card type was configured for the slot, while the following lines contain the line card's interface configuration.

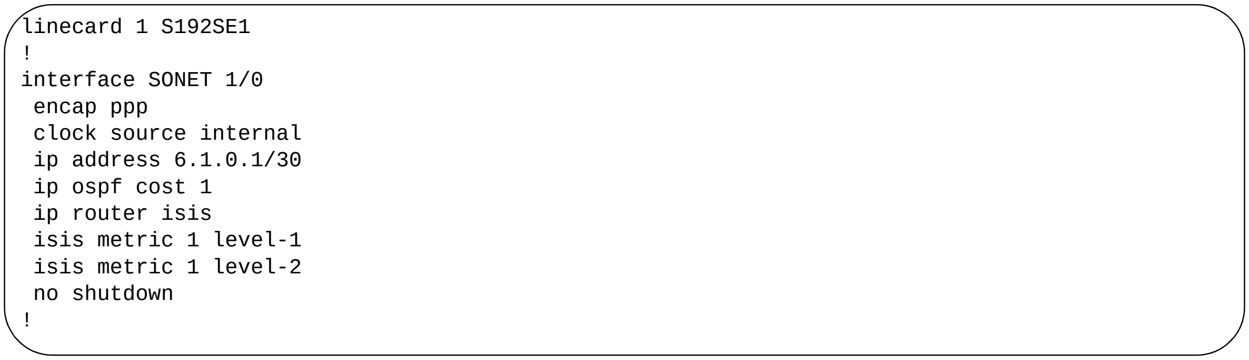

```
linecard 1 S192SE1
!
interface SONET 1/0
 encap ppp
 clock source internal
 ip address 6.1.0.1/30
 ip ospf cost 1
 ip router isis
 isis metric 1 level-1
 isis metric 1 level-2
 no shutdown
!
```

Figure 184 Example of a Pre-configured Line Card and Interface

If you swap line cards of different types (for example, a 24-port 1-Gigabit Ethernet line card for a 2-port 10-Gigabit Ethernet line card), you must change the running-config file to reflect the new line card type information after the line card is removed. If you do not change the line card configuration, FTOS reports a "card mismatch" as the line card status when the new card is installed.

To configure a different line card type, use the following command in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Usage
1	**no linecard** *number*	CONFIGURATION	Remove the old line card configuration. Configure the following parameter: • *number:* Enter a number for the slot number. After entering this command, remove the old line card and insert the new line card.
2	**linecard** *number card-type*	CONFIGURATION	Configure a slot with a new line card type. Use when inserting a line card type in an empty slot Configure the following parameters: • *number:* Enter a number for the slot number. *card-type:* Enter the card type.

Once the system recognizes a line card type or you configure the **linecard** command, the system requires that line card type to be installed in that slot.

Figure 185 is a **show linecard** command example and both the Required Type and Current Type fields must match for the system to correctly access the line card.

If a different line card is inserted, that line card status is “type mismatch.” To clear the “type mismatch” status and bring the line card on-line, use the **linecard** command to change the line card type to match the line card in the slot.

```
Force10>show linecard 2

--  Line card 2 --
Status         : online
Next Boot      : online
Required Type : S48SC2 - 2-port OC48c line card with SR optics (EC)
Current Type  : S48SC2 - 2-port OC48c line card with SR optics (EC)
Hardware Rev  : 1.0
Num Ports     : 2
Up Time       : 2 day, 1 hr, 20 min
FTOS Version  : 4.4.1.0
Jumbo Capable : no
Boot Flash Ver: A: 2.0.0.24     B: 2.0.0.26 [booted]
Memory Size   : 134217728 bytes
Temperature   : 57C
Power Status  : PEM0: absent or down    PEM1: up
Voltage       : ok
Serial Number : 0005422
Part Number   : 7490032200 Rev 1
Vendor Id     : 1
Date Code     : 05222002
Country Code  : 1

Force10>
```

Both must list the same type

Figure 185 show linecard Command Example

In Figure 186, slot 1 does not contain a line card, but there is a card type configured for that slot (EW1YC). If you insert any other line card in that slot, the status of the line card in slot 1 changes to “type mismatch.”

```
Force10>show linecard all

--  Line cards  --
Slot  Status       NxtBoot    ReqTyp   CurTyp   Version     Ports
--------------------------------------------------------------------------
  0   online       online     E24SC    E24SC    3.1.2b2.30  24
  1   not present             EW1YC
  2   online       online     S48SC2   S48SC2   3.1.2b2.30  2
  3   not present
  4   online       online     E24SB    E24SB    3.1.2b2.30  24
  5   online       online     EX1YB    EX1YB    3.1.2b2.30  1
  6   not present
  7   not present
  8   online       online     F12PC    F12PC    3.1.2b2.30  12
  9   not present
 10   online       online     E24SC    E24SC    3.1.2b2.30  24
 11   not present
 12   power off    power off  S12YC12  S12YC12  3.1.2b2.30  12
 13   not present

Force10>
```

This slot was configured for a specific line card (EW1YC).

Figure 186 show linecard all Command Example with ReqType Listed for an Empty Slot

SFMs

In a system with nine SFMs, you can designate which SFM is the Standby SFM, and then hot-swap that SFM with another SFM. At boot time, the system designates, by default, the SFM in slot 8 as the Standby SFM.

Note: Standby SFM is available only in EtherScale systems.

To change an active SFM to the Standby SFM, use the following command:

Command Syntax	Command Mode	Usage
redundancy force-failover sfm *number*	EXEC privilege	Change the status of an SFM from Active to Standby. Enter the following parameter: *number:* 0 to 8 **Note:** On C-Series, this command could affect traffic even for the hot-failover case, since the switch fabric present on the RPM is taken down during the failover.

Use this command only when nine SFMs are present in the system.

Use the **show sfm all** command in the EXEC privilege mode to view which SFM is currently redundant.

```
Force10#show sfm all

Switch Fabric State:  up

--  Switch Fabric Modules  --
Slot  Status
---------------------------------------------------------------------------
  0   active
  1   active
  2   active
  3   active
  4   active
  5   active
  6   active
  7   active
  8   standby

Force10#
```

Figure 187 show sfm all Command Example

```
Force10#redundancy force-failover sfm 0
%TSM-6-SFM_FAILOVER: Standby switch to SFM 8
Standby switch to SFM 0
Force10#
```

Figure 188 redundancy force-failover sfm Command Example

Standby RPM

The E-Series supports the Online Insertion and Removal (OIR) of primary and Standby RPMs. For detailed information, refer to the Implementation Information for RPM Redundancy.

RPM Redundancy

The RPM in the E-Series is the core for routing and control operations. Routing table entries are built on the RPM and directed to the Forwarding (FIB) tables on the line cards. You must install at least one RPM for the E-Series to process packets.

Each RPM contains three CPUs. System control, Layer 2 and Layer 3 functions are divided among the three CPUs.

With two RPMs installed and online, FTOS supports 1+1 RPM redundancy, providing an extra module for failover. The primary RPM (in slot R0 by default) performs all routing and control operations (hardware mastership), and the Standby RPM is on-line and monitoring the primary RPM. Throughout this section, RPM0 is the primary RPM (slot R0) and RPM1 is the Standby RPM (slot R1).

Note: If your system contains two RPMs, both RPMs must contain the same software image. Starting with Release 6.3.1, you can load software on one RPM, and then use the **redundancy force-failover rpm** command to copy the upgrade to the other RPM.

Two RPMs in your E-Series enable it to experience a shorter transition period after an RPM failure. The Standby RPM does not need to reboot and can take over hardware mastership if necessary to return your E-Series to operational status.

FTOS supports the following RPM Redundancy tasks:

- Implementation Information on page 335 (optional)
- Security Considerations on page 336 (optional)
- RPM Failover Example on page 336 (optional)
- RPM High Availability Configuration on page 337 (optional)

Implementation Information

You can boot the system with one RPM and later add a second RPM, which will automatically become the Standby RPM. Force10 Networks recommends that you insert the Standby RPM after the primary RPM is online and stable, and that you copy the running configuration to the startup config file (**copy running-config startup-config** command) after the Standby RPM is online. You can tell when the Standby RPM is online when messages appear indicating that the RPMs have established a connection and the standby prompt appears (see Figure 189).

```
%RPM-2-MSG:CP0 %POLLMGR-2-ALT_RPM_STATE: Alternate RPM is present
%IRC-6-IRC_COMMUP: Link to peer RPM is up
%RAM-6-RAM_TASK: RPM1 is in Standby State.
Force10(standby)>
```

Figure 189 System Messages Indicating the Standby RPM is Online

When you boot the system with two RPMs installed, you can configure one to be the primary RPM or allow the system to select the primary. If you have not entered the **redundancy primary** command, the RPM in slot R0 defaults to the primary RPM. You can trigger an RPM failover through the command line, or it can occur because of one of the following reasons:

- The heartbeat (similar to a keepalive message) between the two RPMs is lost.
- The primary RPM experiences a problem (for example, a task crashed).
- The primary RPM is removed.

Security Considerations

After a failover, the new primary RPM (RPM1) prompts you for a username and password if local authentication was configured and that data was synchronized.

The Standby RPM does not use authentication methods involving client/server protocols, such as RADIUS or TACACS+.

RPM Failover Example

Below are the steps, actions, and results of a typical data synchronization between RPMs in an RPM failover.

Step	Action	Result
1	system boots with 2 RPMs	The system brings up the primary RPM first. If the **redundancy primary** command is not configured, the software automatically makes the RPM in slot 0 the primary RPM.
2	stable system, traffic running	The software performs incremental data synchronizations between the primary RPM and the Standby RPM when data changes on the primary RPM.
3	failover (either user-requested or triggered by an event in the system)	The Standby RPM (RPM1): • notifies all tasks about the RPM failover • transitions the tasks to the active state • reboots RPM0 If user-requested, the software prompts you to save the running configuration to the startup configuration. The process takes approximately 25 seconds.
4	RPM1 is up and traffic is flowing	RPM0 is now the Standby RPM and monitors RPM1.

RPM High Availability Configuration

FTOS provides the following processes to manage RPM failover:

- assign an RPM as primary on page 337 (optional)
- synchronize data between two RPMs on page 338 (optional)
- force an RPM failover on page 339 (optional)
- copy files between RPMs on page 341 (optional)
- specify the auto-failover-limit on page 341 (optional)
- disable auto-reboot on page 341 (optional)

For a complete listing of all commands related to RPM redundancy, refer to High Availability Chapter in the *FTOS Command Line Interface Reference*.

assign an RPM as primary

By default, FTOS assigns the RPM in slot R0 as the primary RPM.

To change this configuration, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Usage
redundancy primary [**rpm0** | **rpm1**]	CONFIGURATION	Assign an RPM as the primary RPM. • **rpm0:** the RPM in slot R0 • **rpm1**: the RPM in slot R1

Use the **show running-config redundancy** command to view the configuration of the redundancy features. Use the **show redundancy** command to view the status of the RPM and its role as a primary or secondary RPM.

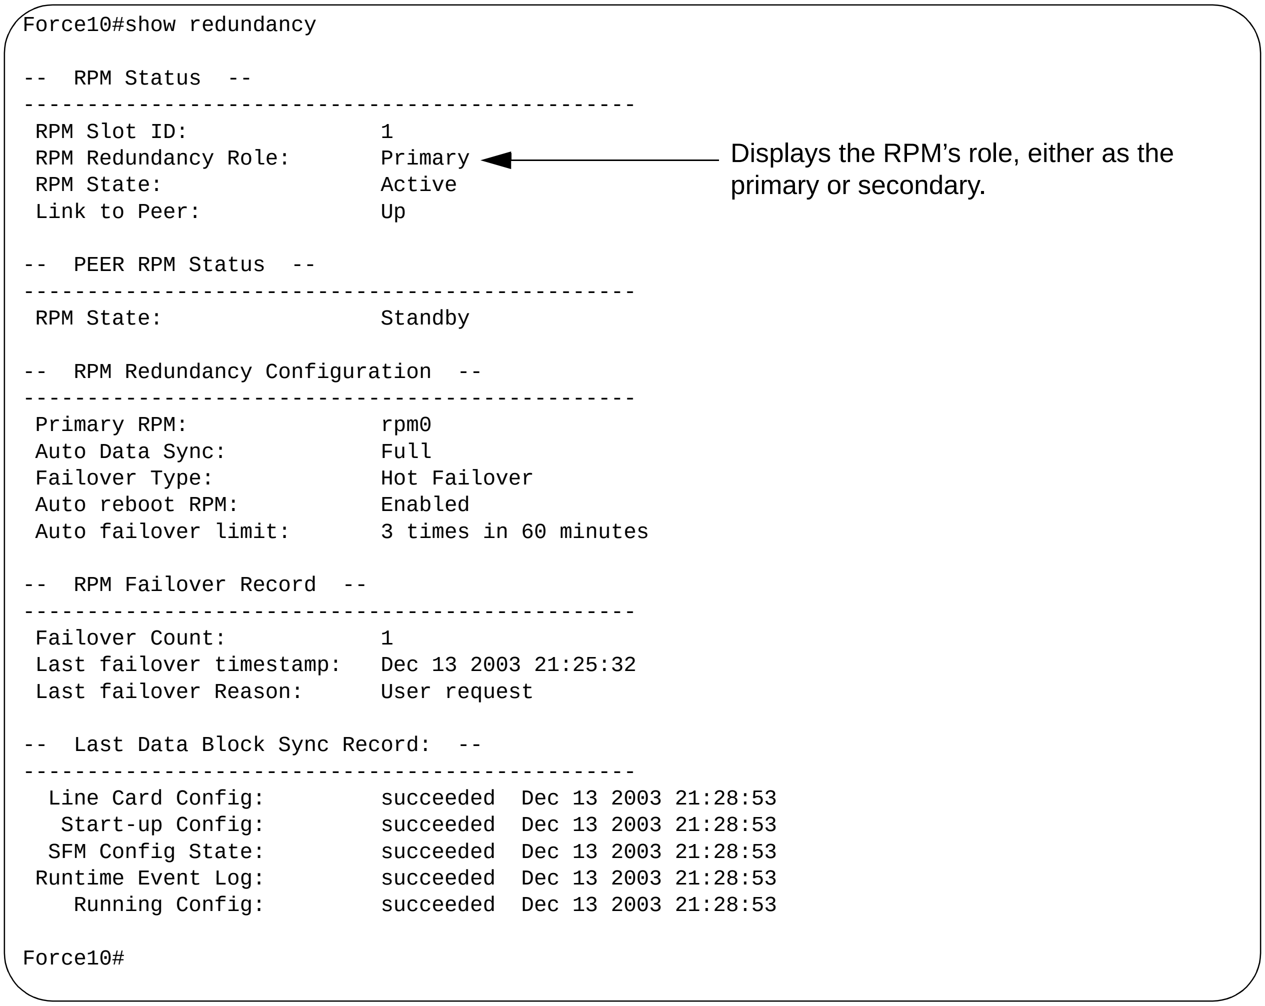

```
Force10#show redundancy

--  RPM Status  --
------------------------------------------------
 RPM Slot ID:               1
 RPM Redundancy Role:       Primary
 RPM State:                 Active
 Link to Peer:              Up

--  PEER RPM Status  --
------------------------------------------------
 RPM State:                 Standby

--  RPM Redundancy Configuration  --
------------------------------------------------
 Primary RPM:               rpm0
 Auto Data Sync:            Full
 Failover Type:             Hot Failover
 Auto reboot RPM:           Enabled
 Auto failover limit:       3 times in 60 minutes

--  RPM Failover Record  --
------------------------------------------------
 Failover Count:            1
 Last failover timestamp:   Dec 13 2003 21:25:32
 Last failover Reason:      User request

--  Last Data Block Sync Record:  --
------------------------------------------------
  Line Card Config:         succeeded  Dec 13 2003 21:28:53
   Start-up Config:         succeeded  Dec 13 2003 21:28:53
  SFM Config State:         succeeded  Dec 13 2003 21:28:53
 Runtime Event Log:         succeeded  Dec 13 2003 21:28:53
    Running Config:         succeeded  Dec 13 2003 21:28:53

Force10#
```

Displays the RPM’s role, either as the primary or secondary.

Figure 190 show redundancy Command Example

synchronize data between two RPMs

By default, all data between the two RPMs is synchronized directly after boot-up. You have several options to synchronize data between the RPMs.

- clock
- preferred primary RPM configuration (**redundancy primary** command)
- boot information
- management port IP address information

Once the two RPMs have done an initial full synchronization, thereafter FTOS only updates changed data.

Table 25 lists the data categories that can be synchronized between the two RPMs.

Table 25 Data Categories

Parameter	Description
full	All operational data. This setting is the default.
persistent-data	The startup-configuration file.
system-data	Includes the data in the persistent-data parameter and the following: • running-configuration file • event log • SFM status • line card status

To change the type of data synchronized and the method of synchronization, use the command:

Command Syntax	Command Mode	Usage
redundancy synchronize [**full** | **persistent-data** | **system-data**]	EXEC privilege	Synchronize data one time between the primary and Standby RPMs.

Use the **show running-config redundancy** command to view the current redundancy configuration.

force an RPM failover

You can trigger a failover between RPMs. This feature is useful in two cases:

- For replacing an RPM: You can fail over to the Standby RPM, and then replace the old primary RPM. After failover, RPM1 (the new primary RPM) reboots RPM0 (the former primary RPM) and is active.
- For upgrading software: You can load the new software on the RPM, and then force a failover instead of doing a full reboot. Some patch releases require only a “hitless upgrade”, which means that a software upgrade does not require a reboot of the line cards. Major releases require a warm upgrade, which is a reset of the line cards and SFMs. For details, see the lastest FTOS release notes.

To force an RPM failover, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Usage
redundancy force-failover rpm	EXEC privilege	Manually trigger the Standby RPM to take mastership of the system. The system recognizes if you have installed a software upgrade on the RPM, and prompts you to boot the RPM, as shown in the following example. **Note:** On C-Series, this command could affect traffic even for the hot-failover case, since the switch fabric present on the RPM is taken down during the failover.

Executing the **redundancy force-failover rpm** command after installing software on the primary RPM causes the system to display the following dialog:

```
Force10#redundancy force-failover rpm
Peer RPM's SW version is different but HA compatible.
Failover can be done by warm or hitless upgrade.
All linecards will be reset during warm upgrade.

Specify hitless upgrade or warm upgrade [confirm hitless/warm]:hitless
Proceed with warm upgrade [confirm yes/no]:
```

Figure 191 Using the redundancy force-failover rpm Command to Copy Software between RPMs

The **show redundancy** command from the primary RPM displays all information on both the primary and secondary:

```
Force10#show redundancy

--  RPM Status  --
------------------------------------------------
 RPM Slot ID:                0
 RPM Redundancy Role:        Primary
 RPM State:                  Active
 Link to Peer:               Up

--  PEER RPM Status  --
------------------------------------------------
 RPM State:                  Standby

--  RPM Redundancy Configuration  --
------------------------------------------------
 Primary RPM:                rpm0
 Auto Data Sync:             Full
 Failover Type:              Hot Failover
 Auto reboot RPM:            Enabled
 Auto failover limit:        3 times in 60 minutes

--  RPM Failover Record  --
------------------------------------------------
 Failover Count:             2
 Last failover timestamp:    Dec 13 2003 21:41:35

Force10#
```

Figure 192 show redundancy Command Example from the Primary RPM

copy files between RPMs

To copy files between RPMs, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Usage
copy *file-url file-url*	EXEC privilege	Use the following *file-url* parameters to copy files between RPMs: • **rpm0flash:**//*filename* (copy a file to/from the internal flash on the RPM in slot R0.) • **rpm0slot0:**//*filename* (copy a file to/from the external flash in the RPM in slot R0.) • **rpm1flash:**//*filename* (copy a file to/from the internal flash in the RPM in slot R1.) • **rpm1slot0:**//*filename* (copy a file to/from the external flash in the RPM in slot R1.)

specify the auto-failover-limit

You can specify an auto-failover limit for RPMs. When a non-recoverable fatal error is detected, an automatic failover occurs and a RPM failover is initiated. By default the auto-failover-limit is enabled. This utility does not impact your ability to initiate manual failovers.

Command Syntax	Command Mode	Usage
redundancy auto-failover-limit [**count** *number* **period** *minutes*]	CONFIGURATION	Use the following parameters to configure an auto-failover limit for your RPMs: • **count***:* The maximum number of times the RPMs can automatically failover within the period you specify. The default is 3. • **period***:* The duration in which to allow a maximum number of automatic failovers. The default is 60 minutes.

To disable the auto-failover-limit, use the **no redundancy** command in CONFIGURATION mode. To re-enable the auto-failover-limit with its default parameters, in CONFIGURATION mode, use the **redundancy auto-failover-limit** command without parameters.

disable auto-reboot

If you wish to keep an RPM in its failed state, you can prevent FTOS from automatically rebooting it and making it the Standby RPM. To do so, use the following command:

Command Syntax	Command Mode	Usage
redundancy disable-auto-reboot rpm	CONFIGURATION	Prevents your E-Series from rebooting a failed RPM. By default, auto-reboot is disabled.

High Availabilty on C-Series

C-Series	✓
E-Series	NO

Platform Specific Feature: The high availability behavior described in this section applies to C-Series only.

Note: All other HA commands available for E-Series are available for C-Series.

C-series line cards and RPMs can be hot-swapped, which means that a line card or RPM can be added, replaced, or removed without interrupting the system. On the C-Series the RPM also contains the switch fabric.

This section covers the following topics:

- Inserting a Line Card into an Operational System on page 342
- Replacing an Existing Line Card on page 345
- Inserting a Second RPM on page 345
- Version Compatibility between RPMs on page 346
- RPM Failover on page 347

Inserting a Line Card into an Operational System

When you insert a line card into an operational system, the software detects the line card type and writes the line card information into the running-configuration. This information is retained even after the line card is removed from the system.

Figure 193 Inserting a Line Card into an Operational C-Series System

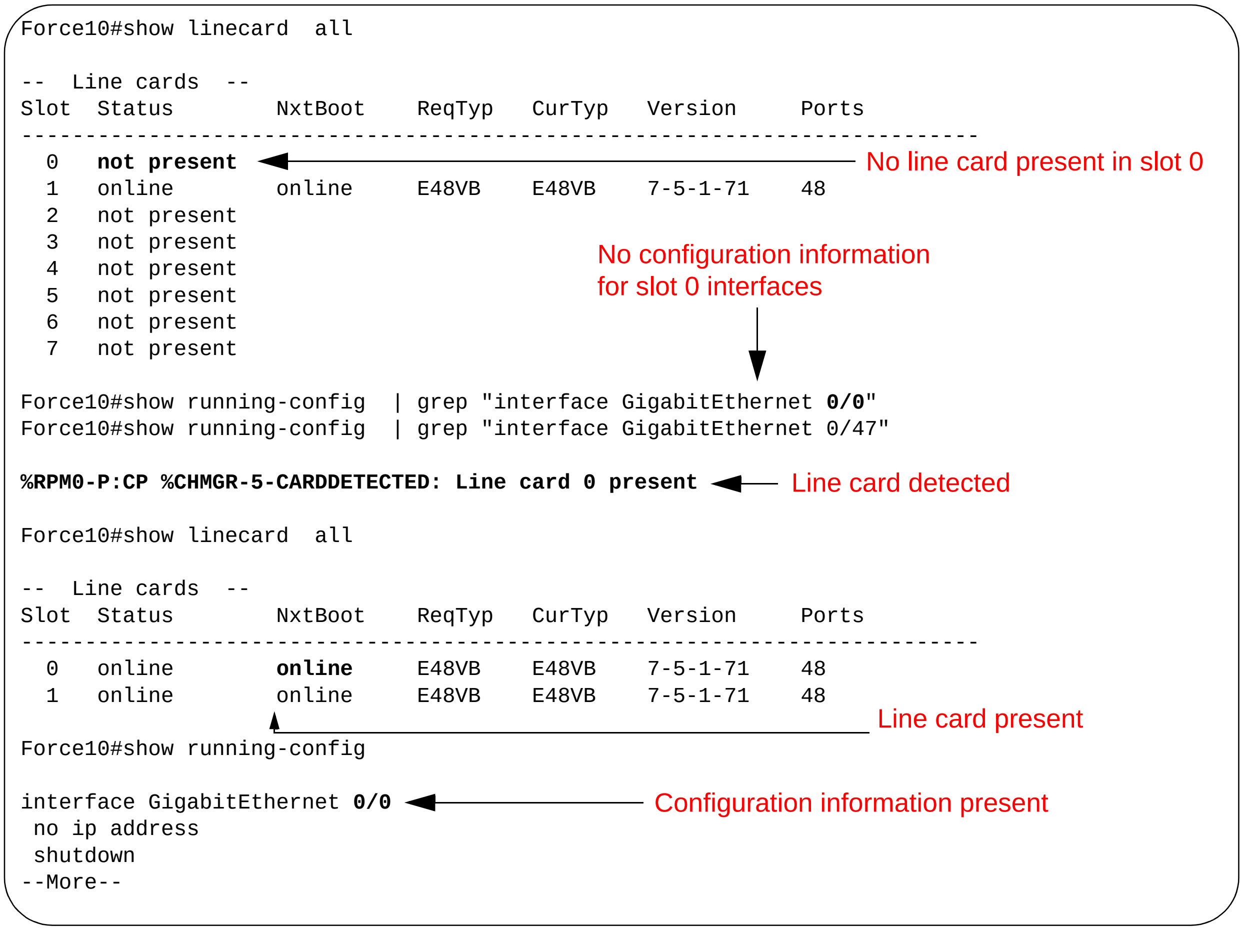

```
Force10#show linecard  all

--  Line cards  --
Slot  Status         NxtBoot    ReqTyp   CurTyp   Version     Ports
---------------------------------------------------------------------------
  0   not present
  1   online         online     E48VB    E48VB    7-5-1-71    48
  2   not present
  3   not present
  4   not present
  5   not present
  6   not present
  7   not present

Force10#show running-config  | grep "interface GigabitEthernet 0/0"
Force10#show running-config  | grep "interface GigabitEthernet 0/47"

%RPM0-P:CP %CHMGR-5-CARDDETECTED: Line card 0 present

Force10#show linecard  all

--  Line cards  --
Slot  Status         NxtBoot    ReqTyp   CurTyp   Version     Ports
---------------------------------------------------------------------------
  0   online         online     E48VB    E48VB    7-5-1-71    48
  1   online         online     E48VB    E48VB    7-5-1-71    48

Force10#show running-config

interface GigabitEthernet 0/0
 no ip address
 shutdown
--More--
```

No line card present in slot 0

No configuration information for slot 0 interfaces

Line card detected

Line card present

Configuration information present

Figure 194 Removing a Line Card from an Operational C-Series System

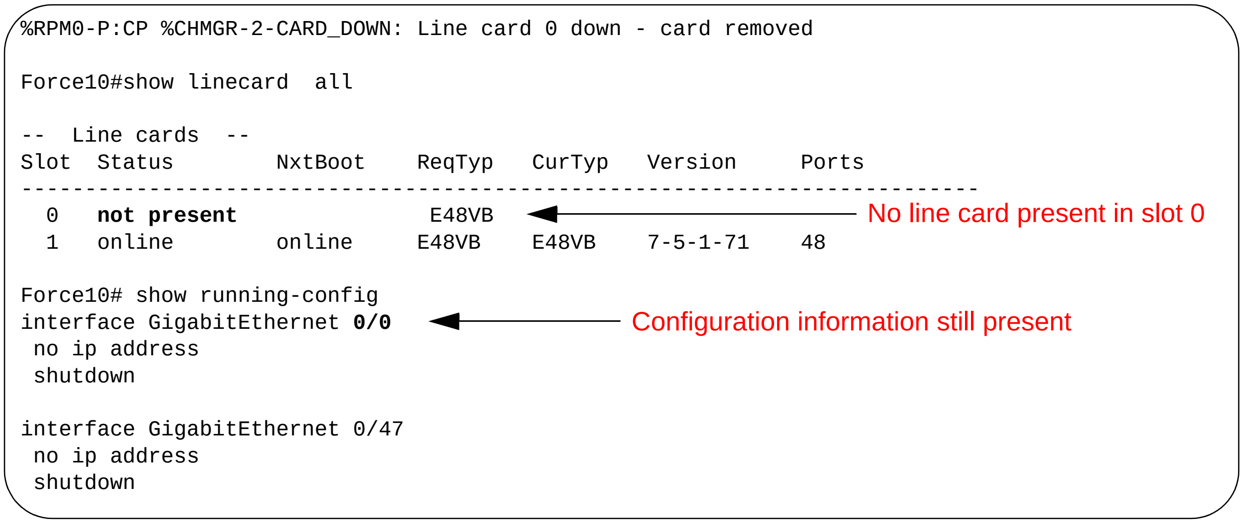

```
%RPM0-P:CP %CHMGR-2-CARD_DOWN: Line card 0 down - card removed

Force10#show linecard  all

--  Line cards  --
Slot  Status         NxtBoot    ReqTyp   CurTyp   Version     Ports
---------------------------------------------------------------------------
  0   not present                E48VB
  1   online         online     E48VB    E48VB    7-5-1-71    48

Force10# show running-config
interface GigabitEthernet 0/0
 no ip address
 shutdown

interface GigabitEthernet 0/47
 no ip address
 shutdown
```

No line card present in slot 0

Configuration information still present

Line card slots can be pre-configured before a line card is inserted. However, if you pre-configure a line card and insert a different type of card into the slot, the system provides an error message (Message 5) and the installed line card has a card mismatch status. To clear this line card mismatch status and bring the correct line card online, re-issue the **linecard** command.

Message 5 Line card Mismatch Error

```
%RPM0-P:CP %CHMGR-3-CARD_MISMATCH: Mismatch: line card 0 is type E48VB - type E48TB required
```

Figure 195 Pre-configuring a Line Card Slot

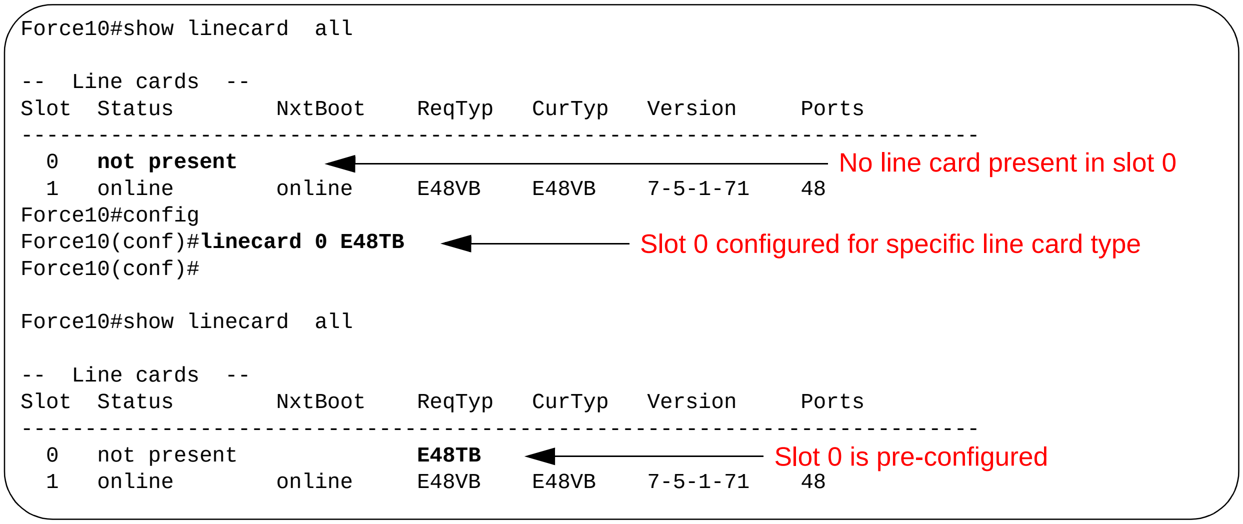

```
Force10#show linecard  all

--  Line cards  --
Slot  Status        NxtBoot    ReqTyp   CurTyp   Version     Ports
---------------------------------------------------------------------------
  0   not present
  1   online        online     E48VB    E48VB    7-5-1-71    48
Force10#config
Force10(conf)#linecard 0 E48TB
Force10(conf)#

Force10#show linecard  all

--  Line cards  --
Slot  Status        NxtBoot    ReqTyp   CurTyp   Version     Ports
---------------------------------------------------------------------------
  0   not present              E48TB
  1   online        online     E48VB    E48VB    7-5-1-71    48
```

No line card present in slot 0

Slot 0 configured for specific line card type

Slot 0 is pre-configured

Figure 196 Troubleshooting a Line Card Mismatch Condition

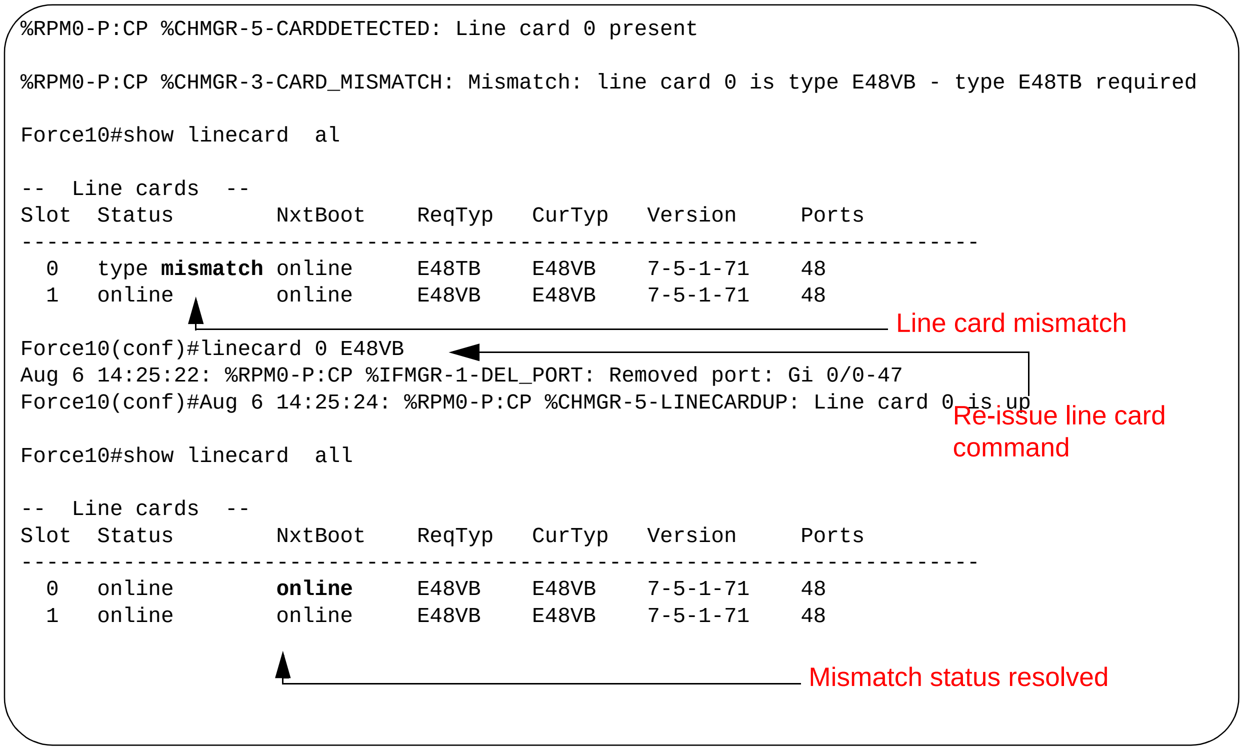

```
%RPM0-P:CP %CHMGR-5-CARDDETECTED: Line card 0 present

%RPM0-P:CP %CHMGR-3-CARD_MISMATCH: Mismatch: line card 0 is type E48VB - type E48TB required

Force10#show linecard  al

--  Line cards  --
Slot  Status        NxtBoot    ReqTyp   CurTyp   Version     Ports
---------------------------------------------------------------------------
  0   type mismatch online     E48TB    E48VB    7-5-1-71    48
  1   online        online     E48VB    E48VB    7-5-1-71    48

Force10(conf)#linecard 0 E48VB
Aug 6 14:25:22: %RPM0-P:CP %IFMGR-1-DEL_PORT: Removed port: Gi 0/0-47
Force10(conf)#Aug 6 14:25:24: %RPM0-P:CP %CHMGR-5-LINECARDUP: Line card 0 is up

Force10#show linecard  all

--  Line cards  --
Slot  Status        NxtBoot    ReqTyp   CurTyp   Version     Ports
---------------------------------------------------------------------------
  0   online        online     E48VB    E48VB    7-5-1-71    48
  1   online        online     E48VB    E48VB    7-5-1-71    48
```

Line card mismatch

Re-issue line card command

Mismatch status resolved

Replacing an Existing Line Card

When a line card is replaced with another, the system checks to ensure that the line card type is identical. If it is not, a mismatch error message (Message 5) is displayed. To clear this error and bring the card online, re-issue the **linecard** command, as shown in Figure 196.

Figure 197 Replacing an Exising Line Card

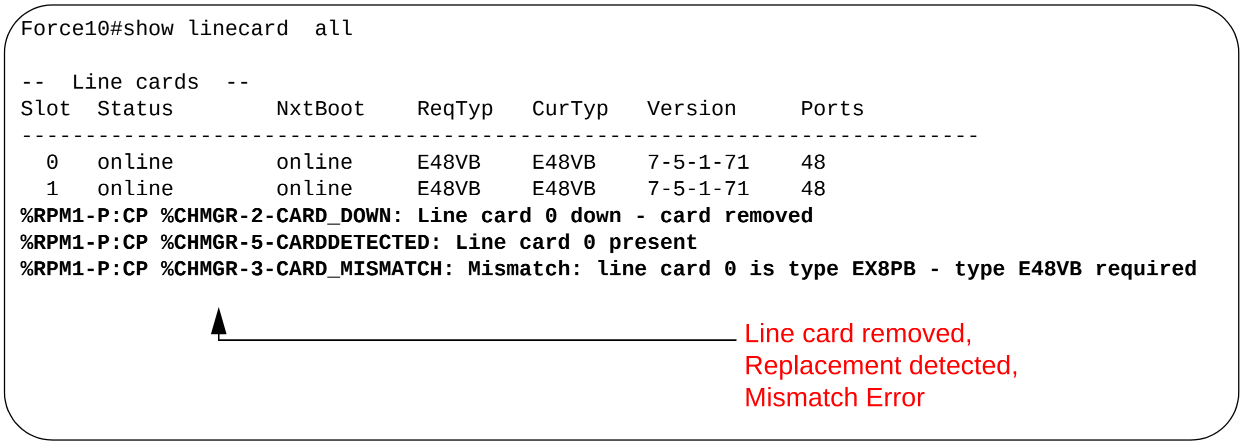

```
Force10#show linecard  all

--  Line cards  --
Slot  Status        NxtBoot    ReqTyp   CurTyp   Version     Ports
---------------------------------------------------------------------------
  0   online        online     E48VB    E48VB    7-5-1-71    48
  1   online        online     E48VB    E48VB    7-5-1-71    48
%RPM1-P:CP %CHMGR-2-CARD_DOWN: Line card 0 down - card removed
%RPM1-P:CP %CHMGR-5-CARDDETECTED: Line card 0 present
%RPM1-P:CP %CHMGR-3-CARD_MISMATCH: Mismatch: line card 0 is type EX8PB - type E48VB required
```

Line card removed,
Replacement detected,
Mismatch Error

Inserting a Second RPM

The C-Series system is functional with only one RPM. If a second RPM is inserted, it comes online as the standby RPM.

Figure 198 Inserting a Second RPM into an Operational System

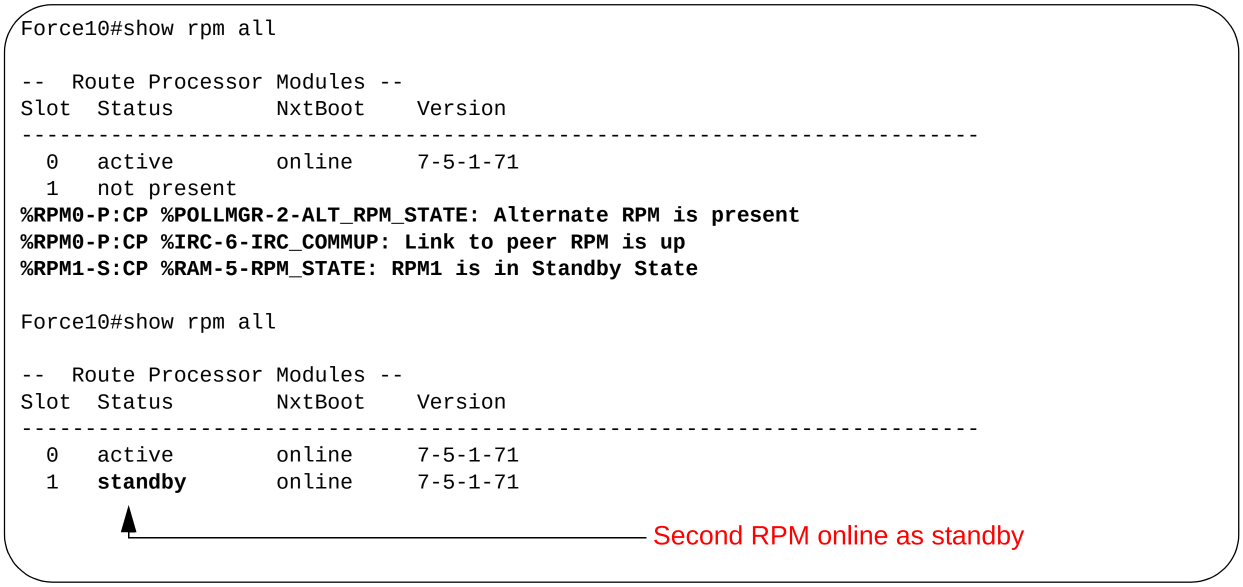

```
Force10#show rpm all

--  Route Processor Modules --
Slot  Status        NxtBoot    Version
---------------------------------------------------------------------------
  0   active        online     7-5-1-71
  1   not present
%RPM0-P:CP %POLLMGR-2-ALT_RPM_STATE: Alternate RPM is present
%RPM0-P:CP %IRC-6-IRC_COMMUP: Link to peer RPM is up
%RPM1-S:CP %RAM-5-RPM_STATE: RPM1 is in Standby State

Force10#show rpm all

--  Route Processor Modules --
Slot  Status        NxtBoot    Version
---------------------------------------------------------------------------
  0   active        online     7-5-1-71
  1   standby       online     7-5-1-71
```

Second RPM online as standby

Only the primary RPM authenticates the user, the standby RPM does not support client/server authentication protocols such as RADIUS or TACACS+.

Version Compatibility between RPMs

The system checks to ensure the high availability compatibility of the software images on the two RPMs. Images are compatible only if the first two digits of the version number match. If they do not, an error message is displayed. View the loaded versions of FTOS on the RPMs using the command **show redundancy** (Figure 199).

Message 6 FTOS Version Incompatibility Error

```
***********************************************
 *
 *     Warning !!!  Warning !!!  Warning !!!
 *
 * ---------------------------------------------
 *
 *       Incompatible SW Version detected !!
 *
 *       This RPM -> E7-5-1-71
 *       Peer RPM -> E7-5-1-56
 *
 ***********************************************
Aug 6 18:39:33: %RPM0-P:CP %IRC-4-IRC_VERSION: Current RPM E7-5-1-71 Peer RPM E7-5-1-56 -
Different software version detected
```

Figure 199 Version Incompatibility in show redundancy Command

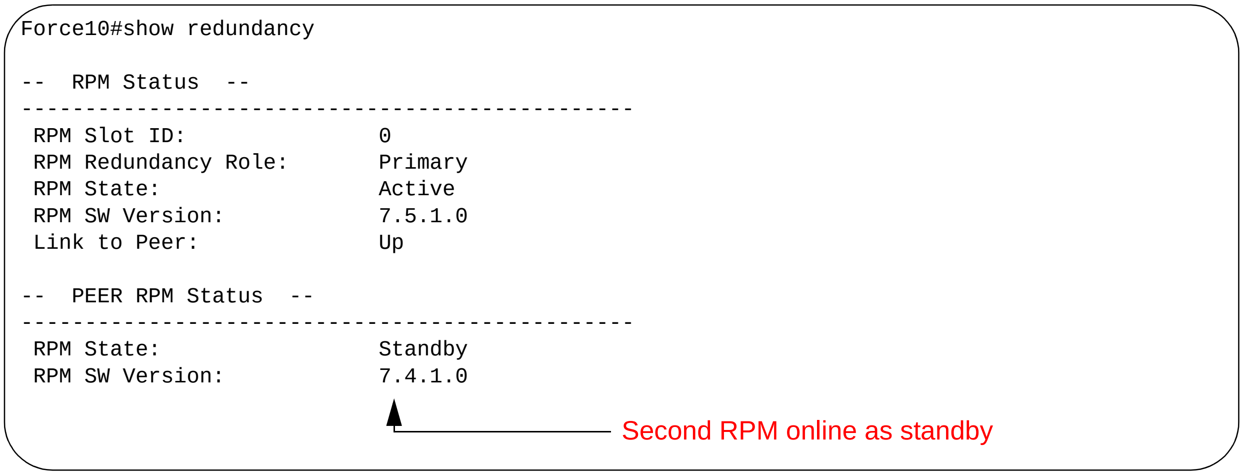

```
Force10#show redundancy

--  RPM Status  --
------------------------------------------------
 RPM Slot ID:                0
 RPM Redundancy Role:        Primary
 RPM State:                  Active
 RPM SW Version:             7.5.1.0
 Link to Peer:               Up

--  PEER RPM Status  --
------------------------------------------------
 RPM State:                  Standby
 RPM SW Version:             7.4.1.0
```

Second RPM online as standby

Since the RPM on the C-Series also contains the switch fabric, even though the second RPM comes online as the standby, the switch fabric is active and is automatically available for routing. Change this behavior using the command **offline sfm standby**. To bring the secondary SFM online, issue the **online sfm standby** command. There is traffic loss anytime an SFM is brought online or taken offline. Use the command **show sfm all** to determine the status of the SFMs on the RPMs.

RPM Failover

When a system with two RPMs boots or reboots, the system brings up the primary RPM first. By default the RPM in slot 0 becomes the primary RPM. To change this, use the command **redundancy primary**. FTOS copies all data from RPM0 to RPM1. Once the system is stable only changes on the primary RPM are copied to the secondary RPM. You can perform a one-time synchronization of data by issuing the command **redundancy synchronize**.

Given the default conditons, when a failover occurs, RPM 1 becomes the primary. RPM 0 reboots and becomes the standby with its switch fabric available for routing.

By default, when a secondary RPM with a logical SFM is inserted or removed, the system must add or remove the backplane links to the switch fabric trunk. Any time such links are changed, traffic is disrupted. Use the command **redundancy sfm standby** to avoid any traffic disruption when the secondary RPM is inserted. When this command is executed, the logical SFM on the standby RPM is immediately taken offline, and the SFM state set as standby. Use the command **show sfm all** to see SFM status information.

Chapter 19 Quality of Service

Force10 Networks' implementation of Quality of Service (QoS) enables customers to configure network traffic conditioning and congestion control into easy-to-use groupings.

Use FTOS to establish QoS configurations that are:

- per-physical-port-based
- policy-based

➡ **Note:** QoS is supported on series ED, EE, EF, and EG series line cards. The Force10 Networks 4-port OC-48c/OC-12c/OC-3c POS (LC-EG-OC48-4P) line card supports only port-based QoS.

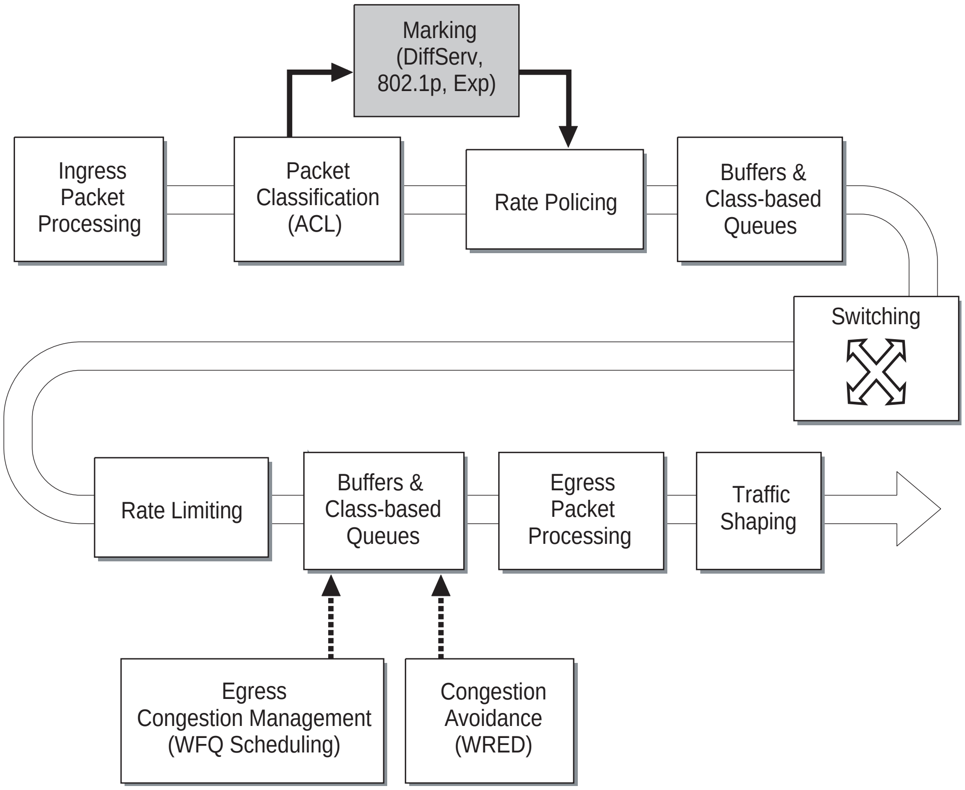

Figure 200 Force10 Networks QoS Architecture

Force10 Networks' QoS implementation complies with IEEE 802.1p *User Priority Bits for QoS Indication*. It also implements these Internet Engineering Task Force (IETF) documents:

- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
- RFC 2475, *An Architecture for Differentiated Services*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*

Control Traffic Prioritization

Control traffic prioritization ensures control traffic priority against data traffic or other low priority traffic during traffic congestion. This prioritization is achieved through the classification, marking, rate limiting, buffering, and scheduling.

Class-Based Queuing and Rate Limiting

To allow packet treatment differentiation, traffic is classified according to the priority and put into different queues. Rate limiting is applied to each class/queue. In the following tables, queue assignment, rate limiting parameters, scheduling and buffering parameters are listed.

Aggregated Shaping

Aggregated shaping helps reduce the rate/burst of traffic after the CPU recovers from high usage and smooths CPU usage under high congestion.

Local Multicast Control Traffic Protection

C-Series	NO
E-Series	✓

Platform Specific Feature: Local Multicast Control Traffic Protection is supported on E-Series only.

When multicast traffic is sent to one multicast egress queue while there is low priority multicast flooding, control traffic may get dropped. Therefore, the multicast traffic and local CPUs generated packets are marked. On egress, buffer space is reserved for multicast traffic using WRED (Weighted Random Early Drop) and tail drop.

When an egress port is congested, the multicast queue builds up. After its threshold is met, only control traffic can be buffered. When buffer occupancy is lower than multicast traffic threshold (when scheduling catches up), multicast traffic is accepted again. By doing this, CPU-generated multicast control traffic is given a strict high priority against pass-through multicast traffic.

Per-Port QoS Configurations

Per-Port QoS allows users to define QoS configuration on a per-physical-port basis. These include:

- dot1p-priority
- rate police
- rate limit
- storm control
- rate shape

Policy QoS configurations can co-exist with port-based configurations if you configure them on different interfaces.

Configuration Task List

Users can configure the following QoS features on an interface:

- set dot1p priorities for incoming traffic on page 352
- apply dot1p priorities to incoming traffic on page 353
- Rate Police for Incoming Traffic on page 353
- Rate Limit for Outgoing Traffic on page 354

- limit or suppress traffic on page 356
- Rate Shape of Outgoing Traffic on page 355

set dot1p priorities for incoming traffic

To assign a value for the IEEE 802.1p bits on traffic received on an interface, use this command:

Command Syntax	Command Mode	Usage
dot1p-priority priority-value	INTERFACE	The dot1p-priority command changes the priority of incoming traffic on the interface. FTOS places traffic marked with a priority in the correct queue and processes that traffic according to its queue. When you set the priority for a port channel, the physical interfaces assigned to the port channel are configured with the same value. You cannot assign dot1p-priority command to individual interfaces in a port channel. Refer to Table 26 for more information about entering dot1p-priority values.

Table 26 dot1p-priority values and queue numbers

dot1p	E-Series Queue Number	C-Series Queue Number
0	2	1
1	0	—
2	1	0
3	3	—
4	4	2
5	5	—
6	6	3
7	7	—

Figure 201 below illustrates how to configure a dot1p-priority:

```
Force10#config t
Force10(conf)#interface gigabitethernet 1/0
Force10(conf-if)#switchport
Force10(conf-if)#dot1p-priority 1
Force10(conf-if)#end
Force10#
```

Figure 201 dot1p-priority Command Example

apply dot1p priorities to incoming traffic

Use this command to honor all incoming 802.1p markings on incoming switched traffic on the interface:

Note: The **service-policy input** *policy-map-name* command and the **service-class dynamic dot1p** command are not allowed simultaneously on a physical interface.

The **service-policy** commands are not allowed on a port channel.

Command Syntax	Command Mode	Usage
service-class dynamic dot1p	INTERFACE	Enter this command to honor all incoming 802.1p markings on incoming switched traffic on the interface. By default, this facility is not enabled (that is, the 802.1p markings on incoming traffic are not honored). This command can now be applied on both physical interfaces and port channels. When you set the service-class dynamic for a port channel, the physical interfaces assigned to the port channel are configured with this configuration. You cannot assign the service-class dynamic command to individual interfaces in a port channel.

To return to the default setting, use the **no service-class dynamic dot1p** command.

Figure 202 shows how to apply dot1p priorities to incoming switched traffic:

```
Force10#config t
Force10(conf)#interface gigabitethernet 1/0
Force10(conf-if)#service-class dynamic dot1p
Force10(conf-if)#end
Force10#
```

Figure 202 service-class dynamic dot1p Command Example

Rate Police for Incoming Traffic

You can configure rate policing for an interface. If you use VLANs for each physical interface, you can configure rate police commands specifying different VLANs.

Note: Do not confuse the INTERFACE QoS “rate police” command with the POLICY QoS “rate-police” command.

Command Syntax	Command Mode	Usage
rate police committed-rate [**burst-KB**] [**peak peak-rate** [**burst-KB**]] [**vlan vlan-id**]	INTERFACE	• **committed-rate**: Enter a number as the bandwidth in Mbps. Range: 0 to 10000 • **burst-KB**: (OPTIONAL) Enter a number as the burst size in KB. Range: 16 to 200000. Default: 50 • **peak peak-rate**: (OPTIONAL) Enter the keyword **peak** followed by a number to specify the peak rate in Mbps. Range: 0 to 10000 • **vlan vlan-id**. (OPTIONAL) Enter the keyword **vlan** followed by a VLAN ID to police traffic to those specific VLANs. Range: 1 to 4094

To remove rate policing, enter the **no rate police committed-rate** [**burst-KB**] [**peak peak-rate** [**burst-KB**]] [**vlan vlan-id**] command.

Figure 203 below demonstrates how to set the rate police for an interface:

```
Force10#config t
Force10(conf)#interface gigabitethernet 1/0
Force10(conf-if)#rate police 100 40 peak 150 50
Force10(conf-if)#end
Force10#
```

Figure 203 rate police Command Example

Rate Limit for Outgoing Traffic

C-Series	NO
E-Series	✓

Platform Specific Feature: Rate Limit for Outgoing Traffic is supported on the E-Series only.

For each interface, you can also rate limit the outgoing traffic. If you use VLANs, for each physical interface, you can configure six rate limit commands specifying different VLANs.:

Note: Do not confuse the INTERFACE QoS “rate limit” command with the POLICY QoS “rate-limit” command.

Command Syntax	Command Mode	Usage
rate limit committed-rate [**burst-KB**] [**peak peak-rate** [**burst-KB**]] [**vlan vlan-id**]	INTERFACE	• **committed-rate**: Enter the bandwidth in Mbps. Range: 0 to 10000 • **burst-KB**: (OPTIONAL) Enter the burst size in KB. Range: 16 to 200000. Default: 50 • **peak peak-rate**: ((OPTIONAL) Enter the keyword **peak** followed by a number to specify the peak rate in Mbps. Range: 0 to 10000 • **vlan vlan-id**: (OPTIONAL) Enter the keyword vlan followed by a VLAN ID to limit traffic to those specific VLANs. Range: 1 to 4094

To remove rate limiting, use the **no rate limit committed-rate** [**burst-KB**] [**peak peak-rate** [**burst-KB**]] [**vlan vlan-id**] command.

Figure 204 shows how to rate limit outgoing traffic:

```
Force10#config t
Force10(conf)#interface gigabitethernet 1/0
Force10(conf-if)#rate limit 100 40 peak 150 50
Force10(conf-if)#end
Force10#
```

Figure 204 rate limit Command Example

Rate Shape of Outgoing Traffic

C-Series	NO
E-Series	✓

Platform Specific Feature: Rate Shape of Outgoing Traffic is supported on the E-Series only.

For each interface, you can also rate shape the outgoing traffic:

Command Syntax	Command Mode	Usage
rate shape rate [**burst-KB**]	INTERFACE	• *rate*: Enter the outgoing rate in multiples of 10 Mbps. Range: 0 to 10000 • *burst-KB*: (OPTIONAL) Enter a number as the burst size in KB. Range: 0 to 10000. The default is 4 KB.

To delete the command, use the **no rate shape rate** command.

Figure 205 shows how to rate shape outgoing traffic:

```
Force10#config t
Force10(conf)#interface gigabitethernet 1/0
Force10(conf-if)#rate shape 500 50
Force10(conf-if)#end
Force10#
```

Figure 205 rate shape Command Example

limit or suppress traffic

Storm-Control allows you to control Unknown-Unicast and Broadcast traffic on Layer 2 and Layer 2/ Layer 3 interfaces. Below are important considerations before using the storm-control feature:

- The interface commands can be applied only on physical interfaces (VLANs and LAG interfaces are not supported).
- An Interface-level command supports only storm-control configuration on the Ingress.
- An Interface-level command overrides any Configuration-level Ingress command for that physical interface, if both are configured.
- The configuration-level storm-control commands can be applied at the Ingress or Egress and are supported on all physical interfaces.
- On E-Series, when storm control is applied on an interface, the percentage of storm control applied is calculated based on the advertised rate of the line card. Not based on the speed setting for the line card.
- Do not apply per-VLAN QoS on an interface that has storm-control enabled (either on an Interface or globally)
- On E-Series, when broadcast storm-control is enabled on an interface or globally on the Ingress and DSCP marking for a DSCP value 1 is configured for the data traffic, the traffic goes to queue 1 instead of queue 0. Similarly, if unicast storm-control is enabled on an interface or globally on the Ingress, and DSCP marking for a DSCP value 2 is configured for the data traffic, the traffic goes to queue 2 instead of queue 0.

Note: On E-Series, Bi-directional traffic (unknown unicast and broadcast) along with egress storm control causes the configured traffic rates to be split between the involved ports. The percentage of traffic that each port receives after the split is not predictable. These ports can be in the same/different port-pipes, or the same/different line cards.

To storm control the unknown-unicast and broadcast traffic on a Layer 2 or Layer 2/Layer 3 interface, use the following command:

Command Syntax	Command Mode	Usage
storm-control {broadcast | unknown-unicast} [*percentage* | *pps*] **in]**	INTERFACE	Enter the amount of traffic allowed into an interface. Percentage: 1 to 100 (E-Series) Packets per second: 1 - 33554431(C-Series)

On E-Series, at the Interface level (the command above) FTOS permits storm-control in the inbound direction only. At the global configuration level (the command below), FTOS permits storm-control in both the inbound and outbound directions. The Interface level configuration overrides the global configuration.

Command Syntax	Command Mode	Usage
storm-control {broadcast | unknown-unicast} [*percentage* | *pps* **in** | **out**]	CONFIGURATION	Enter the amount of traffic allowed in or out of the network. Percentage: 1 to 100 (E-Series) Packets per second: 1 - 33554431(C-Series) Broadcast Storm-Control is valid on Layer 2/ Layer 3 interfaces only. Layer 2 broadcast traffic is treated as unknown-unicast traffic. **out** is supported on E-Series only.

Show Commands

C-Series	NO
E-Series	✓

Platform Specific Feature: The command **show interfaces rate** is supported on the E-Series only.

To view all configured interfaces, use the **show interfaces rate** command:

Command Syntax	Command Mode	Usage
show interfaces [interface] rate [limit | police]	EXEC EXEC privilege	(OPTIONAL) Enter the following keywords and slot/port or number information: • For a 10/100/1000 Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a SONET interface, enter the keyword **sonet** followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. Enter the keyword **limit** to view the outgoing traffic rate. Enter the keyword **police** to view the incoming traffic rate.

See Figure 206 and Figure 207 for some examples of output from this command.

```
Force10#show interfaces gigabitEthernet 1/1 rate limit
  Rate limit 300 (50) peak 800 (50)
    Traffic Monitor 0: normal 300 (50) peak 800 (50)
      Out of profile yellow 23386960 red 320605113
    Traffic Monitor 1: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 2: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 3: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 4: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 5: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 6: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 7: normal NA peak NA
      Out of profile yellow 0 red 0
    Total: yellow 23386960 red 320605113
```

Figure 206 show interfaces rate limit Command Example

```
Force10#show interfaces gigabitEthernet 1/2 rate police
  Rate police 300 (50) peak 800 (50)
    Traffic Monitor 0: normal 300 (50) peak 800 (50)
      Out of profile yellow 23386960 red 320605113
    Traffic Monitor 1: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 2: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 3: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 4: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 5: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 6: normal NA peak NA
      Out of profile yellow 0 red 0
    Traffic Monitor 7: normal NA peak NA
      Out of profile yellow 0 red 0
    Total: yellow 23386960 red 320605113
```

Figure 207 show interfaces rate police Command Example

For more information on port-based QoS commands, refer to the *FTOS Command Line Interface Reference*.

Policy-based QoS Configurations

Policy-based QoS is not supported on logical interfaces, such as port-channels, VLANS, or loopbacks.The goals of the Force10 Networks implementation of policy-based QoS are:

- Provide a flexible and powerful method of provisioning QoS on the E-Series
- Separate classification and QoS functionalities
- Isolate classification changes from QoS policy changes
- Enable customers to apply policy maps to multiple physical interfaces
- Allow customers to apply a single class definition to multiple policy maps

- Provide the maximum amount of compatibility between FTOS policy QoS and that of other major network equipment manufacturers

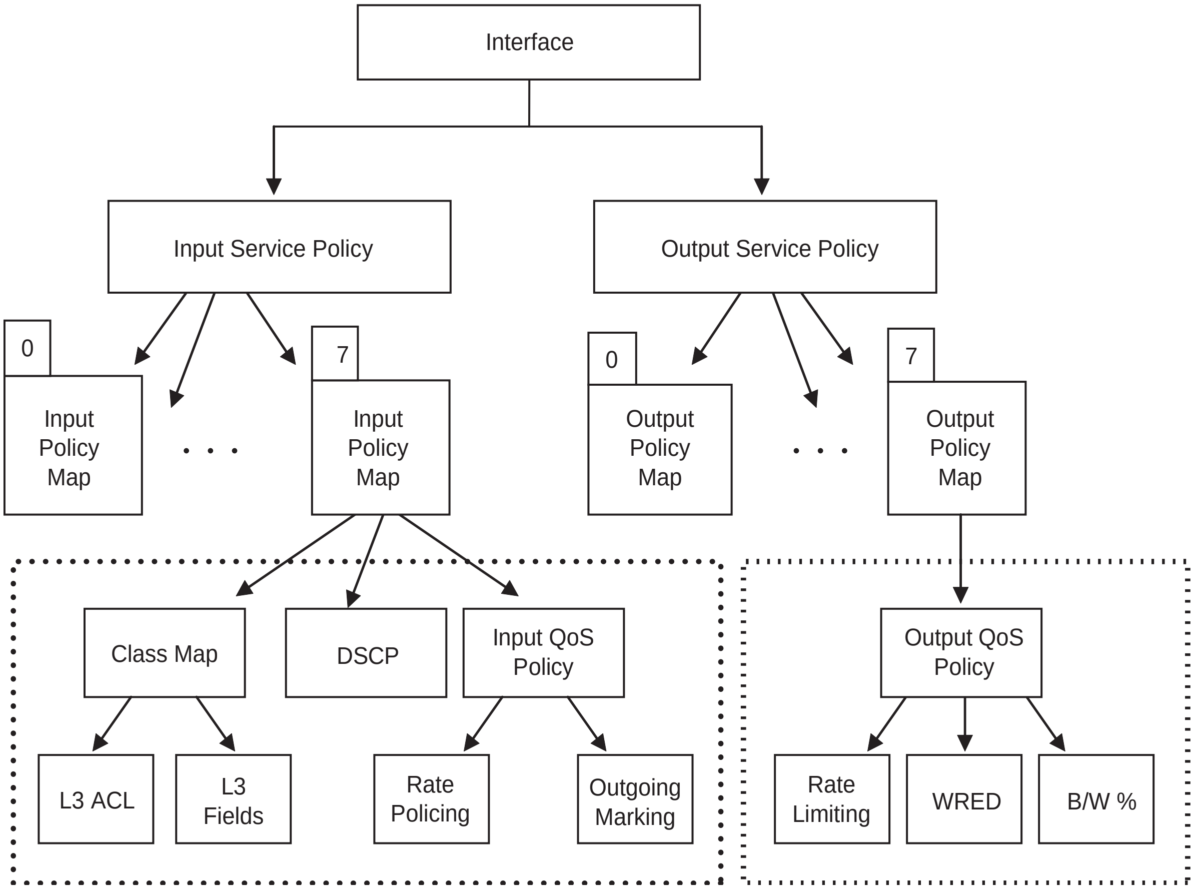

Figure 208 Policy-based QoS CLI Hierarchy

FTOS supports these policy QoS features:

- Traffic Classification on page 359
- Input/Output QoS Policies on page 364
- Input/Output Policy Maps on page 372
- WRED Profile on page 377
- Marking DSCP in Outgoing Packet on page 384

Traffic Classification

The Force10 system handles policy-based traffic classification by using class-maps. These maps classify unicast traffic to one of several classes—eight for E-Series and four for C-Series—used to define network flows. You can set up class-maps for each of the following match criteria:

- IP access lists

- IP DSCP (Differentiated Services Code Point)
- IP precedence

Additionally, FTOS enables you to match multiple class-maps and specify multiple match criteria.

Refer to the QoS Chapter in the *FTOS Command Line Interface Reference* for the commands necessary to set up any or all of the above class matches.

Configuration Task List

The following includes the configuration task list for QoS traffic classification:

- create a class-map on page 360
- configure a class-map on page 361

create a class-map

To create Layer 2 or Layer 3 (the default) class-map to match packets to a specified class, use this command:

Command Syntax	Command Mode	Usage
class-map {**match-all** | **match-any**} *class-map-name* [**layer2**]	CONFIGURATION (config-class-map)	Packets arriving at the input interface are checked against the match criteria, configured using this command, to determine if the packet belongs to that class. This command enables the class-map configuration mode. The parameters are: • **match-all:** Determines how packets are evaluated when multiple match criteria exist. Enter the keyword **match-all** to determine that the packets must meet all the match criteria in order to be considered a member of the class. • **match-any:** Determines how packets are evaluated when multiple match criteria exist. Enter the keyword **match-any** to determine that the packets must meet at least one of the match criteria in order to be considered a member of the class. • **class-map-name:** Enter a name of the class for the class-map in a character format (16 character maximum). • (OPTIONAL) **layer2:** Enter to specify a Layer 2 Class Map. This option is available for E-Series only. Default: Layer 3

To delete an existing class-map, use the **no class-map** {**match-all** | **match-any**} *class-map-name* [**layer2**] command. Figure 209 illustrates how to create a class-map.

```
Force10#config t
Force10(conf)#class-map match-any ClassMap05
Force10(conf)#end
Force10#
```

Figure 209 class-map Command Example

configure a class-map

To configure a class-map, use one of the following commands to set up the match criteria:

- match ip access-group on page 361
- QoS Exception Handling Rules on page 362
- match ip precedence on page 362
- match ip dscp on page 363

match ip access-group

To configure match criteria for a class-map based on the contents of the ACL (access control list), enter this command:

Command Syntax	Command Mode	Usage
match ip access-group *access-group-name*	conf-class-map	Enter the ACL name whose contents are used as the match criteria in determining if packets belong to the class specified by class-map. You must enter the **class-map** command before using this command. Once the class-map is identified, you can configure the match criteria. For **class-map match-any**, a maximum of five ACL match criteria is allowed. For **class-map match-all**, only one ACL match criteria is allowed.

To remove ACL match criteria from a class-map, use the **no match ip access-group** *access-group-name* command.

Figure 210 illustrates how to configure a class-map with an IP access-group as its matching criteria:

```
Force10#config
Force10(conf)#class-map match-any ClassMap05
Force10(conf-class-map)#match ip access-group aclgrp2
Force10(conf-class-map)#end
Force10#
```

Figure 210 class-map with ip access-group Command Example

QoS Exception Handling Rules

FTOS provides an order option (**order** *number*) to specify the QoS order of priority (rule) for an ACL entry. The order option determines where the ACL rule will sit within the CAM; allowing users to control the classifications rules.

The order number range is 0 to 254; where 0 is the highest priority and 254 is the lowest priority. That is, the lower order number, or rule, has a higher priority. The higher priority rule preempts the lower priority rule. If the order option is not configured, the ACL defaults to the lowest priority rule (255). When there are 2 rules having the same order, the Queue is used as the tie breaker. Rules with the same order and queue value are ordered according to their configuration-order.

match ip precedence

You can also use IP precedence for class-map matching criteria. To set up an IP precedence matching criteria, use the following command:.

Command Syntax	Command Mode	Usage
match ip precedence *ip-precedence-list*	CONFIGURATION conf-class-map	Enter the IP precedence value(s) that is to be the match criteria. Separate values by commas—no spaces (1,2,3) or indicated a list of values separated by a hyphen (1-3). Range: 0 to 7. Up to 8 precedence values can be matched in one match statement. For example, to indicate the IP precedence values 0 1 2 3 enter either the command match ip precedence 0-3 or match ip precedence 0,1,2,3.

To remove IP precedence as a match criteria, use the **no match ip precedence** *ip-precedence-list* command.

Note: Only one of the IP precedence values must be a successful match, not all of the specified IP precedence values need to match.

Figure 211 below illustrates how to configure a class-map with an IP precedence and match-any as its matching criteria:

```
Force10#config t
Force10(conf)#class-map match-any ClassMap04
Force10(conf-class-map)#match ip precedence 4
Force10(conf-class-map)#end
Force10#
```

Figure 211 class-map with IP precedence Command Example

match ip dscp

To configure a class-map to use a Differentiated Services Code Point (DSCP) value as its match criteria, enter this command.

Note: Only one of the IP DSCP values must be a successful match criterion, not all of the specified IP DSCP values need to match.

Command Syntax	Command Mode	Usage
match ip dscp *dscp-list*	CONFIGURATION conf-class-map	Enter the IP DSCP value(s) that is to be the match criteria. Separate values by commas—no spaces (1,2,3) or indicate a list of values separated by a hyphen (1-3). Range: 0 to 63. Up to 64 IP DSCP values can be matched in one match statement. For example, to indicate IP DCSP values 0 1 2 3 4 5 6 7 enter either the command **match ip dscp 0,1,2,3,4,5,6,7** or **match ip dscp 0-7**.

To remove a DSCP value as a match criteria, use the **no match ip dscp** *dscp-list* command.

Figure 212 illustrates how to configure a class-map using IP DSCP values as its matching criteria:

```
Force10#config t
Force10(conf)#class-map match-any ClassMap08
Force10(conf-class-map)#match ip dscp 8-11
Force10(conf-class-map)#end
Force10#
```

Figure 212 class-map with ip dscp Command Example

Show Commands

Use the EXEC privilege mode **show qos class-map** [*class-name*]:

```
Force10#show qos class-map

Class-map match-any CM01
  Match ip dscp 34

Force10#
```

Figure 213 show qos class-map Command Example

For more information on QoS classification and honoring commands, please see the *FTOS Command Line Interface Reference.*

Input/Output QoS Policies

You set up policy-based QoS by defining QoS policies and policy maps for both input and output queues, and class-maps for input queues. FTOS input QoS policies enable customers to regulate incoming traffic before scheduling it for processing by the backplane. These policies enable you to set up various input traffic conditioning mechanisms, including ingress rate policing.

FTOS output policies enable you to regulate outgoing traffic after FTOS schedules it for egress. The output policies available to you include output rate limits and congestion control mechanisms such as WFQ and WRED.

Configuration Task List

You can configure policy-based QoS for both ingress (input) and egress (output) queues.

To configure input QoS policy for ingress queues, perform these tasks:

- define input QoS policies on page 364 (optional)
- assign input aggregate policy to input policy maps on page 365 (optional)
- Rate-police Incoming Traffic on page 366 (optional)
- trust DSCP on page 373 (optional)

To configure output QoS policy for egress queues, perform these tasks:

- Output QoS Policies on page 367 (mandatory)
- assign output aggregate policy to output policy maps on page 367 (optional)
- rate-limit outgoing traffic on page 368 (optional)
- define rate shape of outgoing traffic on page 369 (optional)
- configure bandwidth percentages on page 369 (optional)
- specify WRED drop precedence on page 370 (optional)

define input QoS policies

To create an input QoS policy use this command:

Command Syntax	Command Mode	Usage
qos-policy-input qos-policy-name	CONFIGURATION	Enter your input QoS policy name in character format (16 characters maximum). Use this command to specify the name of the input QoS policy. Once input policy is specified, rate-police can be defined. This command enables the qos-policy-input configuration mode—(conf-qos-policy-in).

To remove an existing input QoS policy from the router, use the **no qos-policy-input qos-policy-name** command.

Figure 214 shows how to configure a QoS input policy.

```
Force10#config t
Force10(conf)#qos-policy-input QosPolicy25
Force10(conf-qos-policy-in)#end
Force10#
```

Figure 214 qos-policy-input Command Example

assign input aggregate policy to input policy maps

Use this command to assign an input aggregate policy to input policy-map:

Command Syntax	Command Mode	Usage
policy-aggregate *qos-policy-name*	policy-map-input	*qos-policy-name* Enter the name of the policy map in character format (16 characters maximum). This specifies the input QoS policy assigned in policy-map-input context.

To remove a policy aggregate configuration, use **no policy-aggregate** *qos-policy-name* command.

Figure 215 displays the assigning of an input aggregate policy to input policy maps.

```
Force10#config t
Force10(conf)#policy-map-input PolicyMapInput
Force10(conf-policy-map-in)#policy-aggregate QosPolicyInput
Force10(conf-policy-map-in)#end

Force10#
```

Figure 215 policy-aggregate Command Example

Rate-police Incoming Traffic

C-Series	NO
E-Series	✓

Platform Specific Feature: Rate-police Incoming Traffic is supported on the E-Series only.

FTOS supports policy-based policing of incoming traffic. To implement rate-policing, use this command:

Command Syntax	Command Mode	Usage
rate-police committed-rate [**burst-KB**] [**peak peak-rate** [**burst-KB**]]	conf-qos-policy-in	• *committed rate*: Enter the committed rate in Mbps. Range: 0 to 10000 Mbps. • **burst-KB**: (OPTIONAL) Enter the burst size in KB. Range: 16 to 200000 KB. Default: 50 KB. • **peak peak-rate**: (OPTIONAL) Enter the keyword **peak** followed by the peak rate in Mbps. Range 0 to 10000 Mbps. The default is the same as that for **committed-rate**.

To remove rate policing functionality, use the **no rate-police committed-rate** [**burst-KB**] [**peak peak-rate**] [**burst-KB**]] command.

Figure 216 shows how to configure a QoS input policy's rate-police setting:

```
Force10#config t
Force10(conf)#qos-policy-input QosPolicy25
Force10(conf-qos-policy-in)#rate-police 100 40 peak 150 50
Force10(conf-qos-policy-in)#end
Force10#
```

Figure 216 rate-police Command Example

Output QoS Policies

C-Series	NO
E-Series	✓

Platform Specific Feature: Output QoS Policies are supported on the E-Series only.

To create a output QoS policy on the router, use this command:

Command Syntax	Command Mode	Usage
qos-policy-output qos-policy-name	CONFIGURATION	Enter your output QoS policy name in character format (16 characters maximum). Use this command to specify the name of the output QoS policy. Once output policy is specified, rate-limit, bandwidth-percentage, and WRED can be defined. This command enables the qos-policy-output configuration mode—(conf-qos-policy-out).

To remove an existing output QoS policy from the router, use the **no qos-policy-output qos-policy-name** command.

Figure 217 demonstrates how to create an output QoS policy:

```
Force10#config t
Force10(conf)#qos-policy-output QosPolicy25
Force10(conf-qos-policy-out)#end
Force10#
```

Figure 217 qos-policy-output Command Example

assign output aggregate policy to output policy maps

Use this command to assign an output aggregate policy to output policy-map:

Command Syntax	Command Mode	Usage
policy-aggregate *qos-policy-name*	policy-map-output	*qos-policy-name* Enter the name of the policy map in character format (16 characters maximum). This specifies the output QoS policy assigned in policy-map-output context.

To remove a policy aggregate configuration, use **no policy-aggregate** *qos-policy-name* command.

Figure 218 displays the assigning of an output aggregate policy to output policy maps.

```
Force10#config t
Force10(conf)#policy-map-output PolicyMapOutput
Force10(conf-policy-map-out)#policy-aggregate QosPolicyOutput
Force10(conf-policy-map-out)#end
Force10#
```

Figure 218 policy-aggregate output Command Example

rate-limit outgoing traffic

You can configure FTOS to establish QoS policy-based rate-limits for outgoing traffic. To limit outgoing traffic, use this command:

Command Syntax	Command Mode	Usage
rate-limit committed-rate [**burst-KB**] [**peak peak-rate** [**burst-KB**]]	conf-qos-policy-out	• *committed-rate:* Enter the committed rate in Mbps. Range: 0 to 10000 Mbps. • **burst-KB**: (OPTIONAL) Enter the burst size in kilobytes. Range: 16 to 200000. The default is 50. • **peak peak-rate**: (OPTIONAL) Enter the keyword **peak** followed by the peak rate in Mbps. Range: 0 to 10000 Mbps. The default is the same as that for **committed-rate**.

To remove the rate limiting functionality, use the **no rate-limit committed-rate** [**burst-KB**] [**peak peak-rate** [**burst-KB**]] command.

Figure 219 demonstrates how to establish QoS policy-based rate-limits for outgoing traffic:

```
Force10#config t
Force10(conf)#qos-policy-output QosPolicy25
Force10(conf-qos-policy-out)#rate-limit 100 40 peak 150 50
Force10(conf-qos-policy-out)#end
Force10#
```

Figure 219 rate-limit Command Example

define rate shape of outgoing traffic

For each interface, you can also rate shape the outgoing traffic by configuring rate-shape in qos-policy-output configuration mode and applying it as an aggregate policy:

Command Syntax	Command Mode	Usage
rate-shape *rate* [*burst-KB*]	qos-policy-output configuration mode (conf-qos-policy-out)	*rate* Enter the outgoing rate in multiples of 10 Mbps. Range: 0 to 10000 *burst-KB* (OPTIONAL) Enter a number as the burst size in KB. Range: 0 to 10000 Default: 10

To delete the command, use the **no rate-shape** *rate* command.

Figure 220 is an example rate-shape configuration.

```
Force10#config t
Force10(conf)#qos-policy-output QosPolicyOutput
Force10(conf-qos-policy-out)#rate-shape 100 50
Force10(conf-qos-policy-out)#end

Force10#
```

Figure 220 rate-shape Command Example

configure bandwidth percentages

To assign a percentage of bandwidth to a class or queue in FTOS, use this command: :

Command Syntax	Command Mode	Usage
bandwidth-percentage percentage	conf-qos-policy-out	Enter the percentage assignment of bandwidth to this class or queue. Range: 0 to 100% (granularity 1%). The unit of bandwidth percentage is 1%. A bandwidth percentage of 0 is allowed and disables the scheduling of that class. If the sum of the bandwidth percentages given to all eight classes exceeds 100%, the bandwidth percentage automatically sets to 100%.

To remove the bandwidth percentage, use the **no bandwidth-percentage percentage** command.

Figure 221 demonstrates how configure a bandwidth percentage.

```
Force10#config t
Force10(conf)#qos-policy-output PolicyName25
Force10(conf-qos-policy-out)#bandwidth-percentage 10
Force10(conf-qos-policy-out)#end
Force10#
```

Figure 221 bandwidth-percentage Command Example

specify WRED drop precedence

Use this command to designate a WRED (Weighted Random Early Detection) profile for yellow or/and green traffic:

Command Syntax	Command Mode	Usage
wred {**yellow** | **green**} **wred-profile-name**	conf-qos-policy-out	Use this command to assign drop precedence to green or yellow traffic. If there is no honoring enabled on the input, all the traffic defaults to green drop precedence. • **yellow | green**: Enter the keyword **yellow** for yellow traffic. DSCP value of xxx110 and xxx100 maps to yellow. Enter the keyword **green** for green traffic. DSCP value of xxx010 maps to green. • Enter your WRED profile name in character format (16 character maximum) or enter one of the 5 pre-defined WRED profile names. Pre-defined profiles: wred_drop, wred-ge_y, wred_ge_g, wred_teng_y, wred_teng_g.

To remove the WRED drop precedence, use the **no wred** {**yellow** | **green**} [**profile-name**] command.

Figure 222 below shows how to specify WRED drop precedence:

```
Force10#config t
Force10(conf)#qos-policy-output QosPolicy26
Force10(conf-qos-policy-out)#wred yellow YellowProfile
Force10(conf-qos-policy-out)#wred green GreenProfile
Force10(conf-qos-policy-out)#end
Force10#
```

Figure 222 wred Command Example

Show Commands

To view input QoS policy settings, use the **show running qos-policy-input** [**qos-policy-name**] command in EXEC privilege mode:

```
Force10#
Force10#show running qos-policy-input
!
qos-policy-input QPN25
rate-police 100 40 peak 150 50
```

Figure 223 show running qos-policy-input Command Example

To see an input policy map, use the **show qos policy-map-input** [**policy-map-name**] [**class class-map-name**] [**qos-policy-input qos-policy-name**].

```
Force10#show qos policy-map-input
Policy-map-input PM04

Queue#   Class-map-name          Qos-policy-name
  6      CM03                    QPN1
```

Figure 224 show qos policy-map-input Command Example

To view an output QoS policy, use the **show qos qos-policy-output** [**qos-policy-name**] command.

```
Force10#show qos policy-map-input
Policy-map-input PM04

Queue#   Class-map-name          Qos-policy-name
  6      CM03                    QPN1
Force10#
```

Figure 225 show qos Command Example

To see a summary or detailed view of a QoS policy map, use the **show qos policy-map** {**summary** [**interface-name**] | **detail** [**interface-name**]} command.

```
Force10#show qos policy-map summary

Interface        policy-map-input        policy-map-output
Gi 0/0            -                           QPN1
Gi 0/9           QPN12                -
Gi 0/10          -                           QPN25
```

Figure 226 show qos policy-map summary Command Example

For more information about FTOS' implementation of policy QoS, refer to the *FTOS Command Line Reference*.

Input/Output Policy Maps

To configure input policy maps for ingress queues, perform these tasks:

- Input Policy Maps on page 372 (mandatory)
- trust DSCP on page 373 (optional)
- assign input policy maps to input queues on page 374 (mandatory)

To configure output policy maps for egress queues, perform these tasks:

- Output Policy Maps on page 375 (mandatory)
- assign output policy maps to output queues on page 376 (mandatory)
- apply output policy maps to interfaces on page 376 (mandatory)

Input Policy Maps

To define an input policy map, use this command:

Command Syntax	Command Mode	Usage
policy-map-input policy-map-name	CONFIGURATION	Enter the name for the policy map in character format (16 characters maximum). An input policy map is used to classify incoming traffic to different flows using class-map, QoS policy or simply using incoming packets DSCP. This command enables policy-map-input configuration mode (conf-policy-map-in).

To remove an input policy map, use the **no policy-map-input policy-map-name** command.

Figure 227 shows how to define an input policy map:

```
Force10#config t
Force10(conf)#policy-map-input PolicyMapInput
Force10(conf-policy-map-in)#end
Force10#
```

Figure 227 policy-map-input Command Example

trust DSCP

To define a dynamic classification for an input policy map to trust Differentiated Services Code Point (DSCP), use this command:

Command Syntax	Command Mode	Usage
trust diffserv	conf-policy-map-in	Specify dynamic classification to trust (DSCP). Dynamic mapping honors packets marked according to the standard definitions of DSCP. The default mapping table is detailed in Table on page 373

Note: When trust DSCP is configured, matched bytes/packets counters are not incremented in **show qos statistics**.

Table 27 Standard Default DSCP Mapping Table

DSCP/CP	DSCP Definition	Traditional IP Precedence	E-Series Internal Queue ID	C-Series Internal Queue ID
111XXX		Network Control	7	3
110XXX		Internetwork Control	6	3
101XXX	EF (Expedited forwarding) 101110	CRITIC/ECP	5	2
100XXX	AF4 (Assured Forwarding)	Flash Override	4	2
011XXX	AF3	Flash	3	1
010XXX	AF2	Immediate	2	1
001XXX	AF1	Priority	1	0
000XXX	BE: Best Effort	BE: Best Effort	0	0

To remove the definition, use the **no trust diffserv** command.

Figure 228 illustrates how to specify dynamic classification to trust:

```
Force10#config t
Force10(conf)#policy-map-input PolicyMapInput
Force10(conf-policy-map-in)#trust diffserv
Force10(conf-policy-map-in)#end
Force10#
```

Figure 228 trust diffserv Command Example

assign input policy maps to input queues

Use this command to assign a policy-map or class-map to the ingress queue.:

Command Syntax	Command Mode	Usage
service-queue queue-id [**class-map class-map-name**] [**qos-policy qos-policy-name**]	conf-policy-map-in	This command assigns class-map or qos-policy to different queues.

To remove the queue assignment, use the **no service-queue queue-id** [**class-map class-map-name**] [**qos-policy qos-policy-name**] command.

Figure 229 demonstrates how to assign a policy-map or class-map to the ingress queue:

```
Force10#config t
Force10(conf)#policy-map-input PolicyMapInput
Force10(conf-policy-map-in)#service-queue 1 qos-policy QosPolicy25
Force10(conf-policy-map-in)#end
```

Figure 229 service-queue Command Example

apply input policy maps to interfaces

Use this command to apply an input policy map to an interface.

Note: “service-policy" and "service-class" commands are not allowed simultaneously on an interface.

Note: Service-class input and output commands are not available when the interface is in "vlan-stack access" mode. This command works only when the interface is a normal L2 port.

Command Syntax	Command Mode	Usage
service-policy input policy-map-name	INTERFACE	Enter the name for the policy map in character format (16 characters maximum).

To remove the input policy map from the interface, use the **no service-policy input policy-map-name** command.

Note: You can attach the same input policy-map to one or more interfaces to specify the service-policy for those interfaces. You also can modify policy maps attached to interfaces.

Figure 230 demonstrates how to apply an input policy map to an interface:

```
Force10#config t
Force10(conf)#interface gigabitethernet 0/0
Force10(conf-if)#service-policy input PolicyMapInput
Force10(conf-if)#end
Force10#
```

Figure 230 service-policy input Command Example

Output Policy Maps

C-Series	NO
E-Series	✓

Platform Specific Feature: Output policy maps are supported on the E-Series only.

To set up an output policy map, use this command:.

Command Syntax	Command Mode	Usage
policy-map-output policy-map-name	CONFIGURATION	Enter the name for the policy map in character format (16 characters maximum). Output policy map is used to assign traffic to different flows using QoS policy. This command enables the policy-map-output configuration mode (conf-policy-map-out).

To remove an output policy map, use the **no policy-map-output policy-map-name** command.

Figure 231 shows how to define an output policy map:

```
Force10#config t
Force10(conf)#policy-map-output PolicyMapOutput
Force10(conf-policy-map-in)#end
Force10#
```

Figure 231 policy-map-output Command Example

assign output policy maps to output queues

To apply the output policy map to the egress queue.

Command Syntax	Command Mode	Usage
service-queue queue-id qos-policy qos-policy-name	conf-policy-map-out	• *queue-id*: Enter the value used to identify a queue. There are eight (8) queues per interface. Range: 0 to 7. • **qos-policy qos-policy-name**. (MANDATORY) Enter the keyword **qos-policy** followed by the QoS policy name assigned to the queue in character format (16 character maximum). This specifies the output QoS policy in policy-map-output context.

To remove the queue assignment, use the **no service-queue queue-id** [**qos-policy qos-policy-name**] command.

Figure 232 demonstrates how to apply an output policy to an egress queue:

```
Force10#config t
Force10(conf)#policy-map-output PolicyMapOutput
Force10(conf-policy-map-in)#service-queue 1 class-map ClassMap05 qos-policy QosPolicy25
Force10(conf-policy-map-in)#end
Force10#
```

Figure 232 service-queue Command Example

apply output policy maps to interfaces

Use this command to apply an output policy map to an interface:

Command Syntax	Command Mode	Usage
service-policy output policy-map-name	INTERFACE	Apply an output policy map to the interface. Enter the name for the policy map in character format (16 characters maximum).

To remove the output policy map from the interface, use the **no service-policy output policy-map-name** command.

Note: You can attach the same output policy-map to one or more interfaces to specify the service-policy for those interfaces. You also can modify policy maps attached to interfaces.

Figure 233 demonstrates how to apply a service policy to an output queue:

```
Force10#config t
Force10(conf)#interface gigabitethernet 0/1
Force10(conf-if)#service-policy output PolicyMapOutput
Force10(conf-if)#end
Force10#
```

Figure 233 service-policy output Command Example

WRED Profile

C-Series	NO
E-Series	✓

Platform Specific Feature: WRED Profile is supported on E-Serie only.

WRED (Weighted Random Early Detection) is a congestion avoidance mechanism. It works by monitoring traffic and discards packets if the congestion begins to increase. This, in turn, signals the source to slow down its transmission. The drop decision is based upon drop precedence (internally marked color) and programmed drop probability profiles.

WRED is designed primarily to work with TCP in IP internetwork environments.

See for more information about WRED drop profiles.

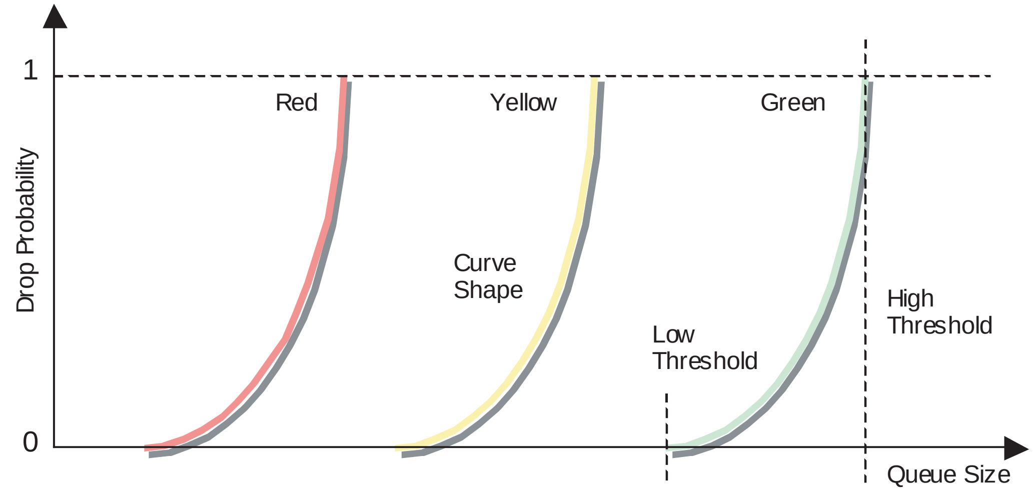

Figure 234 WRED Drop Profiles

Configuration Task List

To configure WRED, perform these tasks:

- define WRED profile on page 378
- specify minimum and maximum WRED thresholds on page 379

define WRED profile

To create a WRED profile, use this command:

Command Syntax	Command Mode	Usage
wred-profile wred-profile-name	CONFIGURATION	Enter your WRED profile name in character format (16 character maximum). Or use one of the pre-defined WRED profile names. You can configure up to 27 WRED profiles plus the 5 pre-defined profiles, for a total of 32 WRED profiles. Pre-defined Profiles: wred_drop, wred_ge_y, wred_ge_g, wred_teng_y, and wred_teng_g. When a new profile is configured, the minimum and maximum threshold defaults to predefined wred_ge_g values.

Table 28 Pre-defined WRED Profile Threshold Values

Default Profile Name	Minimum Threshold	Maximum Threshold
wred_drop	0	0
wred_ge_y	1000	2000
wred_ge_g	2000	4000
wred_teng_y	4000	8000
wred_teng_g	8000	16000

To remove an existing WRED profile, use the **no wred-profile** command.

Note: You cannot delete predefined WRED profiles.

Figure 235 shows how to set up a WRED profile:

```
Force10#config t
Force10(conf)#wred-profile Green-Profile
Force10(conf-wred)#end
Force10#
```

Figure 235 wred-profile wred-profile name Command Example

specify minimum and maximum WRED thresholds

Use this command to configure minimum and maximum threshold values for a user-defined WRED profile. The command can be used to modify the minimum and maximum threshold values for pre-defined WRED profiles.:

Command Syntax	Command Mode	Usage
threshold min min-threshold max max-threshold	conf-wred	Specify the minimum and maximum threshold values for the configured WRED profiles: • **min-threshold:** Enter the minimum threshold for the WRED profile. Range: 1024 to 77824 KB . • **max-threshold:** Enter the maximum threshold for the WRED profile. Range: 1024 to 77824 KB.

To remove the threshold values, use the **no threshold min min-threshold max max-threshold** command.

Figure 236 shows how to configure WRED threshold values:

```
Force10#config t
Force10(conf)#wred-profile Green-Profile
Force10(conf-wred)#threshold min 100 max 1000
Force10(conf-wred)#end
Force10#
```

Figure 236 threshold Command Example

Show Commands

To view the QoS WRED profiles and their threshold values, use the **show qos wred-profile** command (Figure 237) in the EXEC mode.

```
Force10#show qos wred-profile

Wred-profile-name       min-threshold   max-threshold
wred_drop               0               0
wred_ge_y               1000            2000
wred_ge_g               2000            4000
wred_teng_y             4000            8000
wred_teng_g             8000            16000
```

Figure 237 show qos wred-profile Command Example

To view the QoS statistics for WRED drops use the **show qos statistics** [**wred-profile** [interface-name]] command (Figure 238) in the EXEC mode.

```
 Force10#show qos statistics wred-profile
Interface Gi 5/11
Queue#  Drop-statistic  WRED-name        Min      Max      Dropped Pkts

  0     Green           WRED1            10       100      51623
        Yellow          WRED2            20       100      51300
        Out of Profile                                     0
  1     Green           WRED1            10       100      52082
        Yellow          WRED2            20       100      51004
        Out of Profile                                     0
  2     Green           WRED1            10       100      50567
        Yellow          WRED2            20       100      49965
        Out of Profile                                     0
  3     Green           WRED1            10       100      50477
        Yellow          WRED2            20       100      49815
        Out of Profile                                     0
  4     Green           WRED1            10       100      50695
        Yellow          WRED2            20       100      49476
        Out of Profile                                     0
  5     Green           WRED1            10       100      50245
        Yellow          WRED2            20       100      49535
        Out of Profile                                     0
  6     Green           WRED1            10       100      50033
        Yellow          WRED2            20       100      49595
        Out of Profile                                     0
  7     Green           WRED1            10       100      50474
        Yellow          WRED2            20       100      49522
        Out of Profile                                     0
Force10#
```

Figure 238 show qos statistics Command Example

For more information about FTOS' implementation of WRED, refer to the *FTOS Command Line Reference*.

WRED on Storm Control

C-Series	NO
E-Series	✓

Platform Specific Feature: WRED Profile is supported on the E-Series only.

Note: WRED on Storm Control is only available on Egress

Storm control limits the percentage of the total bandwidth that broadcast traffic can consume on an interface (if configured locally) or on all interfaces (if configured globally). For **storm-control broadcast 50 out**, the total bandwidth that broadcast traffic can consume on egress on a 1Gbs interface is 512Mbs. The method by which packets are selected to be dropped is the "tail-drop" method, where packets exceeding the specified rate are dropped.

WRED is a method of randomly dropping packets before congestion occurs. Traffic is a mixture of various kinds of packets. The rate at which some types of packets arrive might be greater than others. In this case, the space on the BTM (ingress or egress) can be consumed by only one or a few types of traffic, leaving no space for other types. A WRED profile can be applied to a policy-map so that specified traffic can be prevented from consuming too much of the BTM resources.

WRED uses a profile to specify minimum and maximum threshold values. The minimum threshold is the allotted buffer space for specified traffic, for example 1000KB on egress. If the 1000KB is consumed, packets will be dropped randomly at an exponential rate until the maximum threshold is reached; this is the "early detection" part of WRED. If the maximum threshold—2000KB, for example—is reached, then all incoming packets are dropped until less than 2000KB of buffer space is consumed by the specified traffic (Figure 239).

Before version 7.5.1.0, WRED could be applied only to unicast traffic using policy-maps. In version 7.5.1.0, WRED can be used in combination with storm control to regulate broadcast and unknown-unicast traffic. This feature is available through an additional option in command **storm-control** [**broadcast** | **unknown-unicast**] at CONFIGURATION. See the *FTOS Command Line Reference* for information on using this command.

Using the command **storm-control broadcast 50 out wred-profile**, for example, first the total bandwidth that broadcast traffic can consume is reduced to 50% of line rate. Even though broadcast traffic is restricted, the rate of outgoing broadcast traffic might be greater than other traffic, and if so, broadcast packets would consume too much buffer space. So, the **wred-profile** option is added to limit the amount of buffer space that broadcast traffic can consume.

Figure 239 Packet Drop Rate for WRED on Storm Control

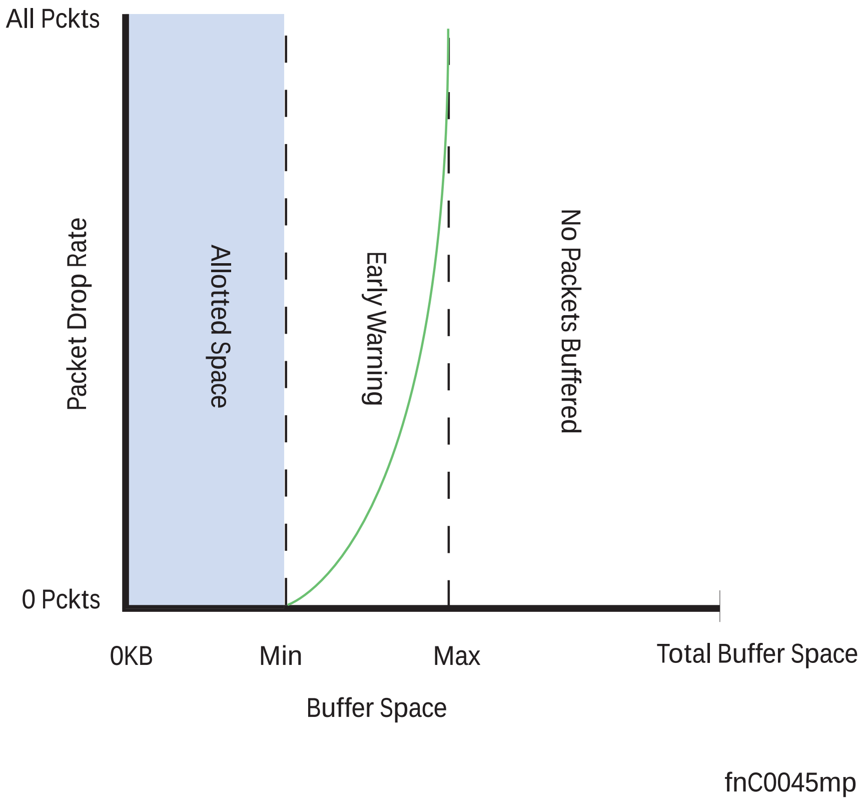

Weighted Fair Queuing for Multicast and Unicast

C-Series	NO
E-Series	✓

Platform Specific Feature: Weighted Fair Queuing for Multicast and Unicast is supported on the E-Series only.

At the egress port, traffic is scheduled using the WFQ algorithm on unicast, multicast, and replication traffic. Once a decision is made between unicast, multicast, or replication traffic, then a decision is made as to which queue is serviced for a selected traffic type. If unicast traffic is selected, unicast queues are serviced using Weighted Round-Robin or one queue can be configured as a Strict Priority—the remaining 7 queues use Weighted Round-Robin. Multicast queues are serviced using Round-Robin. Replication uses one queue.

The multicast queues are not used for QoS and are totally transparent to the user. Instead, these queues are used for internal groupings of egress ports to which multicast packets are transmitted.

Command Syntax

Use **queue egress multicast** command in CONFIGURATION mode:

Command Syntax	Usage
[no] queue egress multicast {linecard *slot number* **port-set** *number* **| all} {[wred-profile** *wred-profile-name* **] | [bandwidth-percent** *percent* **]}**	• Note: the **bandwidth-percentage** configuration is applicable only to egress multicast.

If 70% bandwidth is assigned to multicast and 80% bandwidth is assigned to one queue in unicast and 0% to all remaining queues in unicast, then the first 70% of the bandwidth is assigned to multicast. The 80% bandwidth for unicast is derived from the remaining 30% total bandwidth.

Configuration Steps

- Determine the required unicast bandwidth
- Determine the required multicast bandwidth
- Configure the multicast bandwidth
- Confirm your configuration

```
Force10(config)#do show run | grep mul
queue egress multicast linecard 1 port-set 0 bandwidth-percent 0
queue egress multicast linecard 1 port-set 0 wred-profile abc
```

Using the **no** command removes both the wred-profile and the multicast-bandwidth.

Note: The multicast-bandwidth option is not supported on queue ingress. If you try to use the multicast-bandwidth option for ingress multicast, the following reject error message is generated:

```
% Error:Bandwidth-percent is not allowed for ingress multicast
```

Points to Remember

- Multicast WFQ setting can be applied only on a per port pipe basis.
- If multicast bandwidth configured is 0, then control traffic going through multicast queues is dropped.
- The **no** form of the command without **bandwidth-percent** and **wred-profile**, removes both bandwidth and wred profile configuration.
- The **queue egress multicast** command works with multicast bandwidth configured only if the total unicast bandwidth is more than the multicast bandwidth.
- If strict priority is applied along with multicast bandwidth, the effect of strict priority is on all ports where unicast and multicast bandwidth are applied.

- When multicast bandwidth is assigned along with unicast bandwidth, first multicast bandwidth is reserved for that port, then the remaining unicast bandwidth configured is adjusted according to the bandwidth available after reserving for multicast.
- On 10 Gigabit line cards, if the number of cells in the IWFQ table assigned to multicast is greater than the unicast queues, then the bandwidth assignment does not take effect.
- Strict priority applied to one unicast queue affects traffic on another unicast queue of another port in the same port pipe with multicast bandwidth.
- If two egress ports are in the same port-pipe and have multicast and unicast bandwidth applied to one port only, traffic is not predictable on the other port.
- Unicast and multicast traffic is not calculable for jumbo frames.
- After sending multicast and unicast traffic with jumbo frames and stopping multicast traffic, bandwidth for unicast not to take effect.
- With bidirectional multicast traffic, bandwidth is affected.

Marking DSCP in Outgoing Packet

Marking means that the DSCP value in the outgoing packet is marked based on QoS classification. The 6 bits that are used for DSCP are also used for queue-id to which the traffic is destined. When marking is configured, the CLI generates an informational message advising to which queue the marking should be applied. If applied to a queue *other than* the one specified in the informational message, the first 3 bits in the DSCP are ignored and are replaced with the queue-id.

Command Syntax	Command Mode	Usage
set ip-dscp *dscp-value*	qos-policy-input configuration mode (conf-qos-policy-in)	*dscp-value* Enter the value to set the IP DSCP value. Range: 0 to 63

To remove a previously set IP DSCP value, use the **no set ip-dscp** *dscp-value* command.

Figure 240 displays a sample DSCP marking configuration. Notice the informational message (%Info:) with the queue that this policy should be applied.

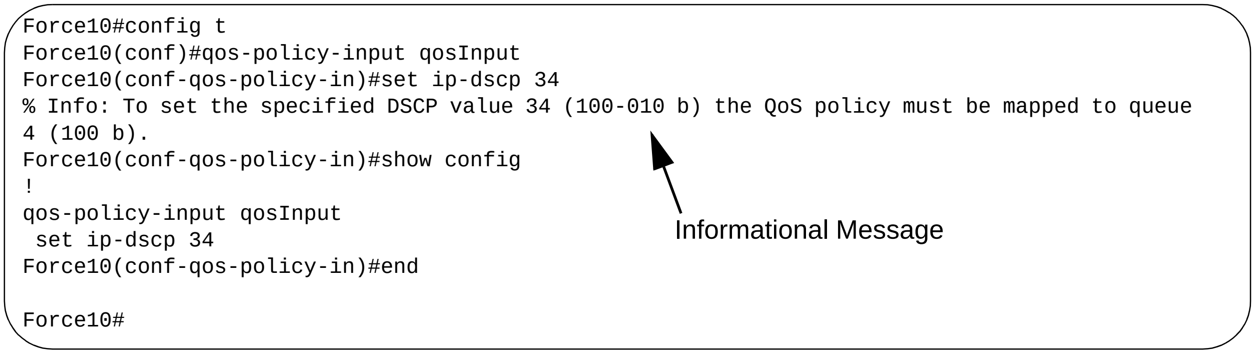

```
Force10#config t
Force10(conf)#qos-policy-input qosInput
Force10(conf-qos-policy-in)#set ip-dscp 34
% Info: To set the specified DSCP value 34 (100-010 b) the QoS policy must be mapped to queue
4 (100 b).
Force10(conf-qos-policy-in)#show config
!
qos-policy-input qosInput
 set ip-dscp 34
Force10(conf-qos-policy-in)#end

Force10#
```

Figure 240 Marking DSCP Configuration Example

Flow-based DSCP Marking

C-Series	NO
E-Series	✓

Platform Specific Feature: Flow-based DSCP Marking is supported on the E-Series only.

With class-maps, ingress traffic is classified according to ACLs, precedence values, or DSCP values. Once packets are classified, they can be marked using QoS input policies. The class-map together with the QoS input policy is applied to an interface and its set of queues.

There are two types of class-map rules, "match all" and "match any." With “match any” rules, multiple flows can be classified into the same queue with the same DSCP value.

Flow-based DSCP marking provides the ability to mark flows that are classified into the same queue of an ingress port with different DSCP values. To use this feature, append a “match any” rule with a DSCP value. This value overrides the QoS input policy DSCP value, and packets matching the rule are marked with the specified value.

To enable flow-based DSCP marking:

Task	Command Syntax	Command Mode
Specify a DSCP value for an access-group.	**match ip access-group** *access-group* **set-ip-dscp** *value*	CLASS-MAP
Specify a DSCP value for matching precedence flow.	**match ip access-group** *precedence* **set ip-dscp** *value*	CLASS-MAP
Specify a DSCP value for a matching DSCP flow.	**match ip dscp** *value* **set-ip-dscp** *value*	CLASS-MAP

Figure 241 Marking Flows in the Same Queue with Different DSCP Values

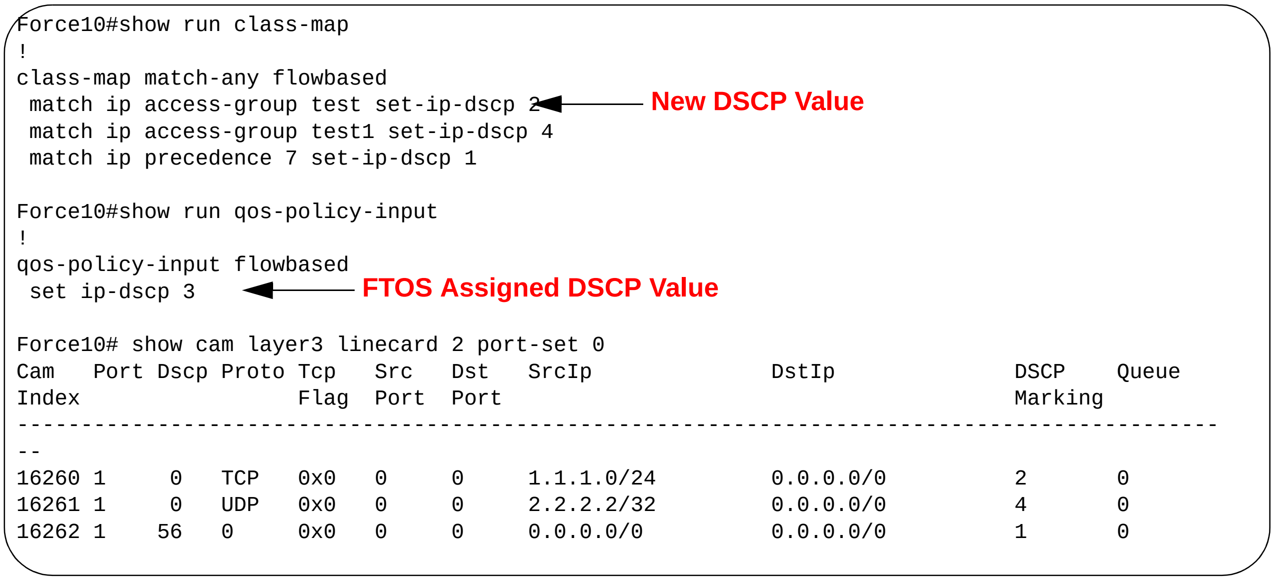

```
Force10#show run class-map
!
class-map match-any flowbased
 match ip access-group test set-ip-dscp 2
 match ip access-group test1 set-ip-dscp 4
 match ip precedence 7 set-ip-dscp 1

Force10#show run qos-policy-input
!
qos-policy-input flowbased
 set ip-dscp 3

Force10# show cam layer3 linecard 2 port-set 0
Cam   Port Dscp Proto Tcp   Src   Dst   SrcIp              DstIp              DSCP    Queue
Index                 Flag  Port  Port                                        Marking
----------------------------------------------------------------------------------------------
--
16260 1     0   TCP   0x0   0     0     1.1.1.0/24         0.0.0.0/0          2       0
16261 1     0   UDP   0x0   0     0     2.2.2.2/32         0.0.0.0/0          4       0
16262 1    56   0     0x0   0     0     0.0.0.0/0          0.0.0.0/0          1       0
```

New DSCP Value

FTOS Assigned DSCP Value

Pre-calculating Available QoS CAM Space

C-Series	NO
E-Series	✓

Platform Specific Feature: Pre-calculating Available QoS CAM Space is supported on the E-Series only.

Before version 7.3.1 there was no way to measure the number of CAM entries a policy-map would consume (the number of CAM entries that a rule uses is not predictable; 1 to 16 entries might be used per rule depending upon its complexity). Therefore, it was possible to apply to an interface a policy-map that requires more entries than are available. In this case, the system writes as many entries as possible, and then generates an CAM-full error message. The partial policy-map configuration might cause unintentional system behavior.

The command **test cam-usage** enables you to verify that there are enough available CAM entries *before* applying a policy-map to an interface so that you avoid exceeding the QoS CAM space and partial configurations. This command measures the size of the specified policy-map and compares it to the available CAM space in a partition for a specified port-pipe.

Test the policy-map size against the CAM space for a specific port-pipe or all port-pipes using these commands:

- **test cam-usage service-policy input** *policy-map* **linecard** *number* **port-set** *number*
- **test cam-usage service-policy input** *policy-map* **linecard** *all*

Note: See the FTOS Command Line Interface Reference for details on using this command and a description of the output.

The ouput of this command, shown in Figure 207, displays:

- the estimated number of CAM entries the policy-map will consume
- whether or not the policy-map can be applied
- the number of interfaces in a port-pipe to which the policy-map can be applied

Specifically:

- **Available CAM** is the available number of CAM entries in the specified CAM partition for the specified line card port-pipe.
- **Estimated CAM** is the estimated number of CAM entries that the policy will consume when it is applied to an interface.
- **Status** indicates whether or not the specified policy-map can be completely applied to an interface in the port-pipe.
    - **Allowed** indicates that the policy-map can be applied because the estimated number of CAM entries is less or equal to the available number of CAM entries. The number of interfaces in the port-pipe to which the policy-map can be applied is given in parenthesis.
    - **Exception** indicates that the number of CAM entries required to write the policy-map to the CAM is greater than the number of available CAM entries, and therefore the policy-map cannot be applied to an interface in the specifed port-pipe.

Note: The command **show cam-usage** provides much of the same information as **test cam-usage**, but whether or not a policy-map can be successfully applied to an interface cannot be determined without first measuring how many CAM entries the policy-map would consume; the command **test cam-usage** is useful because it provides this measurement.

Figure 242 test cam-usage Command Example

```
Force10# test cam-usage service-policy input pmap_l2 linecard 0 port-set 0

Linecard | Port-pipe | CAM Partition | Available CAM | Estimated CAM | Status
===============================================================================
0            0           L2ACL           500            200        Allowed(2)
```

UDP Broadcast

The UDP broadcast feature is a software based method of providing special treatment to forward low throughput (around 200 pps) IP/UDP broadcast traffic arriving on either physical or VLAN interface.

Enable the UDP broadcast feature using the command **ip udp-helper udp-port** from INTERFACE mode. This command enables the feature on an interface either for all UDP ports or a list of UDP ports upto a maximum of 16 ports.

In this section we talk about the behavior of UDP broadcast feature under different configuration and when incoming packet IP DA is IP broadcast or IP subnet broadcast or configured broadcast address. In all cases, source suppression is mandated.

UDP broadcast feature and broadcast address are configured

a)Incoming packet IP DA is as same as IP broadcast address (all FFs):

Send it on all L3 interfaces. However, IP DA will be modified only if

broadcast address is configured on the outgoing interface. The following

example illustrates this case:

Chapter 20 Policy-Based Routing

C-Series	NO
E-Series	✓

Platform Specific Feature: Policy-based Routing is supported on E-Series only.

Policy-Based Routing (PBR) enables you to apply routing policies to a specific interface. To enable PBR, you must first create a redirect list, then apply that list to a specific interface.

Redirect lists are defined by rules, or routing policies. The following parmeters can be defined in the routing policies or rules:

- IP address of the forwarding router (next-hop IP address)
- Protocol as defined in the header
- Source IP address and mask
- Destination IP address and mask
- Source port
- Destination port
- TCP Flags

Once a redirect-list is applied to an interface, all traffic passing through it is subjected to the rules defined in the redirect-list.

The traffic is forwarded based on the following order:

1. The existence of the next hop IP address that is specified in the redirect-list rule is verified. If the specified next hop is reachable, then the traffic is forwarded to the specified next hop.
2. If the specified next hop is not reachable, then the normal routing table is used to forward the traffic.

This chapter covers the following topics:

- Applying a Redirect List to an Interface on page 393
- PBR Exceptions on page 391

Creating a Redirect List

To create a redirect list, use the following command in REDIRECT-LIST mode:

Command Syntax	Command Mode	Purpose
[**no**] **ip-redirect-list** *redirect-list-name*	CONFIGURATION	Configure a redirect list. The **no** version of this command removes the redirect list.

The following example creates a redirect list by the name of “xyz”:

```
Force10(conf)#ip redirect-list ?
WORD                    Redirect-list name (max 16 chars)
Force10(conf)#ip redirect-list xyz
Force10(conf-redirect-list)#
```

Figure 243 Creating a Redirect List Example

Creating a Rule for a Redirect-list

To create a rule for a redirect list, use the following command in REDIRECT-LIST mode:

Command Syntax	Command Mode	Purpose
[**no**] **redirect** {*ip-address* | **sonet** *slot/port*} {*ip-protocol-number* | *protocol-type* [*bit*]} {*source mask* | **any** | **host** *ip-address*} {*destination mask* | **any** | **host** *ip-address*} [*operator*]	CONF-REDIRECT	Configure a rule for the redirect list.

In the following example shows a rule being created by identifying the IP address of the fowarding router, the IP protocol number, the source address with mask information, and the destination address with mask information:

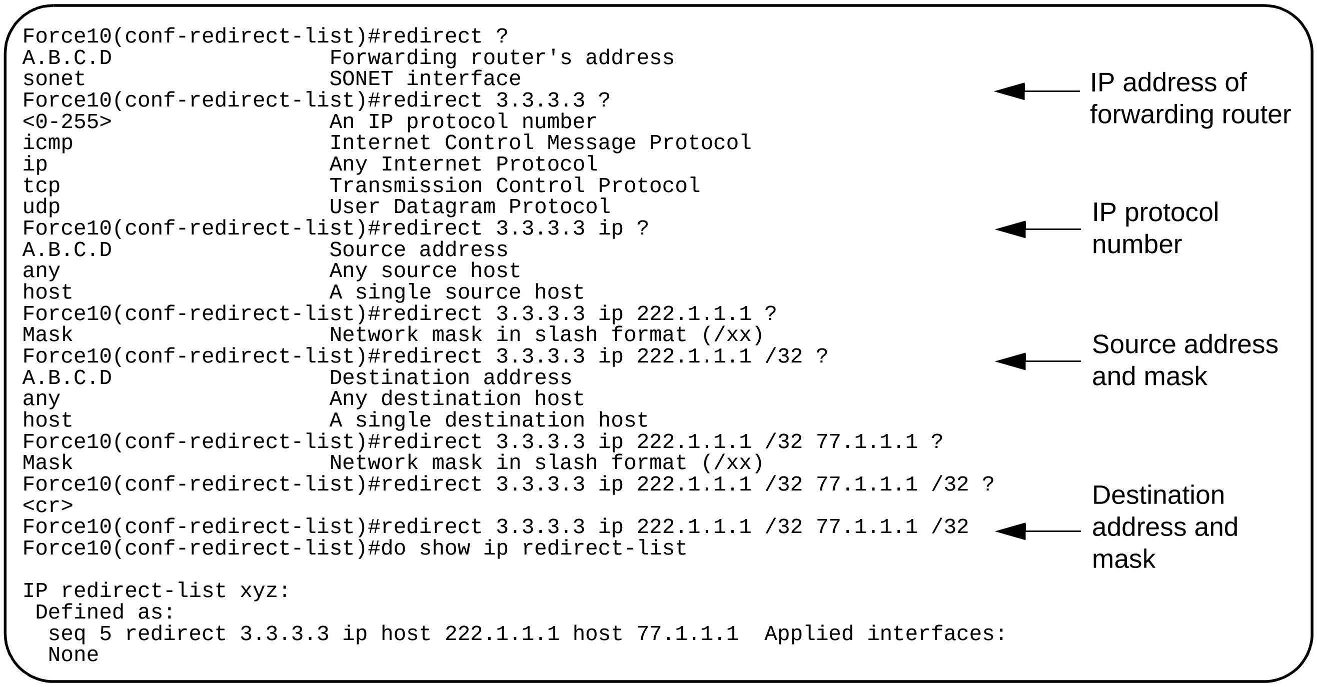

```
Force10(conf-redirect-list)#redirect ?
A.B.C.D                 Forwarding router's address
sonet                   SONET interface
Force10(conf-redirect-list)#redirect 3.3.3.3 ?
<0-255>                 An IP protocol number
icmp                    Internet Control Message Protocol
ip                      Any Internet Protocol
tcp                     Transmission Control Protocol
udp                     User Datagram Protocol
Force10(conf-redirect-list)#redirect 3.3.3.3 ip ?
A.B.C.D                 Source address
any                     Any source host
host                    A single source host
Force10(conf-redirect-list)#redirect 3.3.3.3 ip 222.1.1.1 ?
Mask                    Network mask in slash format (/xx)
Force10(conf-redirect-list)#redirect 3.3.3.3 ip 222.1.1.1 /32 ?
A.B.C.D                 Destination address
any                     Any destination host
host                    A single destination host
Force10(conf-redirect-list)#redirect 3.3.3.3 ip 222.1.1.1 /32 77.1.1.1 ?
Mask                    Network mask in slash format (/xx)
Force10(conf-redirect-list)#redirect 3.3.3.3 ip 222.1.1.1 /32 77.1.1.1 /32 ?
<cr>
Force10(conf-redirect-list)#redirect 3.3.3.3 ip 222.1.1.1 /32 77.1.1.1 /32
Force10(conf-redirect-list)#do show ip redirect-list

IP redirect-list xyz:
 Defined as:
  seq 5 redirect 3.3.3.3 ip host 222.1.1.1 host 77.1.1.1  Applied interfaces:
  None
```

Figure 244 Creating a Rule Example

PBR Exceptions

Use the command **permit** to create an exception to a redirect list. Exceptions are used when a forwarding decision should be based on the routing table rather than a routing policy.

FTOS assigns the first available sequence number to a rule configured without a sequence number and inserts the rule into the PBR CAM region next to the existing entries. Since the order of rules is important, ensure that you configure any necessary sequence numbers.

In Figure 245, the permit statement is never applied because the redirect list covers all source and destination IP addresses.

Figure 245 Ineffective PBR Exception due to Low Sequence Number

```
ip redirect-list rcl0
seq 5 redirect 2.2.2.2 ip any any
seq 10 permit ip host 3.3.3.3 any
```

To ensure that the permit statement or PBR exception is effective, use a lower sequence number, as shown in Figure 246.

Figure 246 Effective PBR Exception due to Proper Sequencing

```
ip redirect-list rcl0
seq 10 permit ip host 3.3.3.3 any
seq 15 redirect 2.2.2.2 ip any any
```

Showing Redirect List Configuration

To view the configuration redirect list configuration, use the following command in EXEC mode:

Command Syntax	Command Mode	Purpose
show ip redirect-list *redirect-list-name*	EXEC	View the redirect list configuration and the associated interfaces.
show cam pbf **show cam-usage**	EXEC	View the redirect list entries programmed in the CAM.

List the redirect list configuration using the **show ip redirect-list** *redirect-list-name* command:

```
Force10#show ip redirect-list xyz

IP redirect-list xyz:
 Defined as:
  seq 5 redirect 3.3.3.3 ip host 222.1.1.1 host 77.1.1.1
Applied interfaces:
  None
```

Figure 247 Showing Redirect List Configuration Example

Use the **show ip redirect-list** (without the list name) to display all the redirect-lists configured on the device:

```
Force10#show ip redirect-list
IP redirect-list xyz:
 Defined as:
  seq 5 redirect 3.3.3.3 ip host 222.1.1.1 host 77.1.1.1
 Applied interfaces:
  None

IP redirect-list abc:
 Defined as:
seq 5 redirect 2.2.2.2 ip host 333.1.1.1 host 88.1.1.1
Applied interfaces:
  None
```

Figure 248 Showing Redirect List Configuration Example 2

Applying a Redirect List to an Interface

IP redirect lists are supported on physical interfaces as well as VLAN and port-channel interfaces.

To apply a redirect list to an interface, use the following command in INTERFACE mode:

Command Syntax	Command Mode	Purpose
[**no**] **ip redirect-group** *redirect-list-name*	INTERFACE	Apply a redirect list (policy-based routing) to an interface. You can apply multiple redirect lists to an interface by entering this command multiple times.

In this example, the list “xyz” is applied to the GigabitEthernet 4/0 interface.

```
Force10(conf-if-gi-4/0)#ip redirect-group xyz
Force10(conf-if-gi-4/0)#
```

Figure 249 Applying a Redirect List to an Interface Example

Note: If, the redirect-list is applied to an interface, the output of **show ip redirect-list** *redirect-list-name* command displays reachability and ARP status for the specified next-hop.

```
Force10#show ip redirect-list xyz
IP redirect-list xyz:
 Defined as:
  seq 5 redirect 3.3.3.3 ip host 222.1.1.1 host 77.1.1.1, Next-hop reachable, ARP resolved
 Applied interfaces:
  Gi 4/0
```

Figure 250 Showing Redirect List Configuration Example after Applying to Interface

Chapter 21 VRRP

Virtual Router Redundancy Protocol (VRRP) is designed to eliminate a single point of failure in a statically routed network. This protocol is defined in RFC 2338 and RFC 3768.

This chapter covers the following topics:

- VRRP Overview on page 395
- VRRP Benefits on page 397
- VRRP Implementation on page 397
- VRRP Configuration on page 398

VRRP Overview

In the most basic terms, VRRP specifies a MASTER router to own the next hop IP and MAC address for end stations on a LAN. The MASTER router is chosen from the virtual routers by an election process and forwards packets sent to the next hop IP address. If the MASTER router fails, VRRP begins the election process to choose a new MASTER router and continues routing traffic.

VRRP uses the Virtual Router Identifier (VRID) to identify each virtual router configured. Each virtual router contains the IP addresses. Of the routers whose IP addresses are configured in a virtual router, one router is elected as the MASTER router. The IP address of the MASTER router is used as the next hop address for all end stations on the LAN. The other routers represented by IP addresses are BACKUP routers. One of the BACKUP routers will transition into the MASTER router if the current MASTER router goes down.

VRRP packets are transmitted with the virtual router MAC address as the source MAC address to ensure that learning bridges correctly determine to which LAN segment the virtual router is attached. The MAC address is in the following format: 00-00-5E-00-01-{VRID}. The first three octets are unchangeable. The next two octets (00-01) indicate the address block assigned to the VRRP protocol, and are unchangeable. The final octet changes depending on the VRRP Virtual Router Identifier and allows for up to 255 VRRP routers on a network.

Figure 251 shows a typical network configuration using VRRP. Instead of configuring the hosts on the network 10.10.10.0 with the IP address of either Router A or Router B as their default router; their default router is the virtual router (IP address 10.10.10.3). When any host on the LAN segment wants to access the Internet, it sends packets to the IP address of the virtual router.

Router A is configured as the MASTER router. It is configured with the IP address of the virtual router and sends any packets addressed to the virtual router through interface GigabitEthernet 1/1 to the Internet. As the BACKUP router, Router B is also configured with the IP address of the virtual router. If for any reason Router A stops transferring packets, VRRP converges, and Router B assumes the duties of Router A and becomes the MASTER router. At that time, Router B responds to the packets sent to the virtual IP address.

All workstations continue to use the IP address of the virtual router to address packets destined to the Internet. Router B receives and forwards them on interface GigabitEthernet 10/1. Until Router A resumes operation, VRRP allows Router B to provide uninterrupted service to the users on the LAN segment accessing the Internet.

While it is the MASTER router, Router B continues to perform its normal function: handling packets between the LAN segment and the Internet.

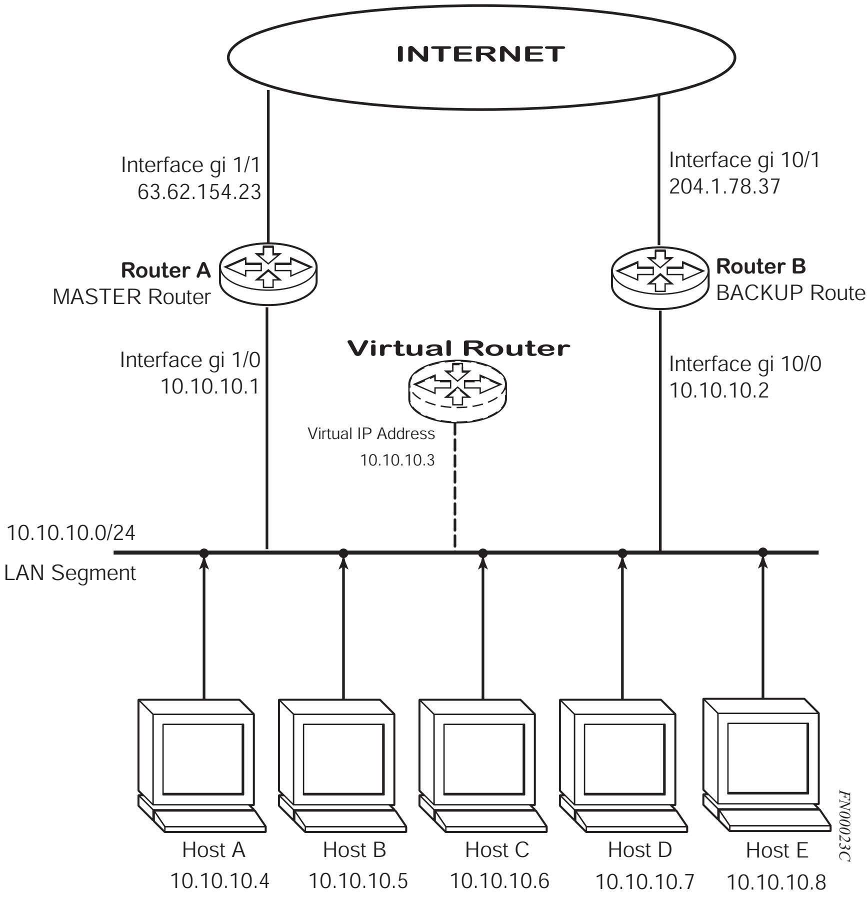

Figure 251 Basic VRRP Configuration

For more information on VRRP, refer to RFC 2338, *Virtual Router Redundancy Protocol*.

VRRP Benefits

With VRRP configured on a network, end-station connectivity to the network is not subject to a single connection. End-station connections to the network are redundant and they are not dependent on IGP protocols to converge or update routing tables.

VRRP Implementation

➡ **Note:** The feature 1500 VRRP groups is not supported in Version 6.2.1.0.

E-Series supports an unlimited number of VRRP groups, and up to 255 VRRP groups on a single interface. C-Series supports 128 VRRP groups; the number of groups per interface varies, as shown in Table 29. Within a single VRRP group, up to 12 virtual IP addresses are supported. Virtual IP addresses can belong to either the primary or secondary IP address' subnet configured on the interface. You can ping all the virtual IP addresses configured on the Master VRRP router from anywhere in the local subnet.

Though FTOS supports unlimited VRRP groups, certain inherent factors affect the maximum number of groups that can be configured and expected to work properly, the main factor being the throttling of VRRP advertisement packets reaching the RP2 processor. To avoid throttling of VRRP advertisement packets, Force10 recommends you to increase the VRRP advertisement interval to a higher value from the default value of 1 second. The recommendations are as follows:

Table 29 Recommended VRRP Advertise Intervals

	Recommended Advertise Interval		**Groups/Interface**	
Total VRRP Groups	**E-Series**	**C-Series**	**E-Series**	**C-Series**
Less than 250	1 second	1 second	255	12
Between 250 and 450	2 second	2 - 3 second	255	24
Between 450 and 600	3 second	4 second	255	36
Between 600 and 800	4 second	5 second	255	48
Between 800 and 1000	5second			
Between 1000 and 1200	6 second			
Between 1200 and 1500	7 second			
Beyond 1500	8 or more seconds			

Please note that the above recommendations are only indicative, and you must take into account the overall traffic pattern in the network (like ARP broadcasts, IP broadcasts, STP, etc.) before changing the advertisement interval. When the number of packets processed by RP2 processor increases or decreases based on the dynamics of the network, the advertisement intervals in Table 29 may increase or decrease accordingly.

Caution: Increasing the advertisement interval increases the VRRP Master dead interval, resulting in an increased failover time for Master/Backup election. The user is implored upon to take extra caution when increasing the advertisement interval, as the increased dead interval causes packets to be dropped during that switch over time.

VRRP Configuration

To configure VRRP, use the commands in the VRRP mode to configure VRRP for 1-Gigabit Ethernet, 10-Gigabit Ethernet, VLAN, and port channel interfaces.

By default, VRRP is not configured on the E-Series.

Configuration Task List for VRRP

The following list includes the configuration tasks for VRRP:

- create a virtual router on page 398 (mandatory)
- assign virtual IP addresses on page 399 (mandatory)
- set priority for the vrrp group on page 401 (optional)
- configure authentication for VRRP on page 402 (optional)
- enable preempt on page 403 (optional)
- change the advertisement interval on page 404 (optional)
- track an interface on page 404 (optional)

For a complete listing of all commands related to VRRP, refer to .

create a virtual router

To enable VRRP, you must create a virtual router. In FTOS, a virtual router is called a VRRP group, and the first step in creating a virtual router is to assign a Virtual Router Identifier (VRID).

The interface containing the VRRP group must be enabled and configured with an primary IP address.

To create a virtual router, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
vrrp-group *vrid*	INTERFACE	Create a virtual router for that interface with a VRID. The VRID will appear in the VRRP mode prompt.

To view the VRRP group or virtual router, use the **show config** command (Figure 252) in either the INTERFACE or VRRP mode.

```
Force10(conf-if-vrid-3)#show config
 vrrp-group 3
Force10(conf-if-vrid-3)#
```

Figure 252 show config Command Example in the VRRP Mode

The **show config** command displays non-default values.

Virtual routers contain virtual IP addresses configured for that VRRP group, in addition to other configuration information. A VRRP group does not transmit VRRP packets until you assign the virtual IP address to the VRRP group.

To delete a VRRP group, use the **no vrrp-group** *vrid* command in the INTERFACE mode.

assign virtual IP addresses

FTOS supports up to 1500 VRRP groups on one system, and 12 VRRP groups on a single interface. Within a single VRRP group, up to 12 virtual IP addresses are supported. Virtual IP addresses can either belong to the primary or secondary IP addresses configured on the interface on which VRRP is enabled. You can ping all the virtual IP addresses configured on the Master VRRP router from anywhere in the local subnet.

To activate a VRRP group on an interface (that is, the VRRP groups starts transmitting VRRP packets), enter the VRRP mode and configure at least one virtual IP address in a VRRP group. The virtual IP address is the IP address of virtual routers and does not include the IP address mask. You can configure up to 12 virtual IP addresses per VRRP group.

You can ping the virtual IP addresses to debug and test reachability.

The following rules apply to virtual IP addresses:

- The virtual IP addresses must belong to either the primary or secondary IP addresses configured on the interface. Though a single VRRP group can contain virtual IP addresses belonging to multiple IP subnets configured on the interface, Force10 recommends you configure virtual IP addresses belonging to the SAME IP subnet for any one VRRP group.

As an example, lets assume an interface (on which VRRP is to be enabled) contains a primary IP address of 50.1.1.1/24 and a secondary IP address of 60.1.1.1/24. And the you intends to configure 4 VRRP groups (VRID 1, VRID 2, VRID 3 and VRID 4) on this interface. VRID 1 should contain virtual addresses belonging to EITHER subnet 50.1.1.0/24 OR subnet 60.1.1.0/24, but NOT from both subnets (though FTOS allows the same). The same rule applies to VRID 2, 3 and 4.

- The virtual IP address assigned in a VRRP group can be the same as the interface's primary or secondary IP address under certain conditions, but the virtual IP address cannot be the same as any other IP address configured on the E-Series, including the virtual IP address for a VRRP group on another interface.
- If the virtual IP address and the interface's primary/secondary IP address are the same, the priority on that VRRP group MUST be set to 255. The interface then becomes the OWNER router of the VRRP group and the interface's physical MAC address is changed to that of the owner VRRP group's MAC address.
- If you have multiple VRRP groups configured on an interface, only one of the VRRP groups can contain the interface primary or secondary IP address.

To configure a virtual IP address, use these commands in the following sequence in the INTERFACE mode.

Step	Command Syntax	Command Mode	Purpose
1	**vrrp-group** *vrrp-id*	INTERFACE	Configure a VRRP group. The range of *vrrp-id* is 1 to 255.
2	**virtual-address** *ip-address1* [...*ip-address12*]	VRRP	Configure up to 12 virtual IP addresses of virtual routers.

To view the VRRP group configuration, use the **show config** command in the VRRP mode or the **show vrrp brief** command (Figure 253) in the EXEC privilege mode.

```
Force10(conf)#do show vrrp brief

Interface Grp Pri  Pre State  Master addr     Virtual addr(s)                    Description
---------------------------------------------------------------------------------------------
Gi 2/1     1   100  Y   Na/If  Unknown         2.2.5.4                            This a description
Force10(conf)#
```

Figure 253 show vrrp brief Command Example

Figure 254 shows the same VRRP group configured on multiple interfaces on different subnets. Note that the virtual addresses are different.

```
Force10>show vrrp
------------------
GigabitEthernet 12/3, VRID: 1, Net: 10.1.1.253
State: Master, Priority: 105, Master: 10.1.1.253 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Adv sent: 1862, Gratuitous ARP sent: 0
Virtual MAC address:
 00:00:5e:00:01:01
Virtual IP address:
 10.1.1.252
Authentication: (none)
Tracking states for 1 interfaces:
 Up  GigabitEthernet 12/17 priority-cost 10
------------------
GigabitEthernet 12/4, VRID: 2, Net: 10.1.2.253
State: Master, Priority: 110, Master: 10.1.2.253 (local)
Hold Down: 10 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Adv sent: 1862, Gratuitous ARP sent: 0
Virtual MAC address:
 00:00:5e:00:01:02
Virtual IP address:
 10.1.2.252
Authentication: (none)
Tracking states for 2 interfaces:
 Up  GigabitEthernet 2/1 priority-cost 10
 Up  GigabitEthernet 12/17 priority-cost 10
Force10>
```

Figure 254 show vrrp Commands Example

When the VRRP process completes its initialization, the State field contains either Master or Backup.

set priority for the vrrp group

When you set the priority of a virtual router to 255 (see assign virtual IP addresses on page 399), that virtual router becomes the OWNER virtual router for the VRRP group. VRRP elects the MASTER router by choosing the router with the highest priority. You configure the priority of the virtual router or you can leave it at the default value of 100.

To configure the priority of a VRRP group, use the following command in the VRRP mode:

Command Syntax	Command Mode	Purpose
priority *priority*	VRRP	Configure the priority for the VRRP group. Default: 100. Range: 1 to 255

If two routers in a VRRP group come up at the same time and contain the same priority value, the interface’s physical IP addresses are used as tie-breakers to decide which is MASTER. The router with the higher IP address will become MASTER.

To view the priority of virtual groups, use the **show vrrp** command in the EXEC privilege mode.

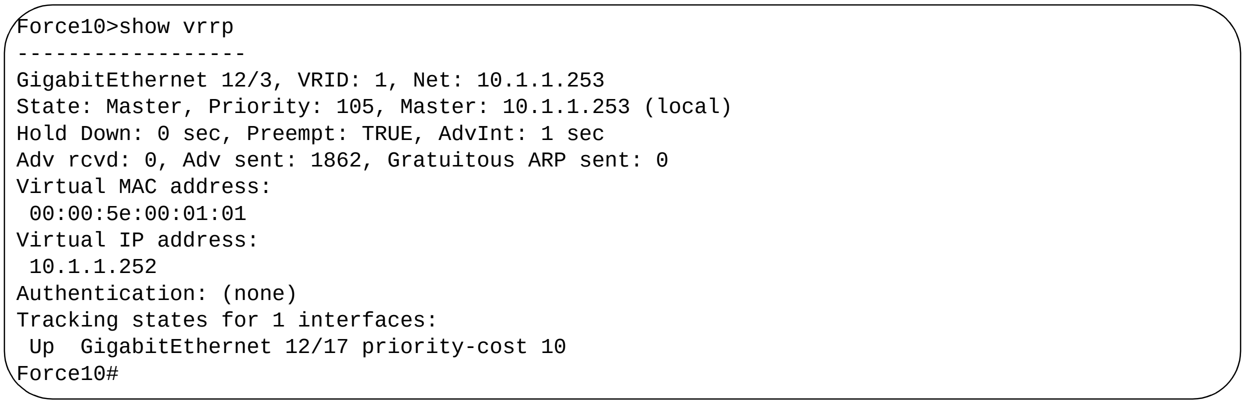

```
Force10>show vrrp
------------------
GigabitEthernet 12/3, VRID: 1, Net: 10.1.1.253
State: Master, Priority: 105, Master: 10.1.1.253 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Adv sent: 1862, Gratuitous ARP sent: 0
Virtual MAC address:
 00:00:5e:00:01:01
Virtual IP address:
 10.1.1.252
Authentication: (none)
Tracking states for 1 interfaces:
 Up  GigabitEthernet 12/17 priority-cost 10
Force10#
```

Figure 255 show vrrp Command Example

configure authentication for VRRP

Simple authentication of VRRP packets ensures that only trusted routers participate in VRRP processes. When you configure authentication, FTOS includes the password in its VRRP transmission and the receiving router uses that password to verify the transmission.

All virtual routers in the VRRP group must all be configured the same; either authentication is enabled with the same password or it is disabled.

To configure simple authentication, use the following command in the VRRP mode:

Command Syntax	Command Mode	Purpose
authentication-type simple [*encryption-type*] *password*	VRRP	Configure a simple text password. You can set the optional encryption-type to encrypt the password.

To view the password, use the **show config** command in the VRRP mode or the **show vrrp** command in EXEC privilege mode (Figure 256).

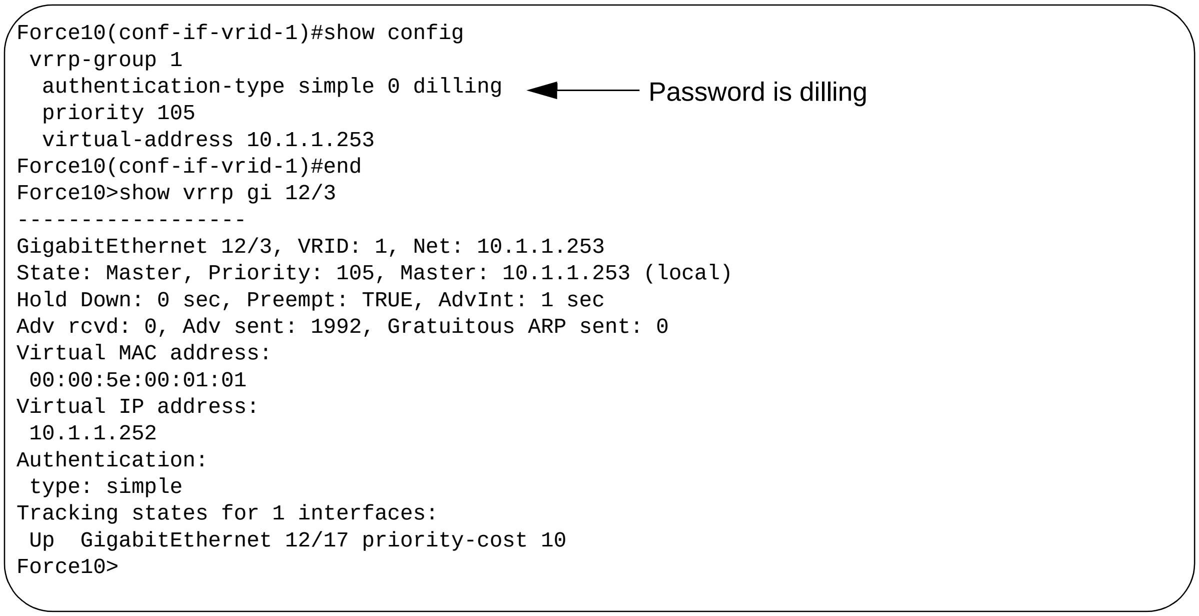

```
Force10(conf-if-vrid-1)#show config
 vrrp-group 1
  authentication-type simple 0 dilling
  priority 105
  virtual-address 10.1.1.253
Force10(conf-if-vrid-1)#end
Force10>show vrrp gi 12/3
------------------
GigabitEthernet 12/3, VRID: 1, Net: 10.1.1.253
State: Master, Priority: 105, Master: 10.1.1.253 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Adv sent: 1992, Gratuitous ARP sent: 0
Virtual MAC address:
 00:00:5e:00:01:01
Virtual IP address:
 10.1.1.252
Authentication:
 type: simple
Tracking states for 1 interfaces:
 Up  GigabitEthernet 12/17 priority-cost 10
Force10>
```

Password is dilling

Figure 256 show config and show vrrp Command Examples with a Simple Password Configured

enable preempt

To force FTOS to change the MASTER router if a BACKUP router with a higher priority comes online, use the **preempt** command. This function is enabled by default.

You can prevent the BACKUP router with the higher priority from becoming the MASTER router by disabling the preempt function.

All virtual routers in a VRRP group must be configured the same; either all configured with preempt enabled or configured with preempt disabled.

To disable the preempt function, use the following command in the VRRP mode:

Command Syntax	Command Mode	Purpose
no preempt	VRRP	Prevent any BACKUP router with a higher priority from becoming the MASTER router.

To view the virtual router's configuration for preempt, use the **show vrrp brief** command in the EXEC privilege mode. If the fourth column from the left contains a Y, then the preempt function is configured for the VRRP group.

```
Force10>show vrrp brief
Interface Grp Pri Pre  State   Master addr   Virtual addr(s)
------------------------------------------------------------
Gi 12/3   1   105  Y Master 10.1.1.253  10.1.1.252
Gi 12/4   2   110  Y Master 10.1.2.253  10.1.2.252
Force10>
```

Figure 257 show vrrp brief Command Example

change the advertisement interval

Every second the MASTER router transmits a VRRP advertisement to all members of the VRRP group indicating it is up and is the MASTER router. If the VRRP group does not receive an advertisement, then election process begins and the BACKUP virtual router with the highest priority transitions to MASTER.

Note: Force10 Networks recommends that you keep the default setting for this command. If you do change the time interval between VRRP advertisements on one router, you must change it on all participating routers.

To change that advertisement interval, use the following command in the VRRP mode:

Command Syntax	Command Mode	Purpose
advertise-interval *seconds*	VRRP	Change the default setting. The range for the *seconds* value is from 1 to 255 seconds. Default: 1 second

track an interface

You can set FTOS to monitor the state of any interface by a virtual group. Each VRRP group can track up to 12 interfaces, which may affect the priority of the VRRP group. If the state of the tracked interface goes down, the VRRP group's priority is decreased by a default value of 10 (also known as cost). If the tracked interface's state goes up, the VRRP group's priority is increased by 10.

The lowered priority of the VRRP group may trigger an election. As the Master/Backup VRRP routers are selected based on the VRRP group's priority, tracking features ensure that the best VRRP router is the Master for that group. The sum of all the costs of all the tracked interfaces should not exceed the configured priority on the VRRP group. If the VRRP group is configured as Owner router (priority 255), tracking for that group is disabled, irrespective of the state of the tracked interfaces. The priority of the owner group always remains at 255.

To track an interface, use the following command in the VRRP mode:

Command Syntax	Command Mode	Purpose
track *interface* [**priority-cost** *cost*]	VRRP	Monitor an interface and, optionally, set a value to be subtracted from the interface's VRRP group priority. **priority-cost** *cost* range: 1 to 254. The default is 10.

The sum of all the costs for all tracked interfaces must be less than or equal to the configured priority of the VRRP group.

To view the current configuration, use the **show config** command in the VRRP or INTERFACE mode.

Chapter 22 Resilient Ring Protocol

C-Series	NO
E-Series	✓

Platform Specific Feature: Resilient Ring Protocol is supported on E-Series only.

Even with optimizations, the Spanning Tree Protocol (STP) can take up to 50 seconds to converge and a topology change (depending on the size of network and node of failure) requires 4 to 5 seconds to reconverge. These times are not acceptable for campus or service provider networks that require a sub-second reconvergence to carry their mission-critical or voice/video/data services. A voice call drops if it is subjected to more than 250 msec delay. Similarly, more than a one-second delay can make the video jitter and several mission-critical applications can suffer.

To operate a deterministic network, a network administrator must run a protocol that converges independently of the network size or node of failure. The Force10 Resilient Ring Protocol (FRRP) is a proprietary protocol that provides this flexibility, while preventing Layer 2 loops. FRRP provides sub-second ring-failure detection and convergence/re-convergence in a Layer 2 network while eliminating the need for running spanning-tree protocol. With its two-way path to destination configuration, FRRP provides protection against any single link/switch failure and thus provides for greater network uptime.

FRRP transmits a high-speed token across a ring to verify the link status. The master node provides all the intelligence—transit nodes require very little intelligence.

FRRP provides a convergence time that can generally range between 150ms and 400ms for Layer 2 networks. The master node originates a high-speed frame that circulates around the ring. This frame, appropriately, sets up or breaks down the ring.

- Ring Health Frames circulate the ring continuously
- Multiple rings can be run on the same switch
- One master node per ring—all other nodes are transit
- Master node states—blocking, pre-forwarding, forwarding
- Transit node states—pre-forwarding, forwarding
- Each node has 2 member interfaces—primary, secondary
- STP disabled on ring interfaces
- Master node secondary port is in blocking state
- Ring Health Frames (RHF)
  - Hello RHF

— Sent @ 50ms (hello interval)

— Processed at master node only

— Topology Change RHF

— Triggered updates

— Processed at all nodes

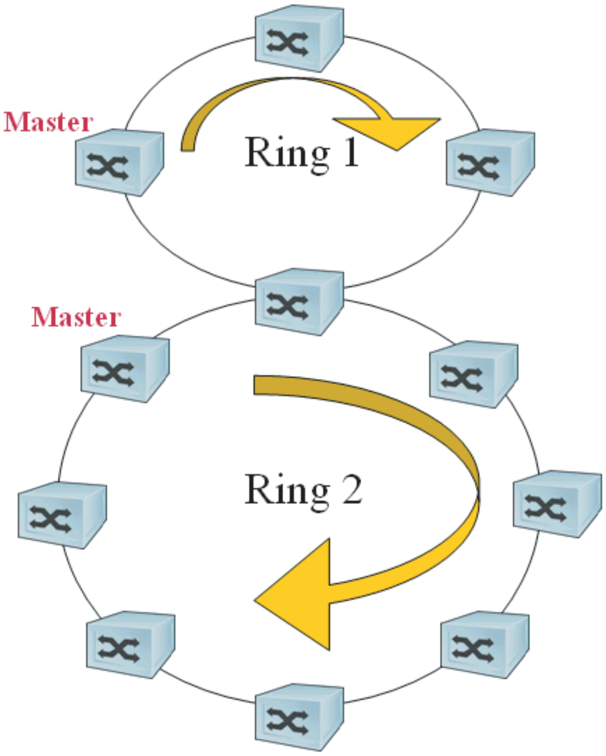

The FRRP basics are:

- **Ring ID**—Each ring has a unique 8-bit ring ID through which the ring is identified.

- **Control VLAN**—Each ring has a unique control VLAN through which tagged ring frames are sent. Control VLANs are used only for sending Ring Health Frames; control VLANs cannot be used for any other purpose.
- **Member VLAN**—Each ring maintains a list of member VLANs. Member VLANs must be consistent across the entire ring.
- **Port Role**—Each node has two interfaces for each ring: primary and secondary. The master node primary interface generates ring health frames (RHF). The master node secondary interface receives the frames. On transit nodes, there is no distinction between a primary and secondary interface.
- **Ring Interface State**—Each interface that is part of the ring maintains one of four states:

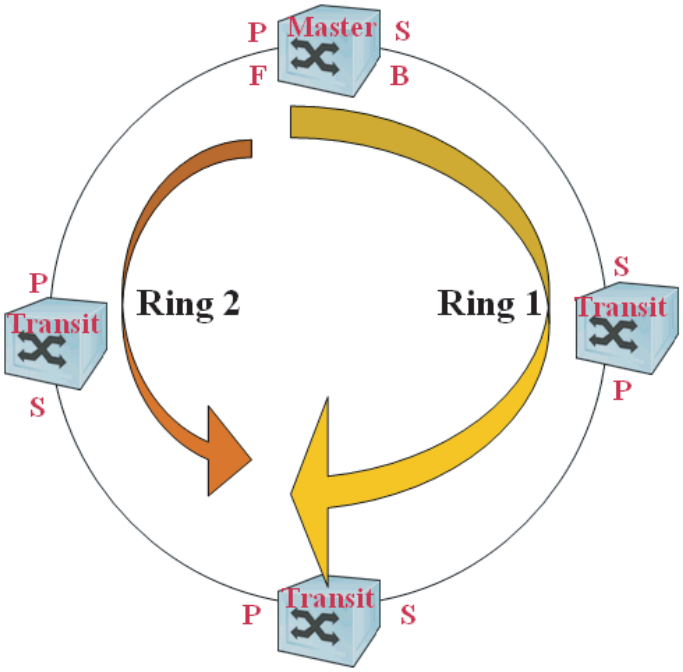

  Blocking State—Accepts ring protocol packets but blocks data plane packets. LLDP, FEFD, or other Layer 2 control packets are accepted. Only the master node secondary interface can enter this state.

  Pre-Forwarding State—A transition state (dead-interval time) before moving to the Forward state. Control plane traffic is forwarded but data plane traffic is blocked. The master node secondary interface transits through this state during ring bring-up. All interfaces transit through this state when an interface comes up after a flap.

  Forwarding State—Both ring protocol (control) and data plane traffic is passed. When the ring is operational, the primary interface on the master node and both primary and secondary interfaces on the transit nodes are in forwarding state. When the ring is broken, all ring interfaces are in this state.

  Disabled State—When the interface is disabled or down or is not on the VLAN, its state is disabled.

- **Ring Protocol Timers**

  Hello Interval—The interval at which ring frames are generated from the primary interface of the master node (default 500 ms). The hello interval is configurable in 50 ms increments from 50 ms to 2000 ms.

  Dead Interval—The default is 3X the hello interval rate. The dead interval is configurable in 50 ms increments from 50 ms to 6000 ms.

- **Ring Status**—During initialization/configuration, the default ring status is disabled. The primary and secondary interfaces, control-vlan, and master/transit mode information must be configured for the ring to be up.

  Ring-Up—Ring is up and operational

  Ring-Down—Ring is broken or not set up

- **Ring Health-check Frame** (**RHF**)—Two types of RHFs are generated by the master node primary interface. Allowing two different RHFs removes additional processing load for the ring from the transit nodes. Because they terminate at the master node, the RHFs never loop in the ring.

  Topology Change RHF (**TCRHF**)—This contains the ring status bit, keepalive bit, and the control and member VLAN hash and is processed at each node of the ring. Multiple TCRHFs are sent out, with the same sequence number, on any topology change to ensure all ring nodes receive it. There is no periodic transmission of TCRHFs. The TCRHFs are sent on triggered events of ring failure or ring restoration only.

  Hello RHF (**HRHF**)—These are processed only on the secondary interface of the master node of the ring. The transit nodes pass the HRHF through the hardware without processing it. A HRHF is sent at every hello interval.

Configuration Task List for FRRP

Important Points to Remember

- FRRP is media and speed independent.
- FRRP is a Force10 proprietary protocol that does not interoperate with any other vendor.
- Spanning Tree must be disabled on both primary and secondary interfaces before Resilient Ring protocol is enabled.
- The port must be a Layer 2 port.
- A VLAN configured as control VLAN for a ring cannot be configured as a control or member VLAN for any other ring.
- The control VLAN should not be used to carry any data traffic—The control VLAN should not have members that are not ring interfaces.
- If multiple rings share one or more member VLANs, they cannot share any links between them.
- Each ring can have only one master node; all others are transit nodes.

Figure 258 is a sample ring topology. The master node is R1 with R2 and R3 the transit nodes. This topology uses VLANs

Figure 258 Ring Topology Example

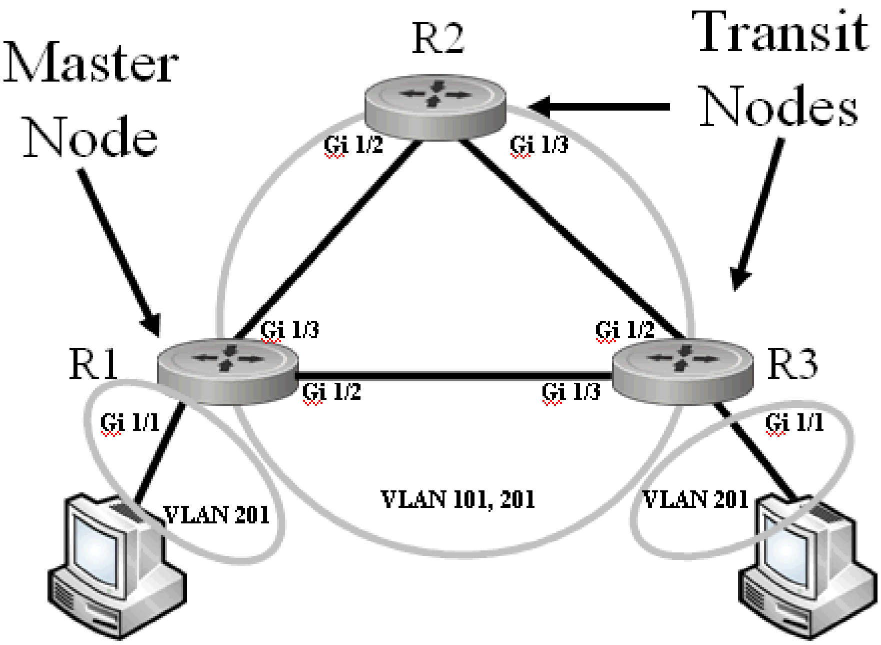

Prerequisites for FRRP Configuration

Although Force10 supports FRRP on a VLAN stack, it's not necessary that the ports or the VLAN must be VLAN-stack compatible. The ports can also be part of a normal VLAN.

configure the physical interface

`Force10#`**`no spanning-tree`**

`Force10#`**`vlan-stack trunk`** (on ports that are part of FRRP rings)

`Force10#`**`vlan-stack access`** (on ports that are connected to VLANs with STP BPDUs to be tunneled)

create traffic vlan, allow trunking, add interfaces

```
Force10#interface vlan 201
```

```
Force10(if)#vlan-stack compatible
```

```
Force10(if)#member gigabitethernet 1/1-3
```

create control vlan, allow trunking, add interfaces

```
Force10#interface vlan 101

Force10(if)#vlan-stack compatible

Force10(if)#member gigabitethernet 1/2-3
```

FRRP Configuration

create and configure the frrp group

Use the commands in the following sequence to create the FRRP group, select primary and secondary interfaces, specify the control VLAN, select the mode for this node, and enable FRRP, starting in CONFIGURATION mode:

Step	Task	Command
1	Enter the FRRP protocol and designate a ring ID	Force10(config)#**protocol frrp** *101*
2	Configure the primary, secondary, and control-vlan interfaces	Force10(config-frrp)#**interface primary gigabitethernet** *1/3* **secondary gigabitethernet** *1/3* **control-vlan** *101*
3	Designate the member VLAN	Force10(config-frrp)#**member-vlan** *201*
4	Set the node mode	Force10(config-frrp)#**mode [master | transit]**
5	Enable the FRRP	Force10(config-frrp)#**no disable**
6	Set the hello and dead interval for the ring control packets	Force10(config-frrp)#**timer hello-interval** *50* **dead-interval** *150*

```
Force10(conf)# protocol frrp 101
Force10(conf-frrp)# interface primary gigabitethernet 1/3 secondary gigabitethernet 1/3 control-vlan
101
Force10(conf-frrp)# member-vlan 201
Force10(conf-frrp)# mode master
Force10(conf-frrp)# no disable
Force10(conf-frrp)# timer hello-interval 50
Force10(conf-frrp)# timer dead-interval 150
```

Figure 259 Force10 resilient ring protocol (frrp) Command Example for Master Node

Configuring Ring Protocol for Transit Node

```
Force10(config)#protocol frrp 101
Force10(config-frrp)#interface primary port-channel 12 secondary port-channel 31 control-vlan 101
Force10(config-frrp)#member-vlan 201
Force10(config-frrp)#mode transit
Force10(config-frrp)#no disable
Force10(config-frrp)#timer dead-interval 150 ms
Force10(config-frrp)#timer hello-interval 50 ms
```

Figure 260 protocol frrp Command Example for Transit Node

Troubleshooting FRRP

Note: To clear the FRRP counters, use the **clear frrp** command.

show frrp command from a Master Node

```
Force10#show frrp 12
Ring protocol  12 is in Master  mode
Ring is UP
Ring Protocol Interface:
Primary  : TenGigabitEthernet 5/1               State: Forwarding
Secondary: TenGigabitEthernet 5/3               State: Blocking
Control Vlan: 998
Ring protocol Timers: Hello-Interval 500 msec          Dead-Interval 1500 msec
Ring Master's MAC Address is 00:01:e8:13:a4:ea
Topology Change Statistics:   Tx:8             Rx:0
Hello Statistics:   Tx:5             Rx:0
Number of state Changes: 3
Member Vlans:

Force10# show frrp summary
Ring ID      State   Mode   Control Vlan    Member Vlans
    12         UP   Master        12         2,3,4-5,2006
Force10#
```

Figure 261 show frrp Command Example from a Master Node

Show FRRP from a Transit Node

```
Force10#show frrp 12
Ring protocol  12 is in Transit mode
Ring Protocol Interface:
Primary  : TenGigabitEthernet 5/1               State: Forwarding
Secondary: TenGigabitEthernet 5/3               State: Forwarding
Control Vlan: 998
Ring protocol Timers: Hello-Interval 500 msec          Dead-Interval 1500 msec
Topology Change Statistics:   Tx:10            Rx:0
Hello Statistics:   Tx:85            Rx:0
Topology change discarded count: 0
Number of state Changes: 9
Member Vlans:
```

Figure 262 show frrp Command Example from a Transit Node

```
Force10#show running-config interface vlan 4000
!
interface Vlan 4000
  no ip address
  tagged TenGigabitEthernet 4/0-1
shutdown
Force10#show running-config interface tengigabitethernet 4/0
!
interface TenGigabitEthernet 4/0
  no ip address
  switchport
  no shutdown
Force10#show running-config interface tengigabitethernet 4/1
!
interface TenGigabitEthernet 4/1
  no ip address
  switchport
no shutdown


Force10# show frrp 1
Ring protocol   1 is in Master  mode
Ring is UP
Ring Protocol Interface:
        Primary  : TenGigabitEthernet 4/0               State: Forwarding
        Secondary: TenGigabitEthernet 4/1               State: Blocking
Control Vlan: 4000
Ring protocol Timers: Hello-Interval 500 msec          Dead-Interval 1500 msec
Ring Master's MAC Address is 00:01:e8:01:f9:b8
Topology Change Statistics:   Tx:10            Rx:0
Hello Statistics:   Tx:11            Rx:11
Number of state Changes: 3
   Member Vlans:
```

Figure 263 Sample Configuration for Control VLAN, and Primary and Secondary Ports for Normal VLAN

FRRP Checks

- Configuration Checks
  - Each ring must use a unique control VLAN
  - Only two interfaces can be members of a control VLAN
  - There can be only one “master” node.
  - Ring protocol configuration on Layer 2 interfaces only
- Ring consistency Auto-Check. Verify that control and member VLANs are consistent across all ring members.

Spanning Tree (if enabled globally) must be disabled on both primary and secondary interfaces when FRRP is enabled. If STP is enabled on the interface, then the **frrp** command is rejected. When the interface ceases to be a part of any ring, if Spanning Tree is enabled globally, it must be enabled explicitly for the interface.

Once a VLAN is configured as a control VLAN for any ring, it cannot be configured as a control or member VLAN for any other ring.

The maximum number of rings allowed on a chassis is 255.

All ring configurations are made from RING mode under CONFIGURATION mode.

Chapter 23 Power over Ethernet

C-Series	✓
E-Series	NO

Platform Specific Feature: Power over Ethernet is supported on C-Series only.

Power over Ethernet (PoE) is a function—described by IEEE 802.3af—that enables power to be transmitted to Ethernet devices over the signal pairs of an Unshielded Twisted Pair (UTP) cable. A maximum of 15.4 Watts can be transmitted over a link. PoE is useful in networks with IP phones and wireless access points because separate power supplies for powered devices (PD) are not needed.

- For a complete listing of commands related to Power over Ethernet, see the *FTOS Command Line Interface Reference*.

The C-Series is AC-only power sourcing equipment (PSE). The chassis transmits power to connected IEEE 802.3af-compliant devices via ports that are enabled with PoE. A minimum of four power supply units (PSU) are required to enable PoE, and 96 ports can be enabled per PSU thereafter, as described in Table 30. The 8th power supply is for PoE redundancy

Table 30 PoE Ports per Power Supply Unit

Power Supply Units	Max PoE Ports
1	—
2	—
3	—
4	77
5	192
6	288
7	384
8	PoE redundancy

Configuring Power over Ethernet

Configuring PoE is a two-step process:

1. Connect an IEEE 802.3af compliant powered device directly to a port.
2. Enable PoE on the port. See page 418.

Related Configuration Tasks

- Managing PoE Ports on page 420
- Monitoring the Power Budget on page 420
- Recovering from a Failed Power Supply on page 422

Enabling PoE on a Port

PoE is disabled by default. Enable PoE on a port using the **power inline** {**auto** [*max_milli-watts*] **| static** [*max milli-watts*] | **never**}command from INTERFACE mode.

- The **power inline auto** command allows the port to determine the amount of power that a Class 1-4 powered device requires, and supply it.
- The **power inline static** command guarantees the powered device the maximum amount of power (15.4 Watts).
- Limit the maximum amount of power (in milli-watts) available to a powered device using the command **power inline auto** *max_milli-watts* or **power inline static** *max_milli-watts*.

Note: Use the **power inline static** *max_milli-watts* command to avoid allocating more power than necessary to a port because allocated power is made unavailable to other ports regardless of whether it is consumed. Typical IP phones use 3-5 Watts.

- Disable PoE on a port using the **power inline never** command.

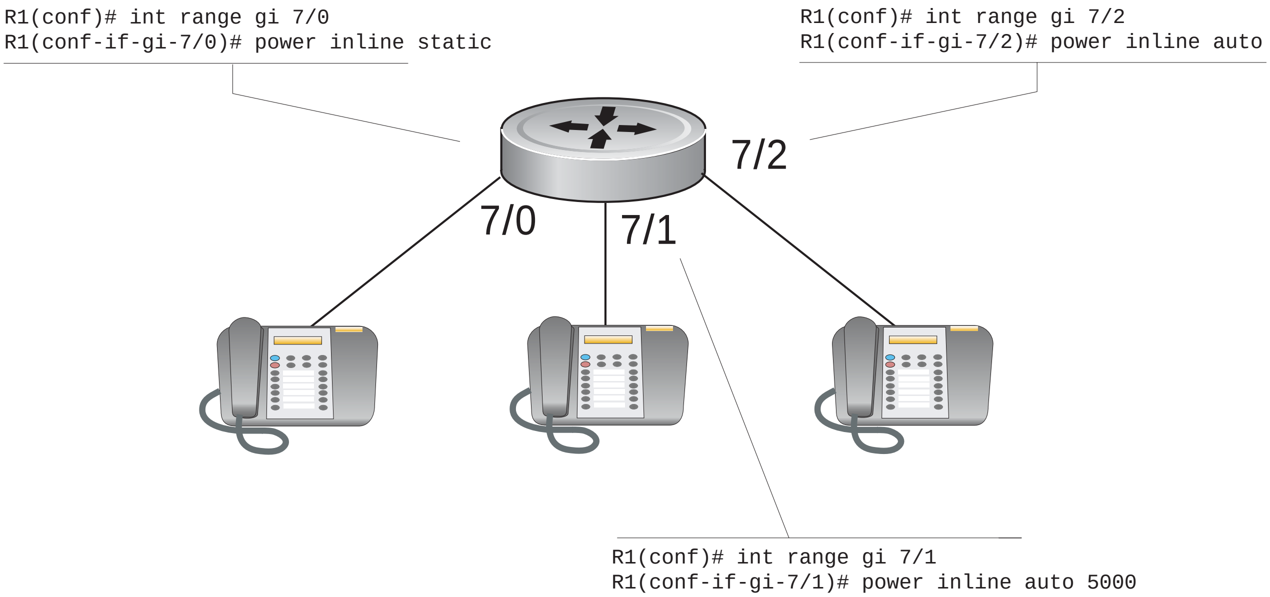

```
R1(conf)# int range gi 7/0
R1(conf-if-gi-7/0)# power inline static
```

```
R1(conf)# int range gi 7/2
R1(conf-if-gi-7/2)# power inline auto
```

```
R1(conf)# int range gi 7/1
R1(conf-if-gi-7/1)# power inline auto 5000
```

Figure 264 Enabling PoE

View the amount of power that a port is consuming using that the **show power inline** command from EXEC privilege mode.

```
R1#show power inline
Interface    Admin     Oper    Inline Power      Inline Power     Class
                                Allocated         Consumed
                                  (Watts)          (Watts)
--------    -------   -----   ------------     --------------    ----
Gi 7/0      static     on        15.40              3.36           2
Gi 7/1      auto       on        5.00               3.36           2
Gi 7/2      auto       on        15.40              3.42           2
```

Figure 265 show power inline Command Example

Table 31 describes the fields the that **show power inline** command displays.

Table 31 show power inline Field Description

Field	Port Number
Interface	Displays all PoE-enabled ports.
Admin	Displays the administrative mode of the inteface: • *auto* indicates that power is supplied according to the requirements of the powered device. • *static* indicates that the maximum configured amount of power is supplied to the powered device.
Oper	Displays the status of the powered device: on or off.
Inline Power Allocated	Displays the amount of power allocated to a port.
Inline Power Consumed	Displays the amount of power that a powered device is consuming.
Class	Displays the type of powered device: Class 0, Class 1, Class 2, Class 3, or Class 4.

View the total power consumption of the chassis using the **show power detail** command from EXEC privilege mode.

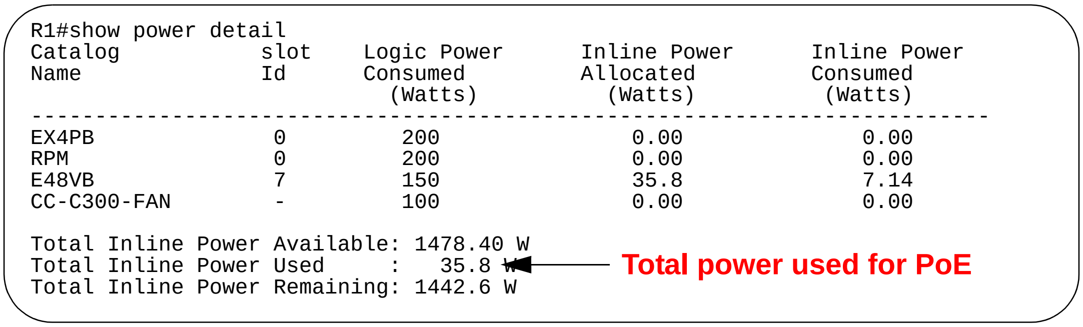

```
R1#show power detail
Catalog             slot    Logic Power      Inline Power      Inline Power
Name                Id      Consumed         Allocated         Consumed
                              (Watts)          (Watts)           (Watts)
---------------------------------------------------------------------------
EX4PB               0          200              0.00              0.00
RPM                 0          200              0.00              0.00
E48VB               7          150              35.8              7.14
CC-C300-FAN         -          100              0.00              0.00

Total Inline Power Available: 1478.40 W
Total Inline Power Used     :   35.8 W
Total Inline Power Remaining: 1442.6 W
```

Total power used for PoE

Figure 266 show power detail Command Example

Table 31 describes the fields that the **show power detail** command displays.

Table 32 show power detail Field Description

Field	Port Number
Catalogue Name	Displays the Force10 catalogue number of the line card, RPM, and fan tray.
Slot ID	Displays the slot number in which the component in installed.
Logic Power Consumed	Displays the total amount of power that the chassis component is consuming for basic functionality.
Inline Power Allocated	Displays the amount of power allocated to a port.
Inline Power Consumed	Displays the amount of power that a powered device is consuming.
Class	Displays the type of powered device: Class 0, Class 1, Class 2, Class 3, or Class 4.

Managing PoE Ports

Monitoring the Power Budget

Every time an interface is enabled in either **auto**, or **static** mode, the configured maximum power (15.4 Watts by default) is allocated to that interface from the power budget.

- The power budget is the amount of power available from the installed PSUs minus the power required to operated the chassis.
- Use the **show power detail** command from EXEC privilege mode to help you determine if power is available for additional PoE ports (1478.40 Watts are supplied per PSU).

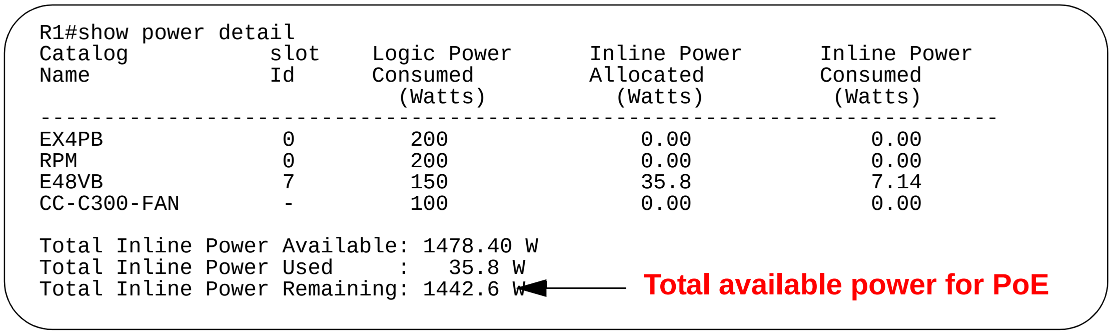

```
R1#show power detail
Catalog             slot     Logic Power       Inline Power       Inline Power
Name                Id       Consumed          Allocated          Consumed
                               (Watts)           (Watts)            (Watts)
----------------------------------------------------------------------------
EX4PB               0          200               0.00               0.00
RPM                 0          200               0.00               0.00
E48VB               7          150               35.8               7.14
CC-C300-FAN         -          100               0.00               0.00

Total Inline Power Available: 1478.40 W
Total Inline Power Used     :   35.8 W
Total Inline Power Remaining: 1442.6 W
```

Total available power for PoE

Figure 267 show power detail Command Example

Enabling PoE on more ports than is supported by power budget produces one of these results:

1. If the newly PoE-enabled port has a lower priority, then the CLI is accepted, but power is not allocated to the port. In this case, Message 7 is displayed.

Message 7 Insufficient Power to Enable PoE

```
%Warning: Insufficient power to enable.POE oper-status set to OFF for port <linecard/
portnumber>
```

2. If the newly PoE-enabled port has a higher priority, then the CLI is accepted, and power is terminated on the lowest priority port in the chassis. If another power supply is added to the system at a later point in time, both ports receive power.
   - If all of the lower priority ports combined cannot meet the power requirements of the newly enabled port, then the CLI is accepted, but power on the lower priority ports is not terminated, and no power is supplied to the port.

The second result in this scenario is true even if a powered device is not connected to the port. Power can be allocated to a port, thus subtracting it from the power budget and making it unavailable to other ports, but that power does not have to be consumed.

Port Priorities

PoE-enabled ports have different priorities based on their configuration, line card number, and port number as described by Table 33.

Table 33 PoE Ports Priorities

Configuration	Port Number	Priority
Ports configured with **power inline static**	Ports with the lowest port numbers in linecards with the lowest slot number	1
	Ports with the lowest port numbers	2

Table 33 PoE Ports Priorities

Configuration	Port Number	Priority
Ports configured with **power inline auto**	Ports with the lowest port numbers in linecards with the lowest slot number	3
	Ports with the lowest port numbers	4

Recovering from a Failed Power Supply

A minimum of four PSUs are required to enable PoE. Three are dedicated to powering the chassis, one of which is redundant, and any remaining PSUs can be allocated to PoE. If ports are PoE-enabled, and a PSU fails, power might be terminated on some ports to compensate for the power loss. This does not affect PoE individual port configurations.

If power must be terminated for some ports, the order in which ports are affected is based on priority. Ports with the lowest priority are terminated first (see Port Priorities).

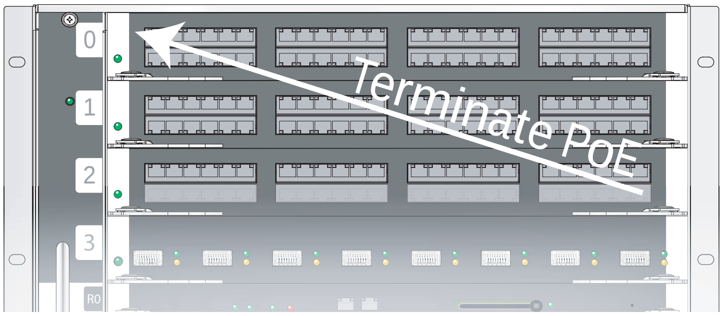

Figure 268 Order of PoE Termination

For the configuration in Figure 264:

- Power for ports 7/1 and 7/2 is terminated first because it is configured with **inline power auto**.
- Power for port 7/1 is terminated before PoE for port 7/2 because port 7/1 has a lower port number.
- Power for port 7/0 is terminated last because it is configured with **inline power static**.

When a failed PSU is replaced and there is sufficient power for PoE, power is automatically re-supplied for previously configured PoE ports, and power is supplied first to ports with the highest priority.

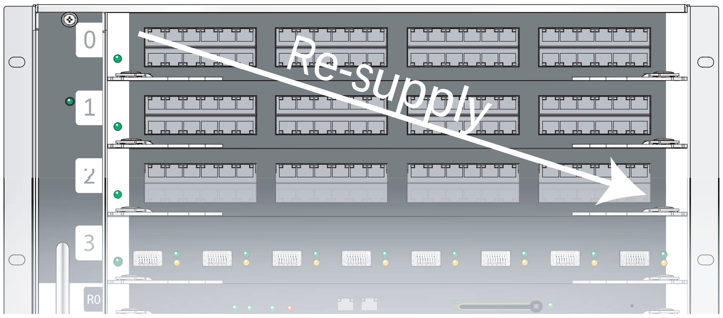

Figure 269 Order of PoE Re-Supply

Chapter 24 Spanning Tree Protocol

Protocol Overview

Spanning Tree Protocol (STP) is a Layer 2 protocol—described by IEEE 802.1d—that eliminates loops in a bridged topology by enabling a single path through the network. By eliminating loops, the protocol improves scalability in a large network and enables you to implement redundant paths, which can be activated upon the failure of active paths.

Implementation Information

C-Series	NO
E-Series	✓

Platform Specific Feature: Layer 2 BPDU Tunneling is supported on E-Series only.

E-Series employs Layer 2 BPDU filtering to reduce flapping in Spanning Tree configurations and add stability to Layer 2 networks. If Spanning Tree is enabled on a remote interface and disabled on the Force10 local interface, the Force10 system drops the BPDUs instead of forwarding them to the RPM CPU for processing.

Layer 2 BPDU filtering is available for STP, RSTP, and MSTP. It is enabled by default and is non-configurable.

Configuring Spanning Tree

Implementing Spanning Tree is a two-step process:

1. Configure interfaces for Layer 2. See page 424.
2. Enable Spanning Tree Protocol. See page 425.

Related Configuration Tasks

- Adding an Interface to the STG on page 428
- Removing an Interface from the STG on page 428
- Modifying Global Parameters on page 429

- Modifying Interface Parameters on page 429
- Enabling PortFast on page 430
- Influencing STP Root Selection on page 431

Important Points to Remember

- Spanning Tree Protocol is disabled by default.
- FTOS supports only one Spanning Tree instance (0).
- All ports in VLANs and all enabled interfaces in Layer 2 mode are automatically added to the Spanning Tree topology.
- Any Layer 2 interfaces that are to be added to the Spanning Tree topology after STP has been enabled globally must be added manually.
- The IEEE Standard 802.1D allows eight bits for port ID and eight bits for priority. However, the eight bits for port ID provide port IDs for only 256 ports and the C-Series can contain 336 ports. To accommodate the increased number of ports, FTOS uses four bits from priority field in the port ID field.This implementation affects the Bridge MIB (RFC 1493), and you must interpret objects such as *dot1dStpPortDesignatedPort* object by using the first four bits as the priority and the last 12 bits as the port ID.

Table 34 displays the default values for Spanning Tree.

Table 34 STP Default Values

STP Parameter		Default Value
Forward Delay		15 seconds
Hello Time		2 seconds
Max Age		20 seconds
Port Cost	100-Mb/s Ethernet interfaces	19
	1-Gigabit Ethernet interfaces	4
	10-Gigabit Ethernet interfaces	2
	Port Channel with 100 Mb/s Ethernet interfaces	18
	Port Channel with 1-Gigabit Ethernet interfaces	3
	Port Channel with 10-Gigabit Ethernet interfaces	1
Port Priority		8

Configuring Interfaces for Layer 2 Mode

All interfaces on all bridges that will participate in Spanning Tree must be in Layer 2 and enabled.

```
R1(conf)# int range gi 1/1 - 4
R1(conf-if-gi-1/1-4)# switchport
R1(conf-if-gi-1/1-4)# no shutdown
R1(conf-if-gi-1/1-4)#show config
!
interface GigabitEthernet 1/1
 no ip address
 switchport
 no shutdown
!
interface GigabitEthernet 1/2
 no ip address
 switchport
 no shutdown
!
interface GigabitEthernet 1/3
 no ip address
 switchport
 no shutdown
!
interface GigabitEthernet 1/4
 no ip address
 switchport
 no shutdown
```

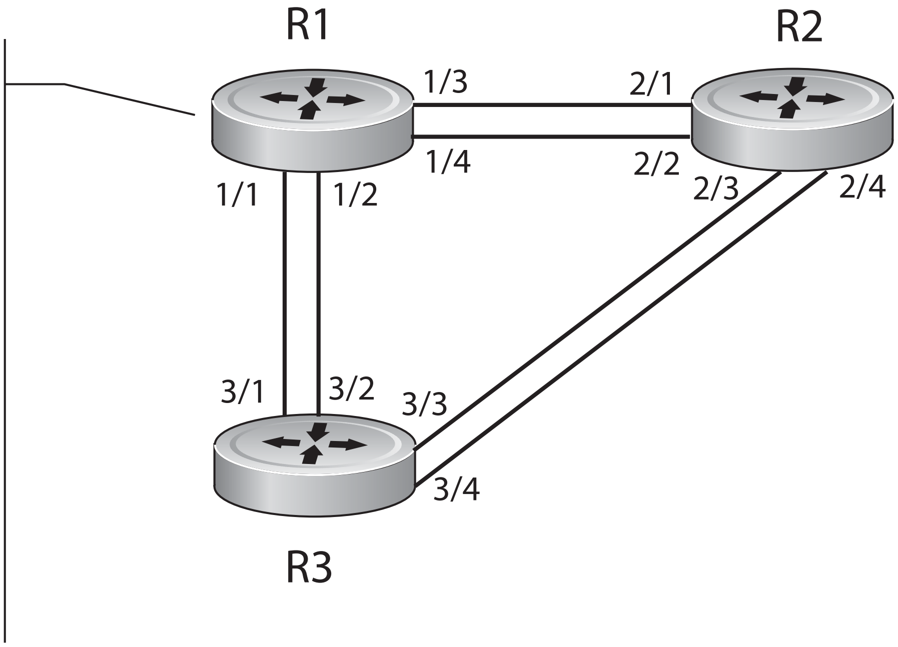

Figure 270 Example of Configuring Interfaces for Layer 2 Mode

To configure the interfaces for Layer 2 and then enable them:

Step	Task	Command Syntax	Command Mode
1	If the interface has been assigned an IP address, remove it.	**no ip address**	INTERFACE
2	Place the interface in Layer 2 mode.	**switchport**	INTERFACE
3	Enable the interface.	**no shutdown**	INTERFACE

Verify that an interface is in Layer 2 mode and enabled using the **show config** command from INTERFACE mode.

```
R1(conf-if-gi-1/1)#show config
!
interface GigabitEthernet 1/1
 no ip address
 switchport
no shutdown
R1(conf-if-gi-1/1)#
```

Indicates that the interface is in Layer 2 mode

Figure 271 Verifying Layer 2 Configuration

Enabling Spanning Tree Protocol Globally

Spanning Tree Protocol must be enabled globally; it is not enabled by default.

To enable Spanning Tree globally for all Layer 2 interfaces:

Step	Task	Command Syntax	Command Mode
1	Enter the PROTOCOL SPANNING TREE mode.	**protocol spanning-tree 0**	CONFIGURATION
2	Enable Spanning Tree.	**no disable**	PROTOCOL SPANNING TREE

Note: To disable STP globally for all Layer 2 interfaces, enter the **disable** command from PROTOCOL SPANNING TREE mode.

Verify that Spanning Tree is enabled using the **show config** command from PROTOCOL SPANNING TREE mode.

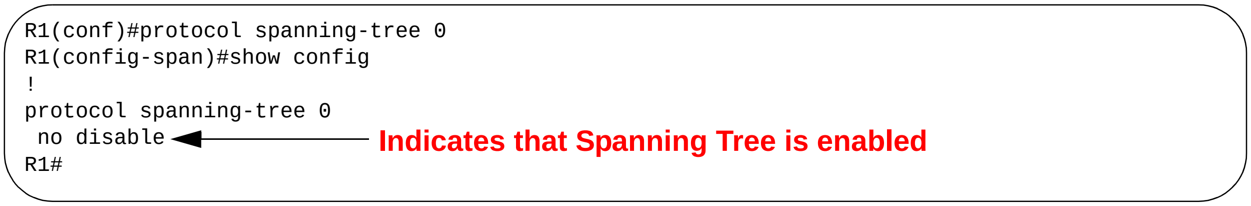

```
R1(conf)#protocol spanning-tree 0
R1(config-span)#show config
!
protocol spanning-tree 0
 no disable
R1#
```

Figure 272 Verifying STP is Enabled

When you enable Spanning Tree, all physical and port-channel interfaces that are enabled and in Layer 2 mode are automatically part of the Spanning Tree topology.

- Only one path from any bridge to any other bridge participating in STP is enabled.
- Bridges block a redundant path by disabling one of the link ports.

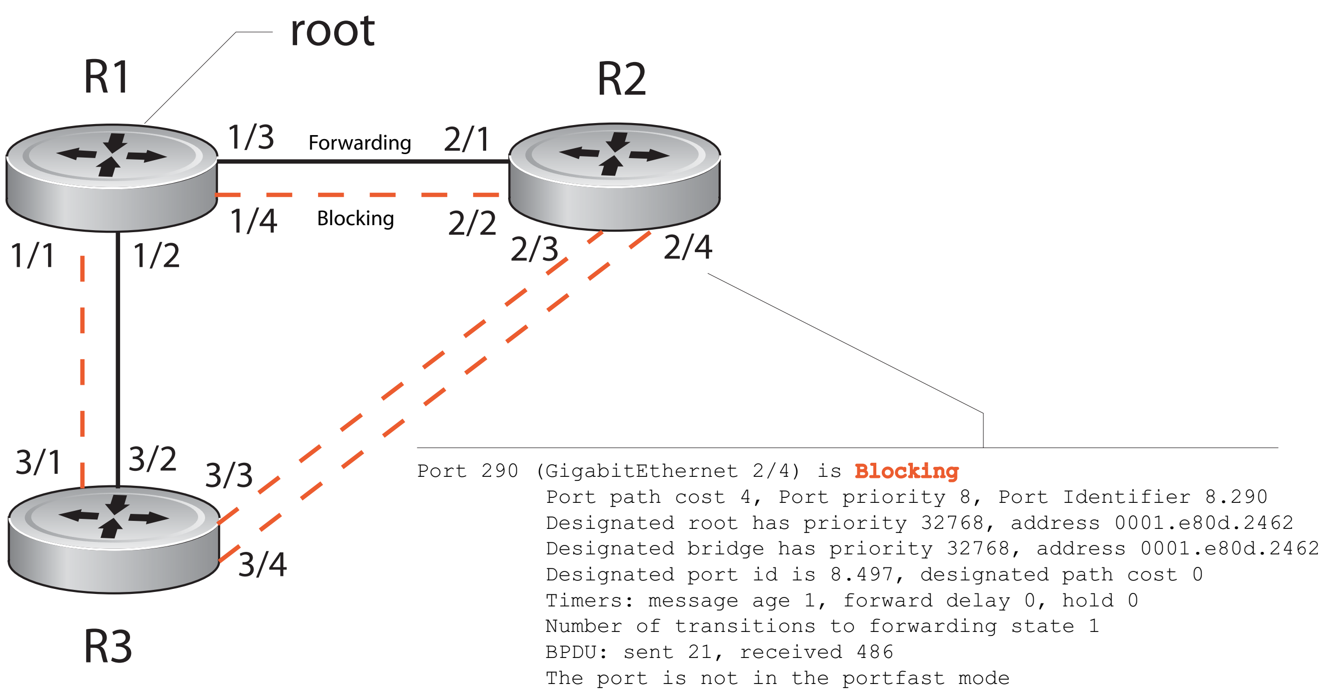

```
Port 290 (GigabitEthernet 2/4) is Blocking
          Port path cost 4, Port priority 8, Port Identifier 8.290
          Designated root has priority 32768, address 0001.e80d.2462
          Designated bridge has priority 32768, address 0001.e80d.2462
          Designated port id is 8.497, designated path cost 0
          Timers: message age 1, forward delay 0, hold 0
          Number of transitions to forwarding state 1
          BPDU: sent 21, received 486
          The port is not in the portfast mode
```

Figure 273 Spanning Tree Enabled Globally

View the Spanning Tree configuration and the interfaces that are participating in STP using the **show spanning-tree 0** command from EXEC privilege mode. If a physical interface is part of a port channel, only the port channel is listed in the command output.

```
R2#show spanning-tree 0
     Executing IEEE compatible Spanning Tree Protocol
          Bridge Identifier has priority 32768, address 0001.e826.ddb7
          Configured hello time 2, max age 20, forward delay 15
          Current root has priority 32768, address 0001.e80d.2462
          Root Port is 289 (GigabitEthernet 2/1), cost of root path is 4
          Topology change flag not set, detected flag not set
          Number of topology changes 3 last change occurred 0:16:11 ago
                  from GigabitEthernet 2/3
          Timers: hold 1, topology change 35
                  hello 2, max age 20, forward delay 15
          Times:  hello 0, topology change 0, notification 0, aging Normal

     Port 289 (GigabitEthernet 2/1) is Forwarding
          Port path cost 4, Port priority 8, Port Identifier 8.289
          Designated root has priority 32768, address 0001.e80d.2462
          Designated bridge has priority 32768, address 0001.e80d.2462
          Designated port id is 8.496, designated path cost 0
          Timers: message age 1, forward delay 0, hold 0
          Number of transitions to forwarding state 1
          BPDU: sent 21, received 486
          The port is not in the portfast mode

     Port 290 (GigabitEthernet 2/2) is Blocking
          Port path cost 4, Port priority 8, Port Identifier 8.290
--More--
```

Figure 274 show spanning-tree 0 Command Example

Confirm that a port is participating in Spanning Tree using the **show spanning-tree 0 brief** command from EXEC privilege mode.

```
R1#show spanning-tree 0 brief
     Executing IEEE compatible Spanning Tree Protocol
          Root ID  Priority 32768, Address 0001.e80d.2462
          We are the root of the spanning tree
          Root Bridge hello time 2, max age 20, forward delay 15
          Bridge ID  Priority 32768, Address 0001.e80d.2462
          Configured hello time 2, max age 20, forward delay 15
Interface                                    Designated
 Name          PortID Prio Cost Sts Cost     Bridge ID            PortID
-------------- ------ ---- ---- --- -----    -----------------    ------
Gi 1/1          8.496    8    4 DIS     0    32768 0001.e80d.2462  8.496
Gi 1/2          8.497    8    4 DIS     0    32768 0001.e80d.2462  8.497
Gi 1/3          8.513    8    4 FWD     0    32768 0001.e80d.2462  8.513
Gi 1/4          8.514    8    4 FWD     0    32768 0001.e80d.2462  8.514
R1#
```

Figure 275 show spanning-tree brief Command Example

Adding an Interface to the STG

To add a Layer 2 interface to the Spanning Tree topology:

Task	Command Syntax	Command Mode
Remove a Layer 2 interface from the Spanning Tree topology.	**spanning-tree 0**	INTERFACE

Removing an Interface from the STG

To remove a Layer 2 interface from the Spanning Tree group:

Task	Command Syntax	Command Mode
Remove a Layer 2 interface to the Spanning Tree group.	**no spanning-tree 0**	INTERFACE

Modifying Global Parameters

You can modify Spanning Tree parameters. The root bridge sets the values for forward-delay, hello-time, and max-age and overwrites the values set on other bridges participating in Spanning Tree.

Note: Force10 Networks recommends that only experienced network administrators change the Spanning Tree parameters. Poorly planned modification of the Spanning Tree parameters can negatively impact network performance.

To change these parameters:

Task	Command Syntax	Command Mode
Change the forward-delay parameter (the wait time before the interface enters the *forwarding* state). • Range: 4 to 30 • Default: 15 seconds	**forward-delay** *seconds*	PROTOCOL SPANNING TREE
Change the hello-time parameter (the BPDU transmission interval). **Note:** With large configurations (especially those with more ports) Force10 Networks recommends that you increase the hello-time. Range: 1 to 10 Default: 2 seconds	**hello-time** *seconds*	PROTOCOL SPANNING TREE
Change the max-age parameter (the refresh interval for configuration information that is generated by recomputing the Spanning Tree topology). Range: 6 to 40 Default: 20 seconds	**max-age** *seconds*	PROTOCOL SPANNING TREE

View the current values for global parameters using the **show spanning-tree 0** command from EXEC privilege mode. See Figure 274.

Modifying Interface Parameters

You can set the port cost and port priority values of interfaces in Layer 2 mode.

- **Port cost** is a value that is based on the interface type. The greater the port cost, the less likely the port will be selected to be a forwarding port.
- **Port priority** influences the likelyhood that a port will be selected to be a forwarding port in case that several ports have the same port cost.

The default values are listed in Table 34.

To change the port cost or priority of an interface:

Task	Command Syntax	Command Mode
Change the port cost of an interface. Range: 0 to 65535 Default: see Table 34.	**spanning-tree 0 cost** *cost*	INTERFACE
Change the port priority of an interface. Range: 0 to 15 Default: 8	**spanning-tree 0 priority** *priority-value*	INTERFACE

View the current values for interface parameters using the **show spanning-tree 0** command from EXEC privilege mode. See Figure 274.

Enabling PortFast

The PortFast feature enables interfaces to begin forwarding traffic approximately 30 seconds sooner. Interfaces forward frames by default until they receive a BPDU that indicates that they should behave otherwise; they do not go through the Learning and Listening states.

Caution: Enable PortFast only on links connecting to an end station. PortFast can cause loops if it is enabled on an interface connected to a network.

To enable PortFast on an interface:

Task	Command Syntax	Command Mode
Enable PortFast on an interface.	**spanning-tree** *stp-id* **portfast**	INTERFACE

Verify that PortFast is enabled on a port using the **show spanning-tree** command from the EXEC privilege mode or the **show config** command from INTERFACE mode; Force10 recommends using the **show config** command, as shown in Figure 276.

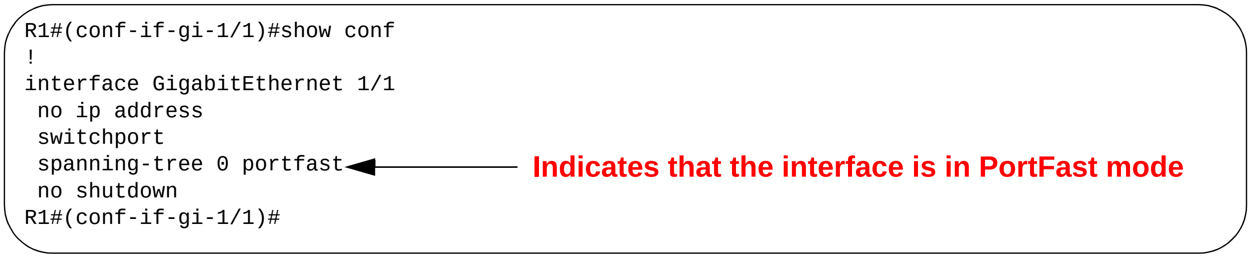

```
R1#(conf-if-gi-1/1)#show conf
!
interface GigabitEthernet 1/1
 no ip address
 switchport
 spanning-tree 0 portfast
 no shutdown
R1#(conf-if-gi-1/1)#
```

Indicates that the interface is in PortFast mode

Figure 276 PortFast Enabled on Interface

Influencing STP Root Selection

The Spanning Tree Protocol determines the root bridge, but you can assign one bridge a lower priority to increase the likelihood that it will be selected as the root bridge. You can also specify that a bridge is the root or the secondary root.

To change the bridge priority or specify that a bridge is the root or secondary root, use the following command:

Task	Command Syntax	Command Mode
Assign a number as the bridge priority or designate it as the root or secondary root. *priority-value* range: 0 to 65535. The lower the number assigned, the more likely this bridge will become the root bridge. The default is 32768. • The primary option specifies a bridge priority of 8192. • The secondary option specifies a bridge priority of 16384.	**bridge-priority** {*priority-value* | **primary** | **secondary**}	PROTOCOL SPANNING TREE

View only the root information using the **show spanning-tree root** command (see Figure 277) from EXEC privilege mode.

```
R1#show spanning-tree 0 root
          Root ID  Priority 32768, Address 0001.e80d.2462
          We are the root of the spanning tree
          Root Bridge hello time 2, max age 20, forward delay 15
R1#
```

Figure 277 show spanning-tree root Command Example

Chapter 25 PVST+

PVST+ (Per-VLAN Spanning Tree plus) Protocol eliminates loops in a bridged topology by designating a single path through the network per VLAN. By eliminating loops, the protocol improves scalability in a large network and provisions redundant paths which can be activated upon the failure of active paths. PVST+ can only be used on Ethernet port-based VLANs: Ethernet, Fast Ethernet, Gigabit Ethernet, and 10Gigabit Ethernet.

PVST+ enables each VLAN in a switch to be loop free. The root switch sends PVST+ information to other switches in the network to maintain the network topology using all links to route traffic a single Spanning Tree Protocol (STP) topology for each VLAN

With PVST+, the user may define a root bridge and an STP topology for each VLAN that allows the switches to use network links instead of creating bridge loops. A loop-free path is selected for each VLAN based on the root and port cost. To prevent loops in the network, PVST+ ports are placed in one of the following states: learning, forwarding, blocking, or disabled.

The following sections describe PVST+ in FTOS:

For more information on PVST+, see the IEEE Standard 802.1D Media Access Control Bridges.

PVST+ Behaviors

EdgePort

FTOS uses PortFast to speed up the connectivity between end stations. When STP discards all user data before it puts a port into forwarding state, it may cause a delay. To eliminate this delay, PortFast, eliminates the STP topology change between learning and forwarding states before moving a port from blocking state to forwarding state.

Configuration Task List for PVST+

The following list includes the configuration tasks for PVST+:

- Enabling PVST+ on page 434
- Configuring Bridge Priority on page 435
- Configuring Forward-delay on page 435
- Configuring Hello-time on page 436
- Configuring Max-age on page 436
- Configuring Port Cost on page 437
- Configuring Port Priority on page 438
- Configuring Interface as Edge Port on page 439

For a complete listing of all commands related to PVST+, see .

Enabling PVST+

By default, PVST+ is not enabled in FTOS. To enable PVST+ globally in FTOS, use these commands in the following sequence, starting in the PVST+ CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**protocol spanning-tree pvst**	CONFIGURATION	Enter the PVST+ mode.
2	**no disable**	CONFIGURATION (conf-pvst)	Enable PVST+. Once PVST+ is enabled, the device runs an STP instance for each VLAN it supports.

You must first enter PVST+ configuration mode before enabling (or disabling) PVST+ on the device.

```
Force10#conf
Force10(conf)#protocol spanning-tree pvst
Force10(conf-pvst)#no disable
```

Figure 278 Enabling PVST+

When PVST+ is enabled, the Force10 device runs an STP instance for each VLAN it supports.

Configuring Bridge Priority

In STP, the algorithm determines the root bridge, but you can assign one bridge a lower priority to increase the likelihood that it will be selected as the STP root bridge. During the STP initialization process, the bridge with the lowest number identifier is elected to be the root; however, you can influence the root selection by designating a bridge as a primary or backup root.

To change the bridge priority, use the following command in the PVST+ CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
vlan *vlan-range* **bridge-priority** *value*	conf-pvst	Enter the keyword **bridge-priority** followed by the bridge priority value in increments of 4096. Range: 0 to 61440 Default: 32768

To configure bridge priority, use the command as shown in the example below:

```
Force10(conf)# protocol spanning-tree pvst
Force10(conf-pvst)# no disable

Context: protocol spanning-tree pvst

Force10(conf-pvst)# vlan 1 bridge-priority 4096
```

Figure 279 PVST+ Configuration Examples

Configuring Forward-delay

Forward-delay is the amount of time an interface waits in the Blocking and Learning States before it transitions to Forwarding State.

To change the forward delay value in seconds, use the following command in the PVST+ CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
vlan *vlan-range* **forward-delay** *seconds*	conf-pvst	Changes the time interval before FTOS transitions to the forwarding state. Enter the keyword **forward-delay** followed by the time interval, in seconds, that FTOS waits before transitioning PVST+ to the forwarding state. Range: 4 to 30 seconds Default: 15 seconds

To configure forward-delay, use the command as shown in the example below:

```
Force10(conf)# protocol spanning-tree pvst
Force10(conf-pvst)# no disable

Context: protocol spanning-tree pvst

Force10(conf-pvst)# vlan 1 forward-delay 4
```

Figure 280 PVST+ Configuration Examples

Configuring Hello-time

Hello-time is the time interval between the generation of PVST+ Bridge Protocol Data Units (BPDUs). To change the hello time value in seconds, use the following command in the PVST+ CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
vlan *vlan-range* **hello-time** *seconds*	conf-pvst	Changes the time interval between BPDUs. Enter the keyword **hello-time** followed by the time interval, in seconds, between transmission of BPDUs. Range: 1 to 10 seconds Default: 2 seconds

To configure hello-time, use the commands as shown in the example below:

```
Force10(conf)# protocol spanning-tree pvst
Force10(conf-pvst)# no disable

Context: protocol spanning-tree pvst

Force10(conf-pvst)# vlan 1 hello-time 6
```

Figure 281 PVST+ Configuration Examples

Configuring Max-age

Max-age is the length of time the PVST+ bridge maintains configuration information before it refreshes that information.

To change the max-age value in seconds, use the follwoing command in the PVST+ CONFIGURATION mode:

Command Syntax	Command Mode	Usage
vlan *vlan-range* **max-age** *seconds*	conf-pvst	Changes the time interval before PVST+ refreshes. Enter the keyword **max-age** followed by the time interval, in seconds, that FTOS waits before refreshing configuration information. Range: 6 to 40 seconds Default: 20 seconds

To configure max-age, use the commands as shown in the example below:

```
Force10(conf)# protocol spanning-tree pvst
Force10(conf-pvst)# no disable

Context: protocol spanning-tree pvst

Force10(conf-pvst)# vlan 1 max-age 10
```

Figure 282 PVST+ Configuration Examples

Configuring Port Cost

Force10 PVST+ implementation uses IEEE MST costs as the default costs. Please be sure to use the appropriate costs in a multi-vendor network as some implementations use IEEE STP costs as the default costs. The default PVST+ costs are listed below.

Table 35 E-Series Port Cost Values

FTOS Default Port Cost Values (IEEE MST)	Values
Range	1 to 200000
100 Mb/s Ethernet interface	200000
1-Gigabit Ethernet interface	20000
10-Gigabit Ethernet interface	2000
Port Channel interface with one 100 Mb/s Ethernet	200000
Port Channel interface with one 1-Gigabit Ethernet	20000
Port Channel interface with one 10-Gigabit Ethernet	2000
Port Channel with two 1-Gigabit Ethernet	18000
Port Channel with two 10-Gigabit Ethernet	1800
Port Channel with two 100-Mbps Ethernet	180000

The following command configures the port cost of the interface:

Command Syntax	Command Mode	Usage
spanning-tree pvst vlan *vlan-range* **cost** *number*	INTERFACE	(OPTIONAL) Enter the keyword **cost** followed by the port cost value. Range: 1 to 200000 Defaults: 100 Mb/s Ethernet interface = 200000 1-Gigabit Ethernet interface = 20000 10-Gigabit Ethernet interface = 2000 Port Channel interface with 100 Mb/s Ethernet = 200000 Port Channel interface with 1-Gigabit Ethernet = 20000 Port Channel interface with 10-Gigabit Ethernet = 2000 Port Channel with 2 1-Gigabit Ethernet = 18000 Port Channel with 2 10-Gigabit Ethernet = 1800 Port Channel with 2 100-Mbps Ethernet = 180000

To change the port cost of an interface, use the following command in the INTERFACE mode, as shown below:

```
Force10#conf
Force10(conf)#interface gigabitethernet 1/1
Force10(conf-if-gi-1/1)#spanning-tree pvst vlan 3 cost 18000
```

Figure 283 Configuring PVST+ Port Cost

Configuring Port Priority

Port priority determines the likelihood that the port will be selected to transmit traffic.

The following command configures the priority of the interface:

Command Syntax	Command Mode	Usage
spanning-tree pvst vlan *vlan-range* **priority** *value*	conf-pvst	(OPTIONAL) Enter the keyword **priority** followed the Port priority value in increments of 16. Range: 0 to 240 Default: 128

To change the port priority of an interface, use the following command in the INTERFACE mode, as shown below:

```
Force10#conf
Force10(conf)#interface gigabitethernet 1/1
Force10(conf-if-gi-1/1)#spanning-tree pvst vlan 3 priority 1
```

Figure 284 Configuring PVST+ Port Priority

Configuring Interface as Edge Port

The following command configures the priority of the interface:

Command Syntax	Command Mode	Usage
spanning-tree rstp edge-port	INTERFACE	(OPTIONAL) Enter the keyword **edge-port** to configure the interface as a Rapid Spanning Tree edge port.

To configure a particular interface as an edge port, issue the following command after entering INTERFACE mode:

```
Force10#conf
Force10(conf)#interface gigabitethernet 1/1
Force10(conf-if-gi-1/1)#spanning-tree pvst edge-port
```

Figure 285 Configuring an Interface as an Edge Port

Disabling PVST+

To disable PVST+ globally in FTOS, simply issue the **disable** command in PVST+ CONFIGURATION mode. When PVST+ is disabled, the **show spanning-tree pvst** command does not return any output.

```
Force10#conf
Force10(conf)#protocol spanning-tree pvst
Force10(conf-pvst)#disable
```

Figure 286 Disabling PVST+ Globally

To disable PVST+ on a particular interface, use the **no spanning-tree pvst** [**edge-port** | **vlan** *vlan-range* {**cost** *number* | **priority** *value*}] command after entering INTERFACE mode:

```
Force10#conf
Force10(conf)#interface gigabitethernet 1/1
Force10(conf-if-gi-1/1)#no spanning-tree
```

Figure 287 Disabling PVST+ on an Interface

The following command disables the port cost or priority of the interface:

Command Syntax	Command Mode	Usage
no spanning-tree pvst [**edge-port** | **vlan** *vlan-range* {**cost** *number* | **priority** *value*}]	INTERFACE	To disable PVST+ Edge port, VLAN, Port priority, and Port cost on an interface.

To disable PVST+ port cost or priority on a particular interface, use the **no spanning-tree pvst** [**vlan** *vlan-range* {**cost** *number*} command after entering INTERFACE mode:

```
Force10#conf
Force10(conf)#interface gigabitethernet 1/1
Force10(conf-if-gi-1/1)#no spanning-tree pvst vlan 3 cost 20000
```

Figure 288 Disabling PVST+ Port Cost on an Interface

To disable PVST+ priority on a particular interface, use the **no spanning-tree pvst** [**edge-port** | **vlan** *vlan-range* {**riority** *value*}] command after entering INTERFACE mode:

```
Force10#conf
Force10(conf)#interface gigabitethernet 1/1
Force10(conf-if-gi-1/1)#no spanning-tree pvst vlan 3 priority 128
```

Figure 289 Disabling PVST+ Priority on an Interface

Viewing PVST+ Configuration

To view any changes to these values, enter the command **show spanning-tree pvst vlan** in EXEC mode.

Command Syntax	Command Mode	Usage
show spanning-tree pvst [**vlan** *vlan-id*] [**brief**] [*Interface*]	EXEC	View the PVST+ configuration.

Figure 290, shows the **brief** output of PVST+ instances configured on VLAN 3:

```
Force10#show spanning-tree pvst vlan 3 brief
VLAN 3
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 4096, Address 0001.e801.6aa8
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 16384, Address 0001.e805.e306
Configured hello time 2, max age 20, forward delay 15

Interface                                        Designated
 Name      PortID   Prio Cost    Sts Cost        Bridge ID        PortID
---------- -------- ---- ------  --- ------- -------------------- --------
Gi 1/0     128.130  128  20000   FWD 20000   4096  0001.e801.6aa8  128.426
Gi 1/1     128.131  128  20000   BLK 20000   4096  0001.e801.6aa8  128.427
Gi 1/16    128.146  128  20000   FWD 20000   16384 0001.e805.e306  128.146
Gi 1/17    128.147  128  20000   FWD 20000   16384 0001.e805.e306  128.147

Interface
 Name      Role   PortID   Prio Cost    Sts Cost    Link-type Edge
---------- ------ -------- ---- ------- --- ------- --------- ----
Gi 1/0     Root   128.130  128  20000   FWD 20000   P2P       No
Gi 1/1     Altr   128.131  128  20000   BLK 20000   P2P       No
Gi 1/16    Desg   128.146  128  20000   FWD 20000   P2P       Yes
Gi 1/17    Desg   128.147  128  20000   FWD 20000   P2P       Yes
```

Figure 290 show spanning-tree pvst brief Command Example

The following example shows viewing PVST+ instances configured on a VLAN 1 on an interface using the command **show spanning-tree pvst** [**vlan** *vlan-id*] [*interface*]:

```
Force10#show spanning-tree pvst vlan 1 interface gigabitethernet 0/0

GigabitEthernet 0/0 of VLAN 1 is designated forwarding

Edge port:yes port guard :none (default)
Link type: point-to-point (auto) bpdu filter:disable (default)
Bpdu guard :disable (default)
Bpdus sent 531, received 0

Interface                                      Designated
 Name     PortID   Prio Cost    Sts Cost       Bridge ID         PortID
--------- -------- ---- ------- --- ------- -------------------- --------
Gi 0/0    128.34   128  20000   FWD 0       32768 0001.e801.6aa8  128.34
```

Figure 291 show spanning-tree pvst vlan Command Example

The following example shows viewing PVST+ information using the **show configuration** and **show running-configuration** commands:

```
Force10#conf
Force10(conf)#interface gigabitethernet 1/1
Force10(conf-if-gi-1/1)#show config
!
interface GigabitEthernet 1/1
 no ip address
 switchport
 spanning-tree pvst vlan 3 cost 18000
 no shutdown
Force10(conf-if-gi-1/1)#end

Force10#show running-config interface gigabitethernet 1/1
!
interface GigabitEthernet 1/1
 no ip address
 switchport
 spanning-tree pvst vlan 3 cost 18000
 no shutdown
Force10#
```

Figure 292 Viewing PVST+ Configuration

Chapter 26 RSTP

Rapid Spanning Tree Protocol (RSTP) is a Layer 2 protocol—described by IEEE 802.1w—that is essentially the same as Spanning-Tree Protocol (STP) but provides faster convergence and interoperability with switches configured with STP and MSTP.

- For a complete listing of commands related to Rapid Spanning Tree Protocol, see the *C-Series FTOS Command Line Interface Reference*.

Implementation Information

C-Series	NO
E-Series	✓

Platform Specific Feature: Layer 2 BPDU Tunneling is supported on E-Series only.

E-Series employs Layer 2 BPDU filtering to reduce flapping in Spanning Tree configurations and add stability to Layer 2 networks. If Spanning Tree is enabled on a remote interface and disabled on the Force10 local interface, the Force10 system drops the BPDUs instead of forwarding them to the RPM CPU for processing.

Layer 2 BPDU filtering is available for STP, RSTP, and MSTP. It is enabled by default and is non-configurable.

Configuring Rapid Spanning Tree

Implementing Spanning Tree is a two-step process:

1. Configure interfaces for Layer 2. See page 444.
2. Enable Rapid Spanning Tree Protocol. See page 445.

Related Configuration Tasks

- Adding and Removing Interfaces on page 449
- Modifying Global Parameters on page 449
- Modifying Interface Parameters on page 450
- Configuring an Edge Port on page 451
- Influencing RSTP Root Selection on page 451

Important Points to Remember

- RSTP is disabled by default.
- FTOS supports only one Rapid Spanning Tree (RST) instance.
- All interfaces in VLANs and all enabled interfaces in Layer 2 mode are automatically added to the RST topology.

Table 36 displays the default values for RSTP.

Table 36 RSTP Default Values

RSTP Parameter		Default Value
Forward Delay		15 seconds
Hello Time		2 seconds
Max Age		20 seconds
Port Cost	100-Mb/s Ethernet interfaces	200000
	1-Gigabit Ethernet interfaces	20000
	10-Gigabit Ethernet interfaces	2000
	Port Channel with 100 Mb/s Ethernet interfaces	180000
	Port Channel with 1-Gigabit Ethernet interfaces	18000
	Port Channel with 10-Gigabit Ethernet interfaces	1800
Port Priority		128

Configuring Interfaces for Layer 2 Mode

All interfaces on all bridges that will participate in Rapid Spanning Tree must be in Layer 2 and enabled.

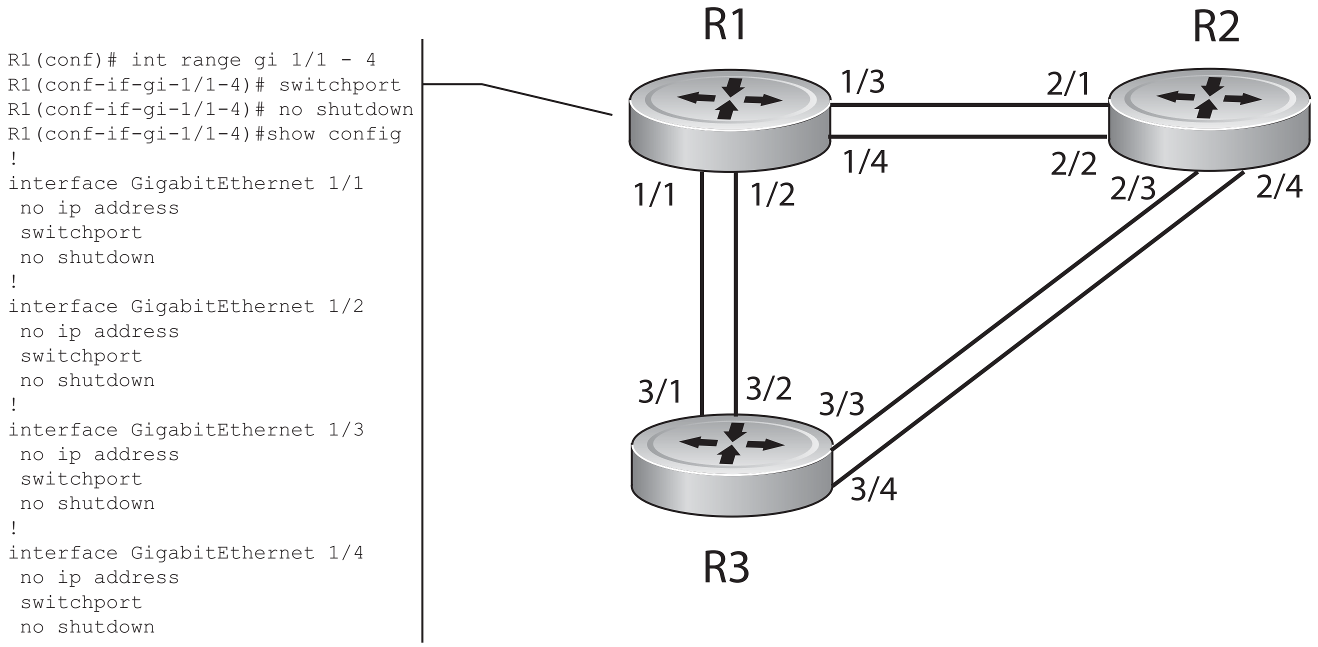

```
R1(conf)# int range gi 1/1 - 4
R1(conf-if-gi-1/1-4)# switchport
R1(conf-if-gi-1/1-4)# no shutdown
R1(conf-if-gi-1/1-4)#show config
!
interface GigabitEthernet 1/1
 no ip address
 switchport
 no shutdown
!
interface GigabitEthernet 1/2
 no ip address
 switchport
 no shutdown
!
interface GigabitEthernet 1/3
 no ip address
 switchport
 no shutdown
!
interface GigabitEthernet 1/4
 no ip address
 switchport
 no shutdown
```

Figure 293 Configuring Interfaces for Layer 2 Mode

To configure the interfaces for Layer 2 and then enable them:

Step	Task	Command Syntax	Command Mode
1	If the interface has been assigned an IP address, remove it.	**no ip address**	INTERFACE
2	Place the interface in Layer 2 mode.	**switchport**	INTERFACE
3	Enable the interface.	**no shutdown**	INTERFACE

Verify that an interface is in Layer 2 mode and enabled using the **show config** command from INTERFACE mode.

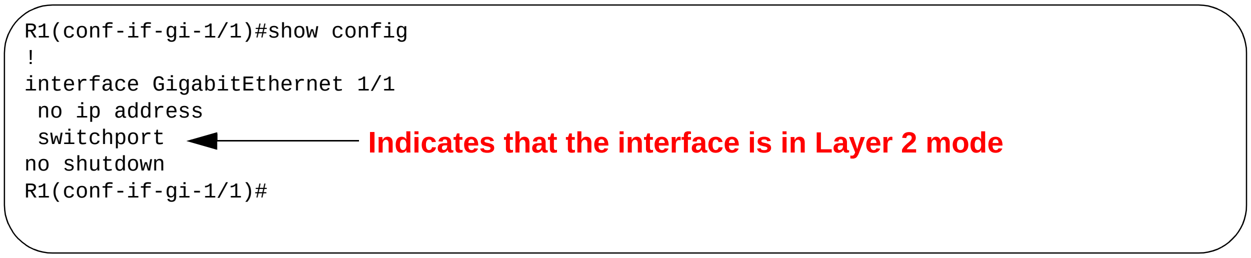

```
R1(conf-if-gi-1/1)#show config
!
interface GigabitEthernet 1/1
 no ip address
 switchport
no shutdown
R1(conf-if-gi-1/1)#
```

Figure 294 Verifying Layer 2 Configuration

Enabling Rapid Spanning Tree Protocol Globally

Rapid Spanning Tree Protocol must be enabled globally on all participating bridges; it is not enabled by default.

To enable Rapid Spanning Tree globally for all Layer 2 interfaces:

Step	Task	Command Syntax	Command Mode
1	Enter the PROTOCOL SPANNING TREE RSTP mode.	**protocol spanning-tree rstp**	CONFIGURATION
2	Enable Rapid Spanning Tree.	**no disable**	PROTOCOL SPANNING TREE RSTP

Note: To disable RSTP globally for all Layer 2 interfaces, enter the **disable** command from PROTOCOL SPANNING TREE RSTP mode.

Verify that Rapid Spanning Tree is enabled using the **show config** command from PROTOCOL SPANNING TREE RSTP mode.

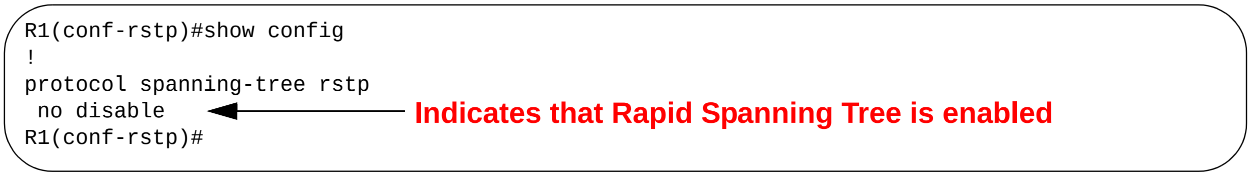

```
R1(conf-rstp)#show config
!
protocol spanning-tree rstp
 no disable
R1(conf-rstp)#
```

Figure 295 Verifying RSTP is Enabled

When you enable Rapid Spanning Tree, all physical and port-channel interfaces that are enabled and in Layer 2 mode are automatically part of the RST topology.

- Only one path from any bridge to any other bridge is enabled.
- Bridges block a redundant path by disabling one of the link ports.

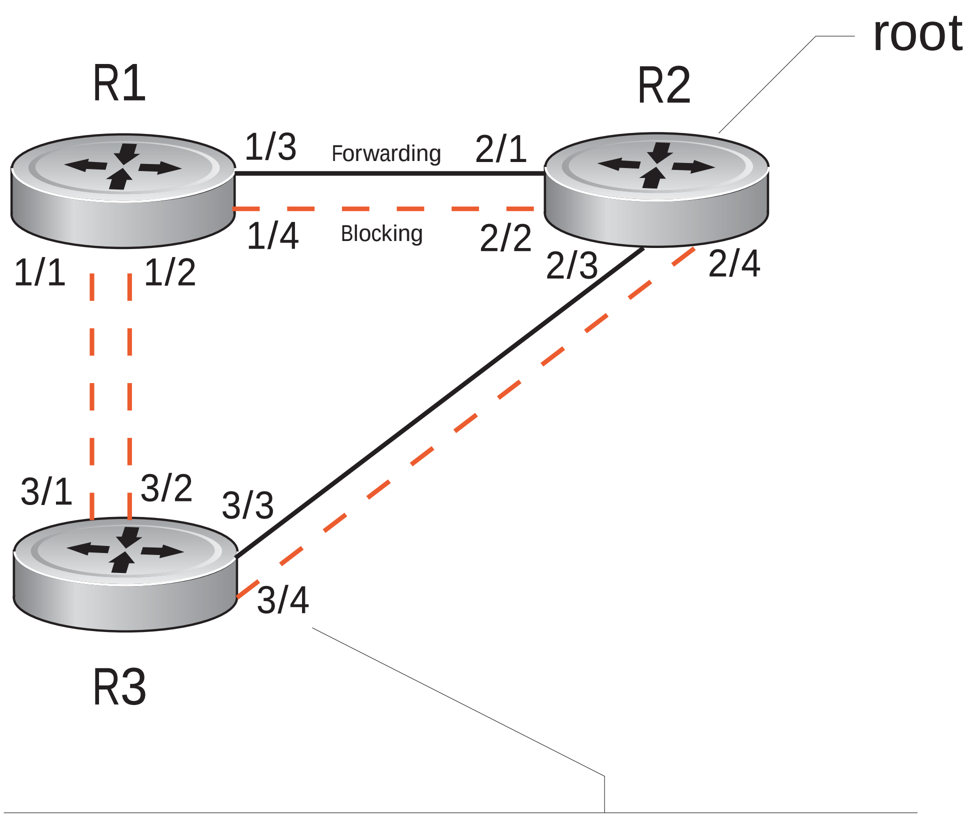

```
Port 684 (GigabitEthernet 4/43) is alternate Discarding
Port path cost 20000, Port priority 128, Port Identifier 128.684
Designated root has priority 32768, address 0001.e801.cbb4
Designated bridge has priority 32768, address 0001.e801.cbb4
Designated port id is 128.684, designated path cost 20000
Number of transitions to forwarding state 0
BPDU : sent 3, received 219
The port is not in the Edge port mode
```

Figure 296 Rapid Spanning Tree Enabled Globally

View the interfaces participating in Rapid Spanning Tree using the **show spanning-tree rstp** command from EXEC privilege mode. If a physical interface is part of a port channel, only the port channel is listed in the command output.

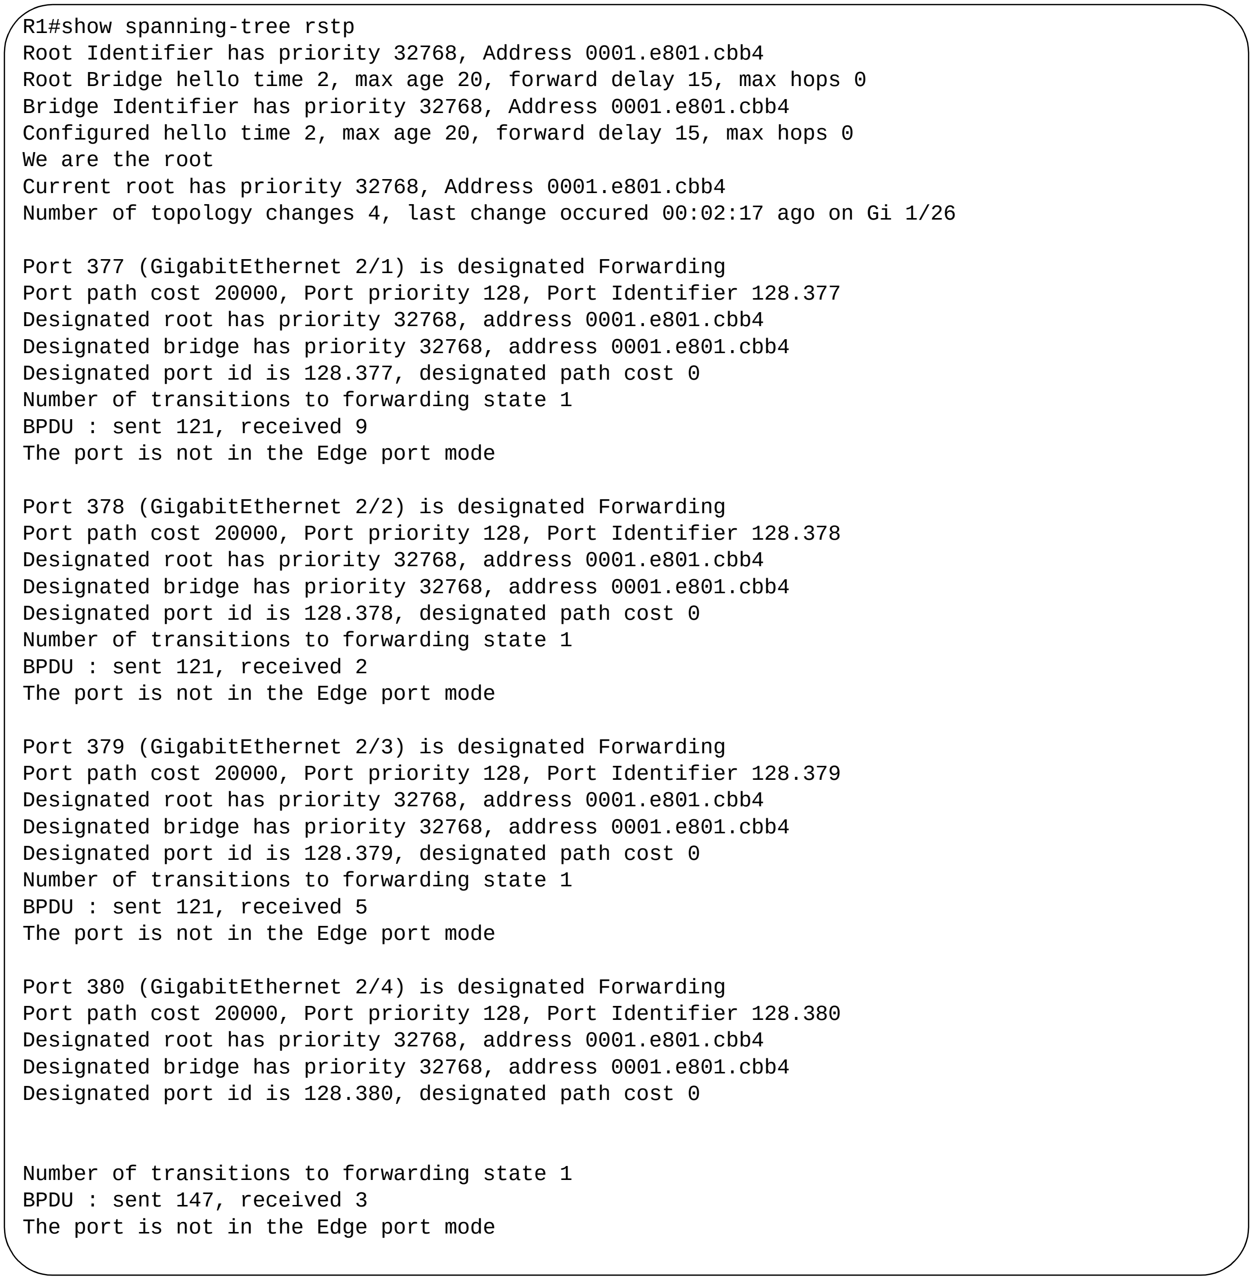

```
R1#show spanning-tree rstp
Root Identifier has priority 32768, Address 0001.e801.cbb4
Root Bridge hello time 2, max age 20, forward delay 15, max hops 0
Bridge Identifier has priority 32768, Address 0001.e801.cbb4
Configured hello time 2, max age 20, forward delay 15, max hops 0
We are the root
Current root has priority 32768, Address 0001.e801.cbb4
Number of topology changes 4, last change occured 00:02:17 ago on Gi 1/26

Port 377 (GigabitEthernet 2/1) is designated Forwarding
Port path cost 20000, Port priority 128, Port Identifier 128.377
Designated root has priority 32768, address 0001.e801.cbb4
Designated bridge has priority 32768, address 0001.e801.cbb4
Designated port id is 128.377, designated path cost 0
Number of transitions to forwarding state 1
BPDU : sent 121, received 9
The port is not in the Edge port mode

Port 378 (GigabitEthernet 2/2) is designated Forwarding
Port path cost 20000, Port priority 128, Port Identifier 128.378
Designated root has priority 32768, address 0001.e801.cbb4
Designated bridge has priority 32768, address 0001.e801.cbb4
Designated port id is 128.378, designated path cost 0
Number of transitions to forwarding state 1
BPDU : sent 121, received 2
The port is not in the Edge port mode

Port 379 (GigabitEthernet 2/3) is designated Forwarding
Port path cost 20000, Port priority 128, Port Identifier 128.379
Designated root has priority 32768, address 0001.e801.cbb4
Designated bridge has priority 32768, address 0001.e801.cbb4
Designated port id is 128.379, designated path cost 0
Number of transitions to forwarding state 1
BPDU : sent 121, received 5
The port is not in the Edge port mode

Port 380 (GigabitEthernet 2/4) is designated Forwarding
Port path cost 20000, Port priority 128, Port Identifier 128.380
Designated root has priority 32768, address 0001.e801.cbb4
Designated bridge has priority 32768, address 0001.e801.cbb4
Designated port id is 128.380, designated path cost 0


Number of transitions to forwarding state 1
BPDU : sent 147, received 3
The port is not in the Edge port mode
```

Figure 297 show spanning-tree rstp Command Example

Confirm that a port is participating in Rapid Spanning Tree using the **show spanning-tree rstp brief** command from EXEC privilege mode.

```
R3#show spanning-tree rstp brief
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 0001.e801.cbb4
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 32768, Address 0001.e80f.1dad
Configured hello time 2, max age 20, forward delay 15
Interface                                        Designated
 Name      PortID   Prio Cost    Sts Cost        Bridge ID        PortID
---------- -------- ---- ------- --- ------- -------------------- --------
Gi 3/1     128.681  128  20000   BLK 20000   32768 0001.e80b.88bd 128.469
Gi 3/2     128.682  128  20000   BLK 20000   32768 0001.e80b.88bd 128.470
Gi 3/3     128.683  128  20000   FWD 20000   32768 0001.e801.cbb4 128.379
Gi 3/4     128.684  128  20000   BLK 20000   32768 0001.e801.cbb4 128.380
Interface
 Name      Role   PortID   Prio Cost    Sts Cost    Link-type Edge
---------- ------ -------- ---- ------- --- ------- --------- ----
Gi 3/1     Altr   128.681  128  20000   BLK 20000   P2P       No
Gi 3/2     Altr   128.682  128  20000   BLK 20000   P2P       No
Gi 3/3     Root   128.683  128  20000   FWD 20000   P2P       No
Gi 3/4     Altr   128.684  128  20000   BLK 20000   P2P       No
R3#
```

Figure 298 show spanning-tree rstp brief Command Example

Adding and Removing Interfaces

- To add an interface to the Rapid Spanning Tree topology, configure it for Layer 2 and it is automatically added.
- To remove an interface from the Rapid Spanning Tree topology, remove its Layer 2 configuration using the **no switchport** command.

Modifying Global Parameters

You can modify Rapid Spanning Tree parameters. The root bridge sets the values for forward-delay, hello-time, and max-age and overwrites the values set on other bridges participating in the Rapid Spanning Tree group.

- **Forward-delay** is the amount of time an interface waits in the Listening State and the Learning State before it transitions to the Forwarding State.
- **Hello-time** is the time interval in which the bridge sends RSTP Bridge Protocol Data Units (BPDUs).
- **Max-age** is the length of time the bridge maintains configuration information before it refreshes that information by recomputing the RST topology.

Note: Force10 Networks recommends that only experienced network administrators change the Rapid Spanning Tree group parameters. Poorly planned modification of the RSTG parameters can negatively impact network performance.

To change these parameters, use the following commands, on the root bridge:

Task	Command Syntax	Command Mode
Change the forward-delay parameter. • Range: 4 to 30 • Default: 15 seconds	**forward-delay** *seconds*	PROTOCOL SPANNING TREE RSTP
Change the hello-time parameter. **Note:** With large configurations (especially those with more ports) Force10 Networks recommends that you increase the hello-time. Range: 1 to 10 Default: 2 seconds	**hello-time** *seconds*	PROTOCOL SPANNING TREE RSTP
Change the max-age parameter. Range: 6 to 40 Default: 20 seconds	**max-age** *seconds*	PROTOCOL SPANNING TREE RSTP

View the current values for global parameters using the **show spanning-tree rstp** command from EXEC privilege mode. See Figure 297.

Modifying Interface Parameters

On interfaces in Layer 2 mode, you can set the port cost and port priority values.

- **Port cost** is a value that is based on the interface type. The default values are listed in Table 36. The greater the port cost, the less likely the port will be selected to be a forwarding port.
- **Port priority** influences the likelyhood that a port will be selected to be a forwarding port in case that several ports have the same port cost.

To change the port cost or priority of an interface, use the following commands:

Task	Command Syntax	Command Mode
Change the port cost of an interface. Range: 0 to 65535 Default: see Table 36.	**spanning-tree rstp cost** *cost*	INTERFACE
Change the port priority of an interface. Range: 0 to 15 Default: 8	**spanning-tree rstp priority** *priority-value*	INTERFACE

View the current values for interface parameters using the **show spanning-tree rstp** command from EXEC privilege mode. See Figure 297.

Configuring an Edge Port

The EdgePort feature enables interfaces to begin forwarding traffic approximately 30 seconds sooner. In this mode an interface forwards frames by default until it receives a BPDU that indicates that it should behave otherwise; it does not go through the Learning and Listening states. This feature is the same as PortFast mode in Spanning Tree.

Caution: Configure EdgePort only on links connecting to an end station. EdgePort can cause loops if it is enabled on an interface connected to a network.

To enable EdgePort on an interface, use the following command:

Task	Command Syntax	Command Mode
Enable EdgePort on an interface.	**spanning-tree rstp edge-port**	INTERFACE

Verify that EdgePort is enabled on a port using the **show spanning-tree rstp** command from the EXEC privilege mode or the **show config** command from INTERFACE mode; Force10 recommends using the **show config** command, as shown in Figure 299.

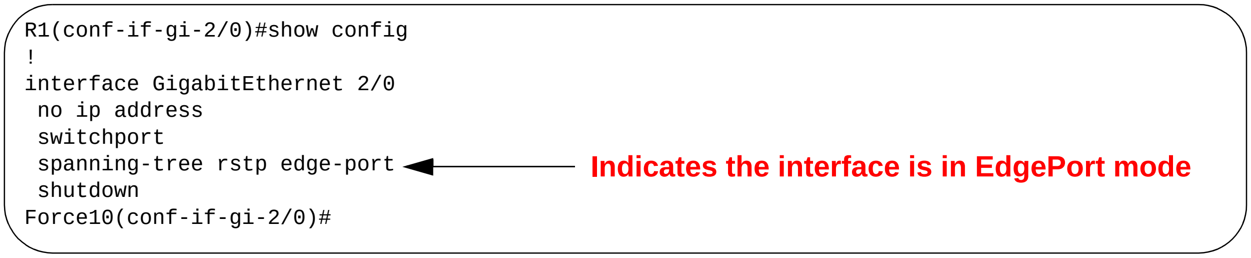

```
R1(conf-if-gi-2/0)#show config
!
interface GigabitEthernet 2/0
 no ip address
 switchport
 spanning-tree rstp edge-port
 shutdown
Force10(conf-if-gi-2/0)#
```

Indicates the interface is in EdgePort mode

Figure 299 EdgePort Enabled on Interface

Influencing RSTP Root Selection

The Rapid Spanning Tree Protocol determines the root bridge, but you can assign one bridge a lower priority to increase the likelihood that it will be selected as the root bridge.

To change the bridge priority, use the following command:

Task	Command Syntax	Command Mode
Assign a number as the bridge priority or designate it as the primary or secondary root. *priority-value* range: 0 to 65535. The lower the number assigned, the more likely this bridge will become the root bridge. The default is 32768. Entries must be multiples of 4096.	**bridge-priority** *priority-value*	PROTOCOL SPANNING TREE RSTP

A console message appears when a new root bridge has been assigned. Figure 300 shows the console message after the **bridge-priorty** command is used to make R2 the root bridge.

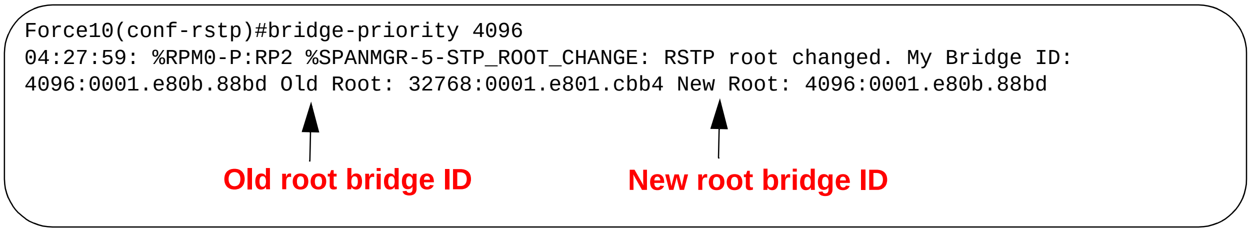

```
Force10(conf-rstp)#bridge-priority 4096
04:27:59: %RPM0-P:RP2 %SPANMGR-5-STP_ROOT_CHANGE: RSTP root changed. My Bridge ID:
4096:0001.e80b.88bd Old Root: 32768:0001.e801.cbb4 New Root: 4096:0001.e80b.88bd
```

Figure 300 bridge-priority Command Example

Chapter 27 MSTP

IEEE 802.1s MSTP (Multiple Spanning Tree Protocol) maps a group of Virtual Local Area Networks (VLANs) to a reduced number of spanning-tree instances. This supplement to IEEE 802.1Q allows VLAN bridges to use multiple spanning trees. This protocol enables network traffic from different VLANs to flow through different potential paths within a bridged VLAN. Because most networks do not need more than a few logical topologies, this feature provides design flexibility as well as better overall network resource utilization. This is because it facilitates the sharing of traffic loads across multiple forwarding paths.

The benefits of MSTP include:

- Network topologies can be designed to be optimal for a VLAN or a set of VLANs
- Traffic loads can be distributed along links
- CPU capacity is used efficiently
- Fault tolerance increases caused by a failure in one MSTP instance do not impact other instances

For more information see Fast Convergence after MSTP-Triggered Topology Changes on page 570.

For complete information on Multiple Spanning Tree Protocol, please see the IEEE Standard 802.1s.

The following sections describe MSTP in FTOS:

- MSTP Interoperability on page 453
- MSTP Implementation on page 454
- Configuration Task List for Multiple Spanning Tree Protocol

MSTP Interoperability

FTOS implementation of MSTP interoperates with any other standard-based implementation, and does not interoperate with routers that have a proprietary configuration digest and PDU. Force10's MSTP is IEEE compliant and interoperates with other implementations of MSTP that is IEEE-compliant.

FTOS MSTP allows users to map between a set of VLANs and an MSTP instance. As per the MSTP standard (IEEE 802.1s), an HMAC-MD5 digest of this mapping is carried in the MSTP BPDU. This standard also specifies a key to be used to generate the digest. FTOS MSTP uses this standard-specified key. However, other vendors may use a different key to generate the digest. A switch that uses a non-standard key, does not interoperate with FTOS.

Hence, the same configuration (mapping between MSTP instances and VLANs) on FTOS is the same as that on other switches, the resulting value from the HMAC-MD5 calculation is different. When the Force10 router receives the MSTP BPDU from another switch, it compares the value of this configuration identifier field with what it expects to receive. As per the MSTP standard, if there is a mismatch in the value of the Configuration Identifier field of the MSTP BPDU, the peer switch is considered to be in a different MSTP region. While the topology will still be loopless, there cannot be fast convergence during failover if the switches are, or appear to be, in different regions.

In order for two MSTP switches to be considered in the same region, the HMAC-MD5 digest carried in the BPDUs must match.

MSTP Implementation

The FTOS implementation of MSTP is compliant with the IEEE specification. When MSTP is enabled, all ports in VLANs and all interfaces that are in Layer-2 mode are added to MSTP.

Important Things to Remember

- By default, MSTP is disabled.
- MSTP is supported only on line card series ED, EE, EF, and above.

The Table displays the default values for MSTP parameters:

Table 37 E-Series MSTP Default Values

MSTP Parameter	Default Value
Forward Delay	15 seconds
Hello Time	2 seconds
Max Age	20 seconds
Port Cost	200000 = 100 Mb/s Ethernet interfaces 20000 = 1-Gigabit Ethernet interfaces 2000 = 10-Gigabit Ethernet interfaces 200000 = Port Channel with 100 Mb/s Ethernet interfaces 20000 = Port Channel with 1-Gigabit Ethernet interfaces 2000 = Port Channel with 10-Gigabit Ethernet interfaces 180000 = Port Channel with 2 100-Mbps Ethernet interfaces 18000 = Port Channel with 2 1-Gigabit Ethernet interfaces 1800 = Port Channel with 2 10-Gigabit Eternity interfaces
Port Priority	128
Bridge Priority for MST Instance	32768

To allow for a larger number of ports in a switch, the port priority field borrows from the port number field that the port identifies. As specified in IEEE Standard 802.1s, the port priority field can range from 0 to 240 in steps of 16. Similarly as described in the standard, the bridge priority can also range from 0 through 61440 in steps of 4096.

Note: SNMP support for MSTP is not available.

FTOS supports the following MSTP features:

- Single region
- 64 instances
- 4000 VLANs with 48 ports
- 100 VLANs with 336 ports

Implementation Information

C-Series	NO
E-Series	✓

Platform Specific Feature: Layer 2 BPDU Tunneling is supported on E-Series only.

E-Series employs Layer 2 BPDU filtering to reduce flapping in Spanning Tree configurations and add stability to Layer 2 networks. If Spanning Tree is enabled on a remote interface and disabled on the Force10 local interface, the Force10 system drops the BPDUs instead of forwarding them to the RPM CPU for processing.

Layer 2 BPDU filtering is available for STP, RSTP, and MSTP. It is enabled by default and is non-configurable.

Configuration Task List for Multiple Spanning Tree Protocol

The following list includes the configuration tasks for Multiple Spanning Tree Protocol:

- enable MSTP globally on page 456
- map VLANs to instances on page 458
- disable or re-enable MSTP on interfaces on page 459
- modify global MSTP parameters on page 459
- set MSTP interface parameters on page 461
- influence MSTP root selection on page 462
- enable edge-ports on page 462

For a complete listing of all commands related to Multiple Spanning Tree Protocol, see the *FTOS Command Line Interface Reference*.

enable MSTP globally

By default, Multiple Spanning Tree Protocol is not enabled in FTOS. To enable MSTP globally in FTOS, use these commands in the following sequence in CONFIGURATION mode:

Step	Command Syntax	Command Mode	Usage
1	**protocol spanning-tree mstp**	CONFIGURATION	Enter the MSTP mode
2	**no disable**	PROTOCOL MSTP	Enable Multiple Spanning Tree Protocol

MSTP runs on all the enabled ports in the VLANs and those running in Layer-2 mode unless you have explicitly disabled the command with **no spanning-tree** keywords in the port configuration. By default, FTOS assigns all VLANs to instance 0.

To view the Multiple Spanning Tree Instance (MSTI) and the interfaces in that instance, use the **show spanning-tree msti** *instance-number* command or the **show spanning-tree msti** *instance-number* **brief** command. The following examples show the **show spanning-tree msti** command for instances 0 and 1.

```
Force10#show spanning-tree msti 0
MSTI 0 VLANs mapped  1-100, 111-4094

Bridge Identifier has priority 32768, Address 0001.e800.0a5c
Configured hello time 2, max age 20, forward delay 15, max hops 20
Current root has priority 32768, Address 0001.e800.0a5c
Number of topology changes 0, last change occurred 47765

Port 58 (GigabitEthernet 1/0) is designated Forwarding
Port path cost 0, Port priority 128, Port Identifier 128.58
Designated root has priority 32768, address 0001.e800.0a:5c
Designated bridge has priority 32768, address 0001.e800.0a:5c
Designated port id is 128.58, designated path cost
Number of transitions to forwarding state 1
BPDU (Mrecords): sent 305, received 0
The port is not in the portfast mode

Port 64 (GigabitEthernet 1/6) is designated Forwarding
Port path cost 0, Port priority 128, Port Identifier 128.64
Designated root has priority 32768, address 0001.e800.0a:5c
Designated bridge has priority 32768, address 0001.e800.0a:5c
Designated port id is 128.64, designated path cost
Number of transitions to forwarding state 1
BPDU (Mrecords): sent 307, received 39
The port is not in the portfast mode

Port 70 (GigabitEthernet 1/12) is designated Forwarding
Port path cost 0, Port priority 128, Port Identifier 128.70
Designated root has priority 32768, address 0001.e800.0a:5c
Designated bridge has priority 32768, address 0001.e800.0a:5c
Designated port id is 128.70, designated path cost
Number of transitions to forwarding state 1
BPDU (Mrecords): sent 307, received 341
The port is not in the portfast mode
```

Figure 301 show spanning-tree msti 0 Command Example

```
Force10#show spanning-tree msti 1
MSTI 1 VLANs mapped  101-110

Bridge Identifier has priority 32768, Address 0001.e802.3506
Configured hello time 2, max age 20, forward delay 15, max hops 20
Current root has priority 16384, Address 0001.e800.0a5c
Number of topology changes 1, last change occurred 60184

Port 82 (GigabitEthernet 2/0) is designated Forwarding
Port path cost 0, Port priority 128, Port Identifier 128.82
Designated root has priority 16384, address 0001.e800.0a:5c
Designated bridge has priority 32768, address 0001.e802.35:06
Designated port id is 128.82, designated path cost
Number of transitions to forwarding state 1
BPDU (Mrecords): sent 413, received 0
The port is not in the portfast mode

Port 88 (GigabitEthernet 2/6) is alternate Discarding
Port path cost 0, Port priority 128, Port Identifier 128.88
Designated root has priority 16384, address 0001.e800.0a:5c
Designated bridge has priority 16384, address 0001.e800.0a:5c
Designated port id is 128.88, designated path cost
Number of transitions to forwarding state 1
BPDU (Mrecords): sent 20, received 399
The port is not in the portfast mode

Port 94 (GigabitEthernet 2/12) is root Forwarding
Port path cost 0, Port priority 128, Port Identifier 128.94
Designated root has priority 16384, address 0001.e800.0a:5c
Designated bridge has priority 16384, address 0001.e800.0a:5c
Designated port id is 128.94, designated path cost
Number of transitions to forwarding state 2
BPDU (Mrecords): sent 810, received 399
The port is not in the portfast mode
```

Figure 302 show spanning-tree msti 1 Command Example

map VLANs to instances

To map VLANs to instances, use the **msti** command. For more information about this command, please see Figure 303 or the *FTOS Command Line Interface Reference*.

```
Force10(conf)#protocol spanning-tree mstp
Force10(conf-mstp)#msti 1 vlan 101-110
Force10(conf-mstp)#show config
!
protocol spanning-tree mstp
no disable
name CustomerSvc
revision 2
MSTI 1 VLAN101-110
```

Figure 303 Example of msti VLAN mapping

disable or re-enable MSTP on interfaces

To disable MSTP for an interface, use the **no spanning-tree** command. Use the **spanning-tree** command to re-enable MSTP if you have disabled it.

Command Syntax	Command Mode	Usage
no spanning-tree	INTERFACE	Disable MSTP.
spanning-tree	INTERFACE	Re-enable MSTP after it has been disabled.

After you enable MSTP globally, to enable physical and port channel interfaces in Layer-2 mode, FTOS includes them in the multiple spanning-tree. When you enable MSTP, the interfaces in Layer-2 mode start sending Bridge Protocol Data Units (BPDUs). MSTP allows VLAN, Loopback, and Null interfaces do not participate in MSTP.

Layer-3 interfaces also do not participate in the spanning-tree protocol and are not listed by the **show spanning-tree msti** *instance-number* commands. FTOS only lists the port-channels in the **show spanning-tree msti** *instance-number* command. It does not list the channel members of the port channels.

Figure 304 demonstrates how to use the **show spanning-tree msti 0 brief** command to verify your configuration for instance 0.

```
Force10#show spanning-tree msti 0 brief
MSTI 0 VLANs mapped  1-100, 111-4094

Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32768, Address 0001.e800.0a5c
Root Bridge hello time 2, max age 20, forward delay 15, max hops 20
Bridge ID    Priority 32768, Address 0001.e800.0a5c
Configured hello time 2, max age 20, forward delay 15, max hops 20

Interface                                     Designated
 Name      PortID  Prio Cost   Sts Cost       Bridge ID         PortID
---------- ------- ---- ------ --- ------ --------------------  -------
Gi 1/0     128.58  128  20000  FWD 0      32768 0001.e800.0a5c  128.58
Gi 1/6     128.64  128  20000  FWD 0      32768 0001.e800.0a5c  128.64
Gi 1/12    128.70  128  20000  FWD 0      32768 0001.e800.0a5c  128.70

Interface
 Name      Role   PortID  Prio Cost   Sts Cost   Link-type
---------- ------ ------- ---- ------ --- ------ -----------
Gi 1/0     Desg   128.58  128  20000  FWD 0      P2P
Gi 1/6     Desg   128.64  128  20000  FWD 0      P2P
Gi 1/12    Desg   128.70  128  20000  FWD 0      P2P
Force10#
```

Figure 304 show spanning-tree msti 0 brief Command Example

modify global MSTP parameters

You can modify MSTP parameters in the PROTOCOL SPANNING-TREE MSTP configuration mode.

The parameters **forward-delay**, **hello-time**, and **max-age** are configurable in PROTOCOL SPANNING-TREE MSTP mode. The root bridge sets these three parameters and overwrites the values set on other bridges participating in Multiple Spanning Tree.

Other parameters that you can modify are **max-hops**, **region-name**, **revision number**, and **MSTI bridge-priority**. Bridge-priority is assigned per MSTP instance and must be assigned in steps of 4096.

Table 38 Example MSTP Configuration and Helps

Prompt	Description
`Force10(conf-mstp)# ?`	
`disable`	`Disable multiple spanning tree protocol globally`
`end`	`Exit from configuration mode`
`exit`	`Exit from multiple spanning tree configuration mode`
`forward-delay`	`Set the forward delay for the spanning tree`
`hello-time`	`Set the hello time for the spanning tree`
`max-age`	`Set the max age for the spanning tree`
`max-hops`	`MST max hop count`
`msti`	`MST instance`
`name`	`MST region name`
`no`	`Negate a command or set its defaults`
`revision`	`MST region revision`
`show`	`Show multiple spanning tree configuration`
`Force10(conf-mstp)#msti 1 ?`	
`vlan`	`VLAN identifier`
`bridge-priority`	`Bridge priority`

`Force10(conf-mstp)#msti 1 bridge priority ?`
`<0-61440>`
	`Bridge priority in increments of 4096 (default = 32768)`

You can view global parameters with the **show spanning-tree msti** *instance-number* command. The table below shows the default values for **forward-delay**, **hello-time**, **max-age**, **max-hops**, the **name** you gave the MSTP region, and the **revision** number assigned to the configuration.

Table 39 Additional Helps for Example MSTP Configuration

Prompt	Description

`Force10(conf-mstp)#forward-delay ?`
`<4-30>`
	`Forward delay in seconds (default = 15)`

`Force10(conf-mstp)#hello-time ?`
`<1-10>`
	`Hello time in seconds (default = 2)`

`Force10(conf-mstp)#max-age ?`
`<6-40>`
	`Max age in seconds (default = 20)`

`Force10(conf-mstp)#max-hops ?`
`<1-40>`
	`Max hop value (default = 20)`

`Force10(conf-mstp)#name ?`
`WORD`
	`Name (32 characters maximum)`

`Force10(conf-mstp)#name DevTestRegion ?`
`<1-10>`
	`Hello time in seconds (default = 2)`

`Force10(conf-mstp)#revision ?`
`<1-10>`
	`Revision`

Together, the MSTP region name, revision number and the instance-to-VLAN mapping determine the region to which the MSTP switch belongs.

To view the changed configuration (non-default), use the **show config** command in **protocol spanning-tree mstp** CONFIGURATION mode. Alternatively, **show running-config spanning-tree mstp** in the EXEC mode gives the same information.

set MSTP interface parameters

For interfaces in Layer-2 mode, you can set the port cost and port priority and also configure a port as an edge port. The default cost is assigned based on the interface speed. The default priority is 128. It can be assigned only in steps of 16.

In FTOS, the interface costs are set based on the IEEE 802.1s standard and are listed in the table below.

Table 40 Port Cost for Interface Types

Interface Type	Port Cost
1-Gigabit Ethernet	20000
10-Gigabit Ethernet	2000
100 Mbps Ethernet	200000
Port Channel with 1-Gigabit Ethernet	20000
Port Channel with 10-Gigabit Ethernet	2000
Port Channel with 100 Mbps Ethernet	200000
Port Channel with 2 1-Gigabit Ethernet	18000
Port Channel with 2 10-Gigabit Ethernet	1800
Port Channel with 2 100-Mbps Ethernet	180000

To change the port cost or priority of an interface, use either the **interface gigabit** *port-number* and **spanning-tree msti** commands in INTERFACE mode.

Table 41 Multiple Spanning Tree Port Cost and Priority Helps

Prompt	Description
Force10(conf-if)#spanning-tree ?	
<0-0> MSTI	STP and RSTP MSTP
Force10(conf-if)#spanning-tree mSTi ?	
<0-62>	Instance
Force10(conf-if)#spanning-tree mSTi 10 ?	
cost priority	Port cost Port priority
Force10(conf-if)#spanning-tree mSTi 10 cost ?	
<1-200000>	Port cost value
Force10(conf-if)#spanning-tree mSTi 10 priority ?	
<0-240>	Port priority value in increments of 16 (default = 128)

To view any changes in these values, enter the **show config** in INTERFACE context or **show running-config interface** command in EXEC mode.

influence MSTP root selection

According to the MSTP root switch selection algorithm, the switch with the lowest value for the bridge priority for a particular MSTP instance in an MSTP region will be chosen as the root switch. If two MSTP switches have the same bridge priority, the switch with a lower MAC address will be selected. To influence the root switch selection for a particular MSTP instance, you can assign one bridge a lower priority for that instance. This increases the likelihood that the switch will be selected as the MSTP root switch.

enable edge-ports

The edge-port feature enables interfaces to begin forwarding packets immediately after they are connected. With an edge-port enabled, an interface does not go through the Blocking and Learning states and forwards traffic sooner.

The edge-port command should be configured only on interfaces connected to end stations. To enable an edge-port on an interface, use the **spanning-tree mstp edge-port** command in INTERFACE context.

Chapter 28 RIP

C-Series	NO
E-Series	✓

Platform Specific Feature: RIP is supported on E-Series only.

Routing Information Protocol (RIP) is a distance-vector routing protocol, which tracks distances or hop counts to nearby routers.

This chapter covers the following topics:

- Protocol Overview on page 463
- Implementation Information on page 464
- Configuration Information on page 464

Protocol Overview

RIP is the oldest interior gateway protocol. There are two versions of RIP: RIP version 1 (RIPv1) and RIP version 2 (RIPv2). These versions are documented in RFCs 1058 and 2453.

RIPv1

RIPv1 uses hop counts as its metric to construct a table of routing information of the network and that routing table is sent as either a request or response message. In RIPv1, the protocol's packets are either one-time requests for all routing information or periodic responses (every 30 seconds) from other routers for routing information. RIP transports its responses or requests by means of UDP, port 520.

RIP must receive regular routing updates to maintain a correct routing table. Response messages containing a router's full routing table are transmitted every 30 seconds. If a router does not send an update within a certain amount of time, the hop count to that route is changed to unreachable (a route hop metric of 16 hops). Another timer sets the amount of time before the unreachable routes are removed from the routing table.

This first RIP version does not support VLSM or CIDR and is not widely used.

RIPv2

RIPv2 adds support for subnet fields in the RIP routing updates, thus qualifying it as a classless routing protocol. The RIPv2 message format includes entries for route tags, subnet masks, and next hop addresses. Another enhancement included in RIPv2 is multicasting for route updates on IP multicast address 224.0.0.9.

Implementation Information

FTOS supports both versions of RIP and allows you to configure one version globally and the other version or both versions on the interfaces. Furthermore, the E-Series supports 1,000 RIP routes.

Table 42 displays the defaults for RIP in FTOS.

Table 42 RIP Defaults in FTOS

Feature	Default
Interfaces running RIP	Listen to RIPv1 and RIPv2 Transmit RIPv1
RIP timers	update timer = 30 seconds invalid timer = 180 seconds holddown timer = 180 seconds flush timer = 240 seconds
Auto summarization	Enabled
ECMP paths supported	16

Configuration Information

To configure RIP, you must use commands in two modes: ROUTER RIP and INTERFACE. Commands executed in the ROUTER RIP mode configure RIP globally on the E-Series while commands executed in the INTERFACE mode configure RIP features on that interface only.

By default, RIP is disabled in FTOS.

RIP is best suited for small, homogeneous networks. All devices within the RIP network must be configured to support RIP if they are to participate in the RIP

Configuration Task List for RIP

The following configuration steps include one mandatory step and several optional steps:

- enable RIP globally on page 465 (mandatory)
- configure RIP on interfaces on page 467 (optional)
- control RIP routing updates on page 467 (optional)
- set send and receive version on page 468 (optional)
- generate default route on page 470 (optional)
- control route metrics on page 471 (optional)
- summarize routes on page 471 (optional)
- control route metrics on page 471
- debug RIP on page 472

For a complete listing of all commands related to RIP, refer to

enable RIP globally

By default, RIP is not enabled in FTOS.

To enable RIP, use the following commands in sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**router rip**	CONFIGURATION	Enter ROUTER RIP mode and enable the RIP process on FTOS.
2	**network** *ip-address*	ROUTER RIP	Assign an IP network address as a RIP network to exchange routing information. You can use this command multiple times to exchange RIP information with as many RIP networks as you want.

After assigning networks with which the E-Series is to exchange RIP information, ensure that all devices on that network are configured to exchange RIP information.

FTOS default is to send RIPv1, and to receive RIPv1 and RIPv2. To change the RIP version globally, use the **version** command in the ROUTER RIP mode. For more information on changing the RIP version defaults, refer to .

When RIP is enabled, you can view the global RIP configuration by using the **show running-config** command in the EXEC mode or the **show config** command (Figure 305) in the ROUTER RIP mode.

```
Force10(conf-router_rip)#show config
!
router rip
 network 10.0.0.0
Force10(conf-router_rip)#
```

Figure 305 show config Command Example in ROUTER RIP mode

When the RIP process has learned the RIP routes, use the **show ip rip database** command in the EXEC mode to view those routes (Figure 306).

```
Force10#show ip rip database
Total number of routes in RIP database: 978
160.160.0.0/16
        [120/1] via 29.10.10.12, 00:00:26, Fa 0/0
160.160.0.0/16         auto-summary
2.0.0.0/8
        [120/1] via 29.10.10.12, 00:01:22, Fa 0/0
2.0.0.0/8              auto-summary
4.0.0.0/8
        [120/1] via 29.10.10.12, 00:01:22, Fa 0/0
4.0.0.0/8              auto-summary
8.0.0.0/8
        [120/1] via 29.10.10.12, 00:00:26, Fa 0/0
8.0.0.0/8              auto-summary
12.0.0.0/8
        [120/1] via 29.10.10.12, 00:00:26, Fa 0/0
12.0.0.0/8             auto-summary
20.0.0.0/8
        [120/1] via 29.10.10.12, 00:00:26, Fa 0/0
20.0.0.0/8             auto-summary
29.10.10.0/24          directly connected,Fa 0/0
29.0.0.0/8             auto-summary
31.0.0.0/8
        [120/1] via 29.10.10.12, 00:00:26, Fa 0/0
31.0.0.0/8             auto-summary
192.162.2.0/24
        [120/1] via 29.10.10.12, 00:01:21, Fa 0/0
192.162.2.0/24         auto-summary
192.161.1.0/24
        [120/1] via 29.10.10.12, 00:00:27, Fa 0/0
192.161.1.0/24         auto-summary
192.162.3.0/24
        [120/1] via 29.10.10.12, 00:01:22, Fa 0/0
192.162.3.0/24         auto-summary
```

Figure 306 show ip rip database Command Example (Partial)

To disable RIP globally, use the **no router rip** command in the CONFIGURATION mode.

configure RIP on interfaces

When you enable RIP globally on the E-Series, interfaces meeting certain conditions start receiving RIP routes. By default, interfaces that are enabled and configured with an IP address in the same subnet as the RIP network address receive RIPv1 and RIPv2 routes and send RIPv1 routes.

Assign IP addresses to interfaces that are part of the same subnet as the RIP network identified in the **network** command syntax.

control RIP routing updates

By default, RIP broadcasts routing information out all enabled interfaces but you can configure RIP to send or to block RIP routing information either from a specific IP address or a specific interface. To control which devices or interfaces receive routing updates, you must configure a direct update to one router and configure interfaces to block RIP updates from other sources.

To control the source of RIP route information, use the following commands, in the ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
neighbor *ip-address*	ROUTER RIP	Define a specific router to exchange RIP information between it and the E-Series. You can use this command multiple times to exchange RIP information with as many RIP networks as you want.
passive-interface *interface*	ROUTER RIP	Disable a specific interface from sending or receiving RIP routing information.

Another method of controlling RIP (or any routing protocol) routing information is to filter the information through a prefix list. Prefix lists are applied to incoming or outgoing routes and routes must meet the conditions of the prefix lists or FTOS drops the route. The prefix lists are globally applied on all interfaces running RIP. Configure the prefix list in the PREFIX LIST mode prior to assigning it to the RIP process.

For configuration information on prefix lists, see Chapter 17, IP Access Control Lists, Prefix Lists, and Route-maps, on page 301.

To apply prefix lists to incoming or outgoing RIP routes, use the following commands in the ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
distribute-list *prefix-list-name* **in**	ROUTER RIP	Assign a configured prefix list to all incoming RIP routes.
distribute-list *prefix-list-name* **out**	ROUTER RIP	Assign a configured prefix list to all outgoing RIP routes.

In addition to filtering routes, you can add routes from other routing instances or protocols to the RIP process. With the **redistribute** command syntax, you can include OSPF, static or directly connected routes in the RIP process.

To add routes from other routing instances or protocols, use any of the following commands in the ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
redistribute {**connected** | **static**} [**metric** *metric-value*] [**route-map** *map-name*]	ROUTER RIP	Include directly connected or user-configured (static) routes in RIP. • *metric* range: 0 to 16 • *map-name*: name of a configured route map.
redistribute isis [**level-1** | **level-1-2** | **level-2**] [**metric** *metric-value*] [**route-map** *map-name*]	ROUTER RIP	Include IS-IS routes in RIP. • *metric* range: 0 to 16 • *map-name*: name of a configured route map.
redistribute ospf *process-id* [**match external** {**1** | **2**} | **match internal**] [**metric** *value*] [**route-map** *map-name*]	ROUTER RIP	Include specific OSPF routes in RIP. Configure the following parameters: • *process-id* range: 1 to 65535 • *metric* range: 0 to 16 • *map-name*: name of a configured route map.

To view the current RIP configuration, use the **show running-config** command in the EXEC mode or the **show config** command in the ROUTER RIP mode.

set send and receive version

To specify the RIP version, use the **version** command in the ROUTER RIP mode. To set an interface to receive only one or the other version, use the **ip rip send version** or the **ip rip receive version** commands in the INTERFACE mode.

To change the RIP version globally in FTOS, use the following command in the ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
version {**1** | **2**}	ROUTER RIP	Set the RIP version sent and received on the E-Series.

You can set one RIP version globally on the E-Series. This command sets the RIP version for RIP traffic on the interfaces participating in RIP unless the interface was specifically configured for a specific RIP version.

Use the **show config** command in the ROUTER RIP mode to see whether the **version** command is configured. You can also use the **show ip protocols** command in the EXEC mode to view the routing protocols configuration.

Figure 307 shows an example of the RIP configuration after the ROUTER RIP mode **version** command is set to RIPv2. When the ROUTER RIP mode **version** command is set, the interface (GigabitEthernet 0/0) participating in the RIP process is also set to send and receive RIPv2.

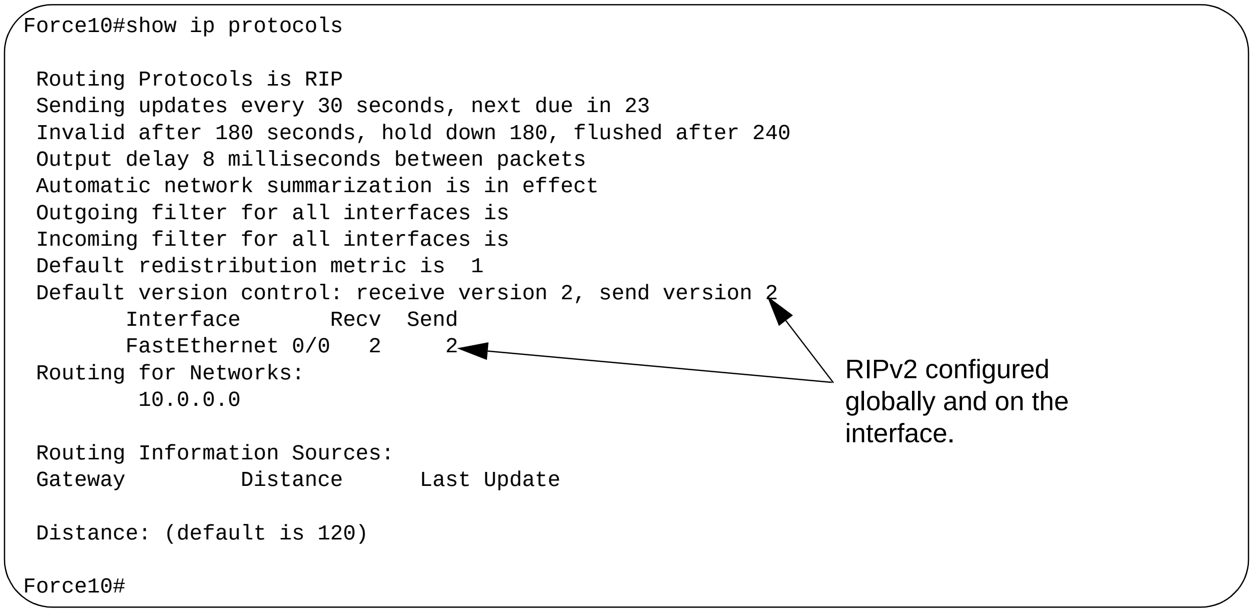

```
Force10#show ip protocols

 Routing Protocols is RIP
 Sending updates every 30 seconds, next due in 23
 Invalid after 180 seconds, hold down 180, flushed after 240
 Output delay 8 milliseconds between packets
 Automatic network summarization is in effect
 Outgoing filter for all interfaces is
 Incoming filter for all interfaces is
 Default redistribution metric is  1
 Default version control: receive version 2, send version 2
        Interface       Recv  Send
        FastEthernet 0/0   2     2
 Routing for Networks:
         10.0.0.0

 Routing Information Sources:
 Gateway         Distance      Last Update

 Distance: (default is 120)

Force10#
```

RIPv2 configured globally and on the interface.

Figure 307 show ip protocols Command Example

To configure the interfaces to send or receive different RIP versions from the RIP version configured globally, use either of the following commands in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip rip receive version [1] [2]	INTERFACE	Set the RIP version(s) received on that interface.
ip rip send version [1] [2]	INTERFACE	Set the RIP version(s) sent out on that interface.

To configure an interface to receive or send both versions of RIP, include 1 and 2 in the command syntax. Figure 308 displays the command syntax for sending both RIPv1 and RIPv2 and receiving only RIPv2.

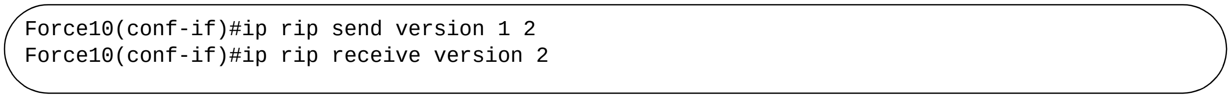

```
Force10(conf-if)#ip rip send version 1 2
Force10(conf-if)#ip rip receive version 2
```

Figure 308 Configuring an interface to send both versions of RIP

The **show ip protocols** command example (Figure 309) confirms that both versions are sent out that interface. This interface no longer sends and receives the same RIP versions as FTOS does globally.

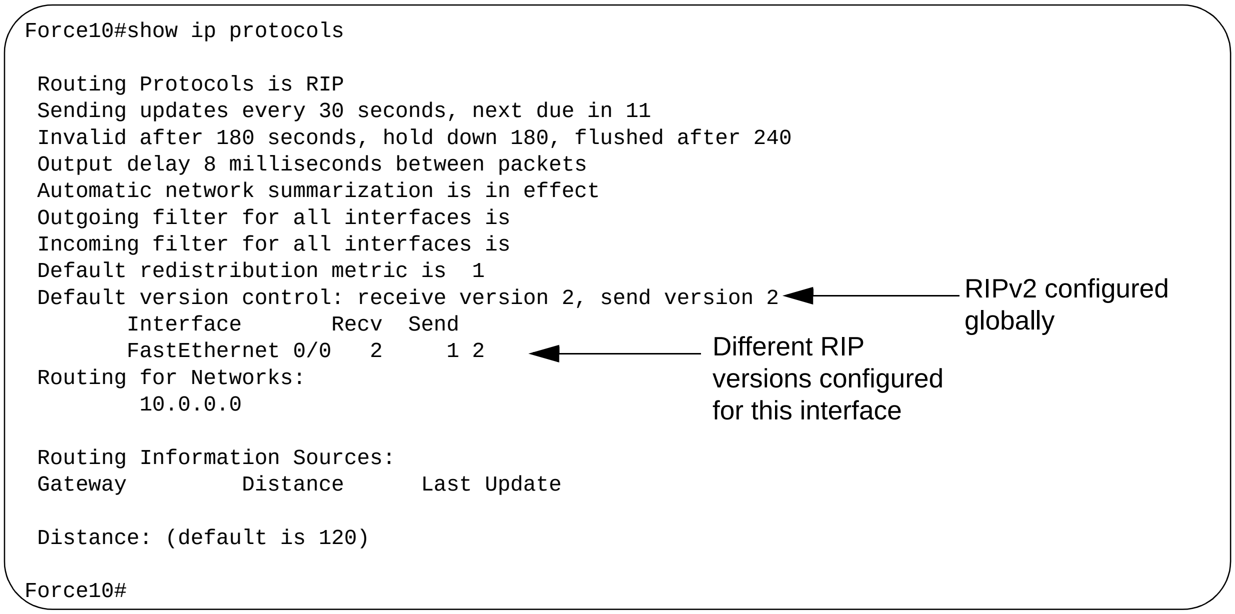

```
Force10#show ip protocols

 Routing Protocols is RIP
 Sending updates every 30 seconds, next due in 11
 Invalid after 180 seconds, hold down 180, flushed after 240
 Output delay 8 milliseconds between packets
 Automatic network summarization is in effect
 Outgoing filter for all interfaces is
 Incoming filter for all interfaces is
 Default redistribution metric is  1
 Default version control: receive version 2, send version 2
        Interface       Recv  Send
        FastEthernet 0/0   2     1 2
 Routing for Networks:
         10.0.0.0

 Routing Information Sources:
 Gateway         Distance      Last Update

 Distance: (default is 120)

Force10#
```

RIPv2 configured globally

Different RIP versions configured for this interface

Figure 309 show ip protocols Command Example

generate default route

Traffic is forwarded to the default route when the traffic's network is not explicitly listed in the routing table. Default routes are not enabled in RIP unless specified. Use the **default-information originate** command in the ROUTER RIP mode to generate a default route into RIP. In FTOS, default routes received in RIP updates from other routes are advertised if the **default-information originate** command is configured.

To configure FTOS to generate a default route, use the following command in the ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
default-information originate [**always**] [**metric** *value*] [**route-map** *route-map-name*]	ROUTER RIP	Specify the generation of a default route in RIP. Configure the following parameters: • **always**: enter this keyword to always generate a default route. • *value* range: 1 to 16. • *route-map-name*: name of a configured route map.

Use the **show config** command in the ROUTER RIP mode to confirm that the default route configuration is completed.

summarize routes

Routes in the RIPv2 routing table are summarized by default, thus reducing the size of the routing table and improving routing efficiency in large networks. By default, the **autosummary** command in the ROUTER RIP mode is enabled and summarizes RIP routes up to the classful network boundary.

If you must perform routing between discontiguous subnets, disable auto summarization. With automatic route summarization disabled, subnets are advertised.

The command **autosummary** requires no other configuration commands. To disable automatic route summarization, in the ROUTER RIP mode, enter **no autosummary**.

Note: If the **ip split-horizon** command is enabled on an interface, then the E-Series does not advertise the summarized address.

control route metrics

RIP is a distance-vector protocol and uses hop counts to determine the best route, but sometimes the shortest hop count is a route over the lowest speed link. To manipulate RIP routes so that the routing protocol prefers a different route, you must manipulate the route by using the **offset** command.

You must exercise caution when applying an **offset** command to routers on a broadcast network since the router using the **offset** command is modifying RIP advertisements before sending out those advertisements.

Another command, **distance,** also allows you to manipulate route metrics. With the **distance** command you assign different weights to routes so that the ones with the lower weight or administrative distance assigned are preferred.

To set route metrics, use either of the following commands in the ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
distance *weight* [*ip-address mask* [*access-list-name*]]	ROUTER RIP	Apply a weight to all routes or a specific route and ACL. Configure the following parameters: • *weight* range: 1 to 255 (default is 120) • *ip-address mask*: the IP address in dotted decimal format (A.B.C.D), and the mask in slash format (/x). • *access-list-name*: name of a configured IP ACL.

Command Syntax	Command Mode	Purpose
offset *access-list-name* {**in** | **out**} *offset* [*interface*]	ROUTER RIP	Apply an additional number to the incoming or outgoing route metrics. Configure the following parameters: • *access-list-name*: the name of a configured IP ACL • *offset* range: 0 to 16. • *interface*: the type, slot, and number of an interface.

Use the **show config** command in the ROUTER RIP mode to view configuration changes.

debug RIP

The **debug ip rip** command enables RIP debugging. When debugging is enabled, you can view information on RIP protocol changes or RIP routes.

To enable RIP debugging, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Purpose
debug ip rip [*interface* | **database** | **events** | **trigger**]	EXEC privilege	Enable debugging of RIP on the E-Series.

When you enable RIP debugging, you see a confirmation that the debug function was enabled (Figure 310).

```
Force10#debug ip rip
RIP protocol debug is ON
Force10#
```

Figure 310 debug ip rip Command Example

To disable RIP, use the **no debug ip rip** command syntax.

Chapter 29 OSPF

Open Shortest Path First (OSPF) is a link-state routing protocol designed to run within a single Autonomous System (AS).

This chapter covers the following topics:

Protocol Overview

OSPF is an Interior Gateway Protocol (IGP) and distributes routing information between routers within an Autonomous System (AS). OSPF is also a link-state protocol in which routers create forwarding tables based on network topology information collected from other routers in the network. Routers create a Link State Database (LSDB) that maintains the best paths between themselves and other routers. OSPF routers initially exchange hello messages to set up adjacencies with neighbor routers. If two routers on the same subnet agree to become neighbors through the hello process, then they will begin to exchange network topology information in the form of Link State Advertisements (LSAs).

To manage the routing information, the AS can be broken up into areas.

This overview is not intended to provide a complete understanding of OSPF; for that, consult the RFC 2328, *OSPF Version 2*.

Implementation Information

FTOS's implementation of OSPF is based on RFC 2328 and supports 10,000 OSPF routes, with 8,000 of those routes as external and 2,000 as inter/intra area routes.

FTOS supports the following LSAs:

- Router (type 1)
- Network (type 2)
- Network Summary (type 3)

- AS Boundary (type 4)
- AS External (type 5)
- NSSA External (type 7)
- Opaque Link-local (type 9)
- Opaque Area-local (type 10)
- Opaque Link-state (type 11)

FTOS also supports Stub areas and Not So Stubby Areas (NSSAs):

- Stub—No Type 5 AS-external LSA allowed.
  - Command: **area 0 stub**
- Totally Stub—No Type 3, 4 or 5 LSAs allowed except the default summary route.
  - Command: **area 0 stub no-summary**
- NSSA—No Type 5 AS-external LSAs allowed, but Type 7 LSAs that convert to Type 5 at the NSSA ABR can traverse.
  - Command: **area 0 nssa**
- NSSA Totally Stub—No Type 3, 4 or 5 LSAs except the default summary route, but Type 7 LSAs that convert to Type 5 at the NSSA ABR are allowed.
  - Command: **area 0 nssa no-summary**

FTOS supports the following RFCs:

- The OSPF NSSA Option (RFC 1587)
- OSPF Version 2 Management Information Base (RFC 1850)
- OSPF Version 2 (RFC 2328)
- The OSPF Opaque LSA Option (RFC 2370)
- Graceful OSPF Restart (RFC 3623)

RFC-2328 Compliant OSPF Flooding

In OSPF, flooding is the most resource-consuming task. The flooding algorithm described in RFC 2328 requires that OSPF flood LSAs on all interfaces, as governed by LSA's flooding scope. (Refer to Section 13 of the RFC.) When multiple direct links connect two routers, the RFC 2328 flooding algorithm generates significant redundant information across all links.

By default, FTOS implements an enhanced flooding procedure which dynamically and intelligently detects when to optimize flooding. Wherever possible, the OSPF task attempts to reduce flooding overhead by selectively flooding on a subset of the interfaces between two routers.

If RFC 2328 flooding behavior is required, the command **flood-2328** can be enabled in ROUTER OSPF mode. When enabled, this command configures FTOS to flood LSAs on all interfaces.

To confirm that this behavior is implemented, use the command **debug ip ospf packet** and look for output similar to the following:

Figure 311 Enabling RFC-2328 Compliant OSPF Flooding

```
00:10:41 : OSPF(1000:00):                            Printed only for ACK packets
Rcv. v:2 t:5(LSAck) l:64 Acks 2 rid:2.2.2.2
        aid:1500 chk:0xdbee aut:0 auk: keyid:0 from:Vl 1000
            LSType:Type-5 AS External id:160.1.1.0 adv:6.1.0.0 seq:0x8000000c
            LSType:Type-5 AS External id:160.1.2.0 adv:6.1.0.0 seq:0x8000000c
00:10:41 : OSPF(1000:00):
Rcv. v:2 t:5(LSAck) l:64 Acks 2 rid:2.2.2.2
        aid:1500 chk:0xdbee aut:0 auk: keyid:0 from:Vl 100
            LSType:Type-5 AS External id:160.1.1.0 adv:6.1.0.0 seq:0x8000000c
            LSType:Type-5 AS External id:160.1.2.0 adv:6.1.0.0 seq:0x8000000c
00:10:41 : OSPF(1000:00):
Rcv. v:2 t:4(LSUpd) l:100 rid:6.1.0.0       No change in update packets
        aid:0 chk:0xccbd aut:0 auk: keyid:0 from:Gi 10/21
            Number of LSA:2
            LSType:Type-5 AS External(5) Age:1 Seq:0x8000000c id:170.1.1.0 Adv:6.1.0.0
                Netmask:255.255.255.0 fwd:0.0.0.0 E2, tos:0 metric:0
            LSType:Type-5 AS External(5) Age:1 Seq:0x8000000c id:170.1.2.0 Adv:6.1.0.0
                Netmask:255.255.255.0 fwd:0.0.0.0 E2, tos:0 metric:0
```

In 7.5.1.0 you can use the command **show ip ospf** to confirm that RFC-2328 compliant OSPF flooding is enabled, as shown below.

Figure 312 Enabling RFC-2328 Compliant OSPF Flooding

```
Force10#show ip ospf
Routing Process ospf 1 with ID 2.2.2.2
Supports only single TOS (TOS0) routes
It is an Autonomous System Boundary Router
It is Flooding according to RFC 2328
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Number of area in this router is 1, normal 0 stub 0 nssa 1
--More--
```

OSPF ACK Packing

The OSPF ACK Packing feature bundles multiple LS acknowledgements in a single packet, significantly reducing the number of ACK packets transmitted when the number of LSAs increases. This feature also enhances network utilization and reduces the number of small ACK packets sent to a neighboring router. OSPF ACK packing is enabled by default, and non-configurable.

Configuration Information

To configure OSPF, you may use commands in two modes: ROUTER OSPF and INTERFACE. Commands in the ROUTER OSPF mode configure OSPF globally, while commands executed in the INTERFACE mode configure OSPF features on that interface only.

By default, OSPF is disabled.

Configuration Task List for OSPF

The following configuration steps include two mandatory steps and several optional ones:

- enable OSPF globally on page 476 (mandatory)
- enable OSPF on interfaces on page 477 (mandatory)
- configure stub areas on page 480 (optional)
- enable passive interfaces on page 480
- enable OSPF authentication on page 483 (optional)
- enable graceful restart on page 484 (optional)
- configure virtual links on page 486 (optional)
- filter routes on page 486 (optional)
- redistribute routes on page 487 (optional)
- troubleshooting OSPF on page 487

For a complete listing of all commands related to OSPF, refer to

enable OSPF globally

Before enabling OSPF globally, you must first assign an IP address to an interface (physical or Loopback) to enable Layer 3 routing. By default, the routing protocols, including OSPF, are disabled.

To enable routing, use these commands in the following sequence in the INTERFACE mode:

Step	Command Syntax	Command Mode	Usage
1	**ip address** *ip-address mask*	INTERFACE	Assign an IP address to an interface.
2	**no shutdown**	INTERFACE	Enable the interface.

After an IP address is assigned to an interface, enter the ROUTER OSPF mode and enable OSPF. To enter the ROUTER OSPF mode, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Usage
router ospf *process-id*	CONFIGURATION	Enables OSPF globally on the E-Series.

FTOS supports one OSPF routing process.

To view the current OSPF status, use the **show ip ospf** command in the EXEC mode (Figure 313).

```
Force10>show ip ospf
Routing Process ospf 1 with ID 11.1.2.1
Supports only single TOS (TOS0) routes
It is an autonomous system boundaryrouter
SPF schedule delay 0 secs, Hold time between two SPFs 5 secs
Number of area in this router is 1, normal 1 stub 0 nssa 0
   Area BACKBONE (0.0.0.0)
       Number of interface in this area is 2
       SPF algorithm executed 4 times
       Area ranges are
Force10>
```

Figure 313 show ip ospf Command Example

After OSPF is enabled, you must assign the interface to an OSPF area.

To disable OSPF, use the **no router ospf** *process-id* command syntax in the CONFIGURATION mode.

To reset the OSPF process, use the **clear ip ospf** command syntax.

enable OSPF on interfaces

You enable OSPF on an interface with the **network** command. You also set up OSPF areas with this command.

OSPF areas are a logical grouping of OSPF routers and links. An area is identified by an integer or dotted-decimal number. Each OSPF network consists of multiple OSPF areas and each area is connected, either directly or virtually to one area (Area ID 0.0.0.0). Area ID 0.0.0.0 is reserved for the OSPF backbone, which summarizes the other areas topologies and passes that information on to the other areas.

As a link-state protocol, OSPF sends routing information to other OSPF routers by means of the interfaces or links. The state (up or down) of those links is important. First, the interfaces must be in Layer-3 mode (that is, assigned an IP address) and enabled so that they can send and receive traffic. Second, the OSPF process must know about these links. To make the OSPF process aware of these interfaces, they must be assigned to OSPF areas.

The OSPF process evaluates the **network** commands in the order they are configured. Assign the network address that is most explicit first to include all subnets of that address. For example, if you assign the network address 90.0.0.0 /8, you cannot assign the network address 90.1.0.0 /16 since it is already included in the first network address.

When configuring the **network** command, you must configure a network address and mask that is a superset of the IP subnet configured on the Layer-3 interface to be used for OSPF.

If your OSPF network contains more than one area, you also must configure a backbone area (Area ID 0.0.0.0).

To enable OSPF on an interface, use the following command in the ROUTER OSPF mode:

Command Syntax	Command Mode	Usage
network *ip-address mask* **area** *area-id*	ROUTER OSPF	Enable OSPF on an interface and assign an network address range to a specific OSPF area.

Figure 314 presents an example of assigning an IP address to an interface and then assigning an OSPF area that includes that Layer-3 interface's IP address.

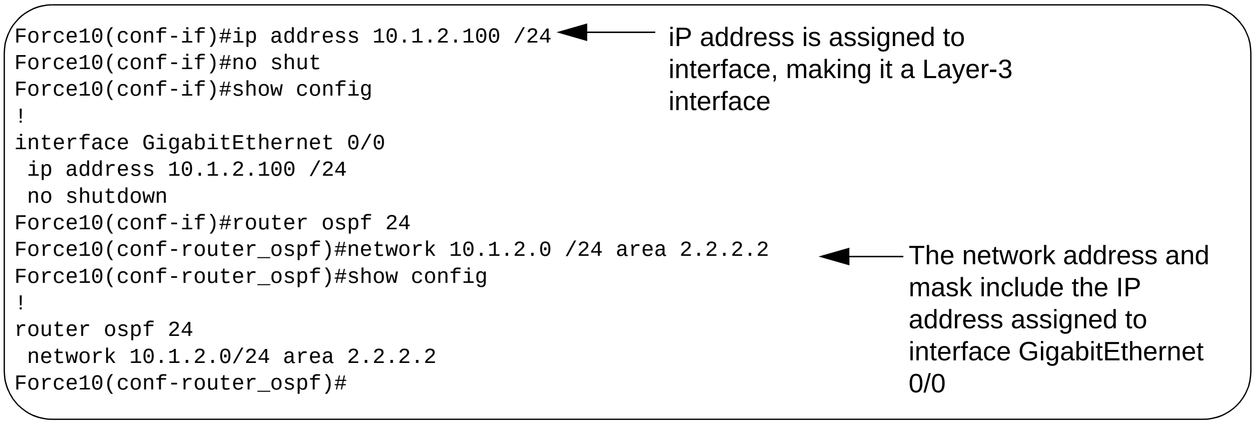

```
Force10(conf-if)#ip address 10.1.2.100 /24
Force10(conf-if)#no shut
Force10(conf-if)#show config
!
interface GigabitEthernet 0/0
 ip address 10.1.2.100 /24
 no shutdown
Force10(conf-if)#router ospf 24
Force10(conf-router_ospf)#network 10.1.2.0 /24 area 2.2.2.2
Force10(conf-router_ospf)#show config
!
router ospf 24
 network 10.1.2.0/24 area 2.2.2.2
Force10(conf-router_ospf)#
```

iP address is assigned to interface, making it a Layer-3 interface

The network address and mask include the IP address assigned to interface GigabitEthernet 0/0

Figure 314 Configuring an OSPF Area Example

The OSPF router ID is derived from the interface IP addresses. FTOS prefers the highest IP address assigned to a Loopback interface, even if the Loopback interface is not included in an OSPF network statement. If a Loopback interface with an IP address is not configured, then FTOS uses the highest IP address configured on an interface as the OSPF router ID. If you delete the interface with the IP address used to determine the OSPF router ID, the OSPF process resets.

To view the configuration, use the **show config** command in ROUTER OSPF mode.

OSPF, by default, sends hello packets out to all physical interfaces assigned an IP address that are a subset of a network on which OSPF is enabled. Use the **show ip ospf interface** command (Figure 315) to view the interfaces currently active and the areas assigned to the interfaces.

```
Force10>show ip ospf interface

GigabitEthernet 12/17 is up, line protocol is up
  Internet Address 10.2.2.1/24, Area 0.0.0.0
  Process ID 1, Router ID 11.1.2.1, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 11.1.2.1, Interface address 10.2.2.1
  Backup Designated Router (ID) 0.0.0.0, Interface address 0.0.0.0
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:04
  Neighbor Count is 0, Adjacent neighbor count is 0

GigabitEthernet 12/21 is up, line protocol is up
  Internet Address 10.2.3.1/24, Area 0.0.0.0
  Process ID 1, Router ID 11.1.2.1, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 13.1.1.1, Interface address 10.2.3.2
  Backup Designated Router (ID) 11.1.2.1, Interface address 10.2.3.1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 13.1.1.1 (Designated Router)
Force10>
```

Figure 315 show ip ospf interface Command Example

Loopback interfaces also assist in the OSPF process. OSPF will pick the highest interface address as the router-id and a loopback interface address has a higher precedence than other interface addresses.

Figure 316 gives an example of the **show ip ospf interface** command with a Loopback interface.

```
Force10#show ip ospf int

GigabitEthernet 13/23 is up, line protocol is up
  Internet Address 10.168.0.1/24, Area 0.0.0.1
  Process ID 1, Router ID 10.168.253.2, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DROTHER, Priority 1
  Designated Router (ID) 10.168.253.5, Interface address 10.168.0.4
  Backup Designated Router (ID) 192.168.253.3, Interface address 10.168.0.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:08
  Neighbor Count is 3, Adjacent neighbor count is 2
    Adjacent with neighbor 10.168.253.5 (Designated Router)
    Adjacent with neighbor 10.168.253.3 (Backup Designated Router)

Loopback 0 is up, line protocol is up
  Internet Address 10.168.253.2/32, Area 0.0.0.1
  Process ID 1, Router ID 10.168.253.2, Network Type LOOPBACK, Cost: 1
Loopback interface is treated as a stub Host.
Force10#
```

Figure 316 show ip ospf interface Command Example with Loopback Interface

configure stub areas

OSPF supports different types of LSAs to help reduce the amount of router processing within the areas. Type 5 LSAs are not flooded into stub areas, instead the Area Border Router (ABR) advertises a default route into the stub area to which it is attached. Stub area routers use the default route to reach external destinations

To ensure connectivity in your OSPF network, never configure the backbone area as a stub area.

To configure a stub area, use these commands in the following sequence, starting in the EXEC privilege mode:

Step	Command Syntax	Command Mode	Usage
1	**show ip ospf database database-summary**	EXEC privilege	Review all areas after they were configured to determine which areas are NOT receiving type 5 LSAs (listed in the S-ASBR column in Figure 317).
2	**configure**	EXEC privilege	Enter the CONFIGURATION mode.
3	**router ospf** *process-id*	CONFIGURATION	Enter the ROUTER OSPF mode.
4	**area** *area-id* **stub** [**no-summary**]	ROUTER OSPF	Configure the area as a stub area. Use the **no-summary** keywords to prevent transmission in to the area of summary ASBR LSAs.

To view which LSAs are transmitted, use the **show ip ospf database database-summary** command syntax (Figure 317) in the EXEC privilege mode.

```
Force10#show ip ospf database database-summary

             OSPF Router with ID (10.1.2.100) (Process ID 34)

Area ID          Router    Network S-Net    S-ASBR   Type-7    Subtotal
2.2.2.2          1         0       0        0        0         1
3.3.3.3          1         0       0        0        0         1
Force10#
```

Figure 317 show ip ospf database database-summary Command Example

To view information on areas, use the **show ip ospf** command in the EXEC privilege mode (Figure 313).

enable passive interfaces

The OSPF process always advertises the IP address of an interface participating in the OSPF process, but you can suppress the OSPF process on an interface.

To suppress the interface's participation in the OSPF process, use the following command in the ROUTER OSPF mode:

Command Syntax	Command Mode	Usage
passive-interface *interface*	ROUTER OSPF	Specify the physical interface type, slot, and number. • For a Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a port channel, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN, enter the keyword **vlan** followed by a number from 1 to 4094.

When you configure a passive interface, the **show ip ospf interface** command (Figure 318) adds the words "`passive interface`" to indicate that hello packets are not transmitted on that interface.

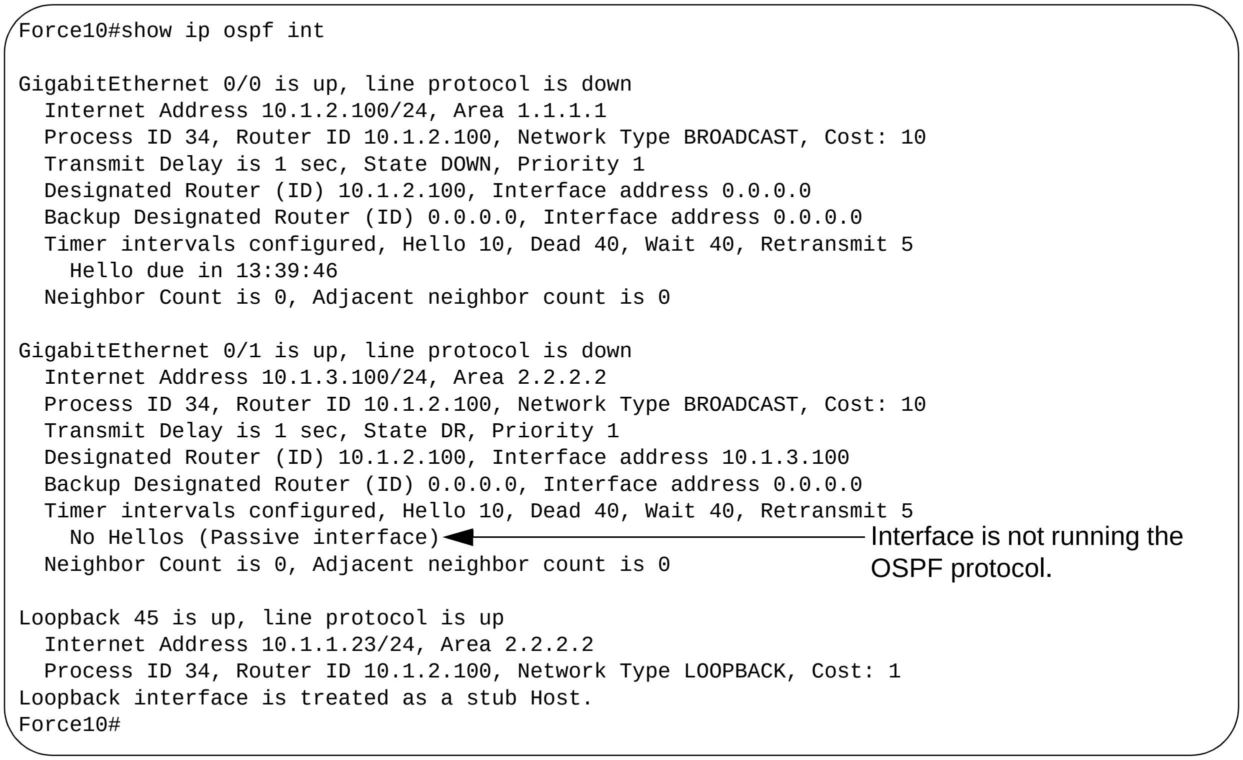

```
Force10#show ip ospf int

GigabitEthernet 0/0 is up, line protocol is down
  Internet Address 10.1.2.100/24, Area 1.1.1.1
  Process ID 34, Router ID 10.1.2.100, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State DOWN, Priority 1
  Designated Router (ID) 10.1.2.100, Interface address 0.0.0.0
  Backup Designated Router (ID) 0.0.0.0, Interface address 0.0.0.0
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 13:39:46
  Neighbor Count is 0, Adjacent neighbor count is 0

GigabitEthernet 0/1 is up, line protocol is down
  Internet Address 10.1.3.100/24, Area 2.2.2.2
  Process ID 34, Router ID 10.1.2.100, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 10.1.2.100, Interface address 10.1.3.100
  Backup Designated Router (ID) 0.0.0.0, Interface address 0.0.0.0
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    No Hellos (Passive interface)
  Neighbor Count is 0, Adjacent neighbor count is 0

Loopback 45 is up, line protocol is up
  Internet Address 10.1.1.23/24, Area 2.2.2.2
  Process ID 34, Router ID 10.1.2.100, Network Type LOOPBACK, Cost: 1
Loopback interface is treated as a stub Host.
Force10#
```

Interface is not running the OSPF protocol.

Figure 318 show ip ospf interface Command Example

In FTOS, you can modify the OSPF settings on the interfaces. Some interface parameter values must be consistent across all interfaces or routing errors will occur. For example, you must set the same time interval for the hello packets on all routers in the OSPF network to prevent misconfiguration of OSPF neighbors.

To change OSPF parameters on the interfaces, use any or all of the following commands in the INTERFACE mode:

Command Syntax	Command Mode	Usage
ip ospf cost *cost*	INTERFACE	Change the cost associated with OSPF traffic on the interface. Configure a *cost* from 1 to 65535 (default depends on the interface speed).
ip ospf dead-interval *seconds*	INTERFACE	Change the time interval the router waits before declaring a neighbor dead. Configure the number of *seconds* from 1 to 65535 (default is 40 seconds). The dead interval must be four times the hello interval. The dead interval must be the same on all routers in the OSPF network.
ip ospf hello-interval *seconds*	INTERFACE	Change the time interval between hello-packet transmission. Configure the number of *seconds* from 1 to 65535 (the default is 10 seconds). The hello interval must be the same on all routers in the OSPF network.
ip ospf message-digest-key *keyid* **md5** *key*	INTERFACE	Use the MD5 algorithm to produce a message digest or key, which is sent instead of the key. Configure the following parameters: • *keyid* range: 1 to 255 • *key*: a character string You cannot learn the key once it is configured. You must be careful when changing this key. For more information on this command, refer to
ip ospf priority *number*	INTERFACE	Change the priority of the interface, which is used to determine the Designated Router for the OSPF broadcast network. Configure the *number* from 0 to 255 (the default is 1).
ip ospf retransmit-interval *seconds*	INTERFACE	Change the retransmission interval between LSAs. Configure the number of *seconds* from 1 to 65535 (the default is 5 seconds). The retransmit interval must be the same on all routers in the OSPF network.

Command Syntax	Command Mode	Usage
ip ospf transmit-delay *seconds*	INTERFACE	Change the wait period between link state update packets sent out the interface. Configure the number of *seconds* between 1 and 65535 (the default is 1 second). The transmit delay must be the same on all routers in the OSPF network.

To view interface configurations, use the **show config** command in the INTERFACE mode (Figure 319). To view interface status in the OSPF process, use the **show ip ospf interface** command in the EXEC mode (Figure 319).

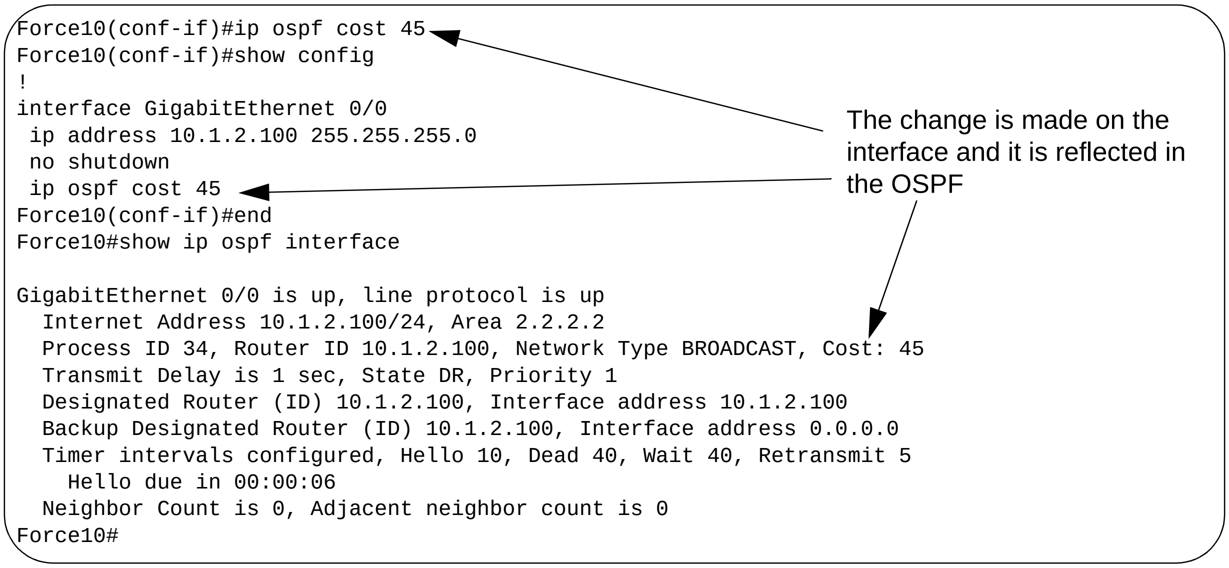

```
Force10(conf-if)#ip ospf cost 45
Force10(conf-if)#show config
!
interface GigabitEthernet 0/0
 ip address 10.1.2.100 255.255.255.0
 no shutdown
 ip ospf cost 45
Force10(conf-if)#end
Force10#show ip ospf interface

GigabitEthernet 0/0 is up, line protocol is up
  Internet Address 10.1.2.100/24, Area 2.2.2.2
  Process ID 34, Router ID 10.1.2.100, Network Type BROADCAST, Cost: 45
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 10.1.2.100, Interface address 10.1.2.100
  Backup Designated Router (ID) 10.1.2.100, Interface address 0.0.0.0
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:06
  Neighbor Count is 0, Adjacent neighbor count is 0
Force10#
```

The change is made on the interface and it is reflected in the OSPF

Figure 319 Changing the OSPF Cost Value on an Interface Example

enable OSPF authentication

You can also enable or change various OSPF authentication parameters. To do so, use the following commands in INTERFACE mode:

Command Syntax	Command Mode	Usage
ip ospf authentication-key *key*	INTERFACE	Set clear text authentication scheme on the interface. Configure a *key* that is a text string no longer than eight characters. All neighboring routers must share the same password to exchange OSPF information.

Command Syntax	Command Mode	Usage
ip ospf auth-change-wait-time *seconds*	INTERFACE	Set the authentication change wait time in *seconds* between 0 and 300 for the interface. This is the amount of time OSPF has available to change its interface authentication type. During the auth-change-wait-time, OSPF sends out packets with both the new and old authentication schemes. This transmission stops when the period ends. The default is 0 seconds.

enable graceful restart

Use this feature to configure OSPF graceful restart. This feature enables you to set up an OSPF router to stay on a forwarding path during both planned and unplanned restarts.

During OSPF graceful restart, OSPF advertises Link-scope Opaque LSA (Grace LSA). Before the restart process commences, the restarting router sends Grace LSA to its neighbors (the helper routers) to request that they cooperate in the restart process.

The Force10 Networks implementation of OSPF graceful restart enables you to specify:

- **grace period**—the length of time the graceful restart process can last before OSPF terminates it.
- **helper-reject neighbors**—the router ID of each restart router that does not receive assistance from the configured router.
- **mode**—the situation or situations that that trigger a graceful restart.
- **role**—the role or roles the configured router can perform.

Note: By default, OSPF graceful restart is disabled.

You enable OSPF graceful restart in OSPF configuration mode. The table below shows the command and its available options:

Command Syntax	Command Mode	Usage
graceful-restart grace-period seconds	ROUTER OSPF	Use this command to enable OSPF graceful-restart. To do so, enter the command followed by the number of *seconds* between 40 and 3000 that this OSPF router’s neighbors will advertise it as fully adjacent, regardless of the synchronization state, during a graceful restart. OSPF terminates this process when the grace period ends.
graceful-restart helper-reject *router-id*	ROUTER OSPF	Enter the router ID of the OSPF helper router from which the router does not accept graceful restart assistance.

Command Syntax	Command Mode	Usage
graceful-restart mode [**planned-only | unplanned-only**]	ROUTER OSPF	Specify the operating mode or modes in which graceful-restart functions. FTOS supports the following options: • Planned-only. The OSPF router supports graceful-restart for planned restarts only. A planned restart is when the user manually enters a fail-over command to force the primary RPM over to the secondary RPM. During a planned restart, OSPF sends out a Grace LSA before the E-Series switches over to the secondary RPM. OSPF also is notified that a planned restart is happening. • Unplanned-only. The OSPF router supports graceful-restart for only unplanned restarts. During an unplanned restart, OSPF sends out a Grace LSA once the secondary RPM comes online. By default, OSPF supports both planned and unplanned restarts.
graceful-restart role [**helper-only** | **restart-only**]	ROUTER OSPF	Configure the graceful restart role or roles that this OSPF router performs. FTOS supports the following options: • Helper-only. The OSPF router supports graceful-restart only as a helper router. • Restart-only. The OSPF router supports graceful-restart only during unplanned restarts. By default, OSPF supports both roles: as a restarting router and as a helper.

When you configure a graceful restart, the **show run ospf** command (Figure 320) displays information such as the following example for router OSPF 1:

```
Force10#show run ospf
!
router ospf 1
 graceful-restart grace-period 300
 graceful-restart role helper-only
 graceful-restart mode unplanned-only
 graceful-restart helper-reject 10.1.1.1
 graceful-restart helper-reject 20.1.1.1
 network 10.0.2.0/24 area 0
Force10#
```

Figure 320 show run ospf Command Example

To disable OSPF graceful-restart after you have enabled it, use the following command:

Command Syntax	Command Mode	Usage
no graceful-restart grace-period	ROUTER OSPF	Disable OSPF graceful-restart. Returns OSPF graceful-restart to its default state.

For more information on OSPF graceful restart, refer to the *FTOS Command Line Interface Reference*.

configure virtual links

Areas within OSPF must be connected to the backbone area (Area ID 0.0.0.0), and if the OSPF area does not have a direct connection to the backbone, at least one virtual link is required. Virtual links must be configured on an ABR connected to the backbone.

To configure virtual links, use the following command in the ROUTER OSPF mode:

Command Syntax	Command Mode	Usage
area *area-id* **virtual-link** *router-id* [**hello-interval** *seconds* | **retransmit-interval** *seconds* | **transmit-delay** *seconds* | **dead-interval** *seconds* | **authentication-key** *key* | **message-digest-key** *keyid* **md5** *key*]	ROUTER OSPF	Configure the optional parameters of a virtual link: • hello-interval • retransmit-interval • dead-interval • authentication-key • message-digest-key

To view the virtual link, use the **show ip ospf virtual-links** command (Figure 321) in the EXEC mode:

```
Force10#show ip ospf virtual-links

Virtual Link to router 192.168.253.5 is up
    Run as demand circuit
    Transit area 0.0.0.1, via interface GigabitEthernet 13/16, Cost of using 2
    Transmit Delay is 1 sec, State POINT_TO_POINT,
        Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in 00:00:02
Force10#
```

Figure 321 show ip ospf virtual-links Command Example

filter routes

To filter routes, use prefix lists. OSPF applies prefix lists to incoming or outgoing routes. Incoming routes must meet the conditions of the prefix lists, and if they do not, OSPF does not add the route to the routing table. Configure the prefix list in the PREFIX LIST mode prior to assigning it to the OSPF process.

For configuration information on prefix lists, refer to Chapter 17, IP Access Control Lists, Prefix Lists, and Route-maps, on page 301.

To apply prefix lists to incoming or outgoing OSPF routes, use the following commands in the ROUTER OSPF mode:

Command Syntax	Command Mode	Usage
distribute-list *prefix-list-name* **in** [*interface*]	ROUTER OSPF	Apply a configured prefix list to incoming OSPF routes.
distribute-list *prefix-list-name* **out** [**connected** | **isis** | **rip** | **static**]	ROUTER OSPF	Assign a configured prefix list to outgoing OSPF routes.

redistribute routes

You can add routes from other routing instances or protocols to the OSPF process. With the **redistribute** command syntax, you can include RIP, static, or directly connected routes in the OSPF process.

To redistribute routes, use the following command in the ROUTER OSPF mode:

Command Syntax	Command Mode	Usage
redistribute {**connected** | **rip** | **static**} [**metric** *metric-value* | **metric-type** *type-value*] [**route-map** *map-name*] [**tag** *tag-value*]	ROUTER OSPF	Specify which routes will be redistributed into OSPF process. Configure the following required and optional parameters: • **connected**, **rip**, or **static**: enter one of the keyword to redistribute those routes. **rip** is supported only on E-Series. • **metric** *metric-value* range: 0 to 4294967295. • **metric-type** *metric-type*: 1 for OSPF external route type 1 or 2 for OSPF external route type 2. • **route-map** *map-name*: enter a name of a configured route map. • **tag** *tag-value* range: 0 to 4294967295.

(Figure 322)

```
Force10(conf-router_ospf)#show config
!
router ospf 34
 network 10.1.2.32 0.0.0.255 area 2.2.2.2
 network 10.1.3.24 0.0.0.255 area 3.3.3.3
 distribute-list dilling in
Force10(conf-router_ospf)#
```

Figure 322 show config Command Example in ROUTER OSPF mode

troubleshooting OSPF

When a routing problem occurs in the OSPF process, use the **show ip route summary** and the **show ip ospf database** commands in EXEC privilege mode to examine the routes. Other options include the **show ip ospf neighbor** and **debug ip ospf** commands.

To view the OSPF configuration for a neighboring router, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Usage
show ip ospf neighbor	EXEC privilege	View the configuration of OSPF neighbors.

To configure the debugging options of the OSPF process, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Usage
debug ip ospf [**event** | **packet** | **spf**]	EXEC privilege	View debug messages. To view all debug message, enter **debug ip ospf**. To view debug messages for a specific operation, enter one of the optional keywords: • **event**: view OSPF event messages • **packet**: view OSPF packet information. • **spf**: view shortest path first (spf) information.

Chapter 30 IS-IS

C-Series	NO
E-Series	✓

Platform Specific Feature: IS-IS is supported on E-Series only.

Intermediate System to Intermediate System (IS-IS) protocol is an interior gateway protocol (IGP) that uses a shortest-path-first algorithm. Force10's implementation of the IPv4 IS-IS is detailed in this chapter.

This chapter covers the following topics:

Protocol Overview

The intermediate system to intermediate system (IS-IS) protocol, developed by the International Organization for Standardization (ISO), is an interior gateway protocol (IGP) that uses a shortest-path-first algorithm. This protocol supports routers passing both IP and OSI traffic, though the Force10 Networks implementation only supports IP traffic.

IS-IS is organized hierarchally into routing domains, and each router or system resides in at least one area. In IS-IS, routers are designated as Level 1, Level 2 or Level 1-2 systems. Level 1 routers only route traffic within an area, while Level 2 routers route traffic between areas. At its most basic, Level 1 systems route traffic within the area and any traffic destined for outside the area is sent to a Level 1-2 system. Level 2 systems manage destination paths for external routers. Only Level 2 routers can exchange data packets or routing information directly with external routers located outside of the routing domains. Level 1-2 systems manage both inter-area and intra-area traffic by maintaining two separate link databases; one for Level 1 routes and one for Level 2 routes. A Level 1-2 router does not advertise Level 2 routes to a Level 1 router.

To establish adjacencies, each IS-IS router sends different Protocol Data Units (PDU). For IP traffic, the IP addressing information is included in the IS-IS hello PDUs and the Link State PDUs (LSPs).

This brief overview is not intended to provide a complete understanding of IS-IS; for that, consult the documents listed in IS-IS Standards on page 490.

IS-IS Addressing

IS-IS PDUs require ISO-style addressing called Network Entity Title (NET). For those familiar with NSAP addresses, the composition of the NET is identical to an NSAP address, except the last byte is always 0. The NET is composed of IS-IS area address, system ID, and the N-selector. The last byte is the N-selector. All routers within an area have the same area portion. Level 1 routers route based on the system address portion of the address, while the Level 2 routers route based on the area address.

The NET length is variable, with a maximum of 20 bytes and a minimum of 8 bytes. It is composed of the following:

- area address. Within your routing domain or area, each area must have a unique area value. The first byte is called the authority and format indicator (AFI).
- system address. This is usually the router’s MAC address.
- N-selector. This is always 0.

Figure 323 is an example of the ISO-style address to illustrate the address format used by IS-IS. In this example, the first five bytes (47.0005.0001) are the area address. The system portion is 000c.000a.4321 and the last byte is always 0.

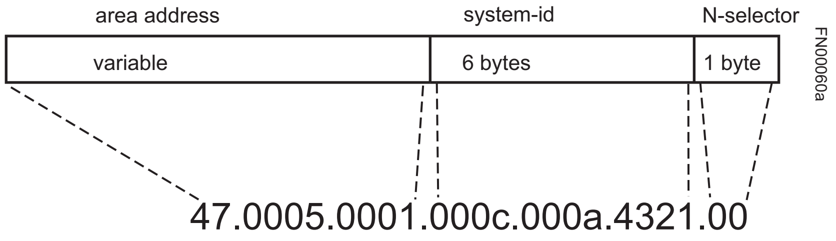

Figure 323 ISO Address Format

IS-IS Standards

The IS-IS protocol is defined in the following documents:

- ISO/IEC 10589, *Information Technology—Telecommunication and information exchange between systems—Intermediate system to Intermediate system intradomain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service* (ISO 8473)
- RFC 1142, *OSI IS-IS Intra-Domain Routing Protocol* (This is an ASCII version of ISO/IEC 10589)
- RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*

- RFC 2763, *Dynamic Hostname Exchange Mechanism for IS-IS*
- RFC 2966, *Domain-wide Prefix Distribution with Two-Level IS-IS*
- RFC 3373, Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies

Implementation Information

The E-Series's implementation of IS-IS is based on RFC 1195 and supports one instance of IS-IS and six areas. The E-Series can be configured as a Level 1 router, a Level 2 router, or a Level 1-2 router.

By default, FTOS supports dynamic hostname exchange to assist with troubleshooting and configuration. By assigning a name to an IS-IS NET address, you can track IS-IS information on that address easier. FTOS does not support ISO CLNS routing, however, the ISO NET format is supported for addressing.

Table 43 displays the default values for IS-IS.

Table 43 E-Series IS-IS Default Values

IS-IS Parameter	Default Value
Complete Sequence Number PDU (CSNP) interval	10 seconds
IS-to-IS hello PDU interval	10 seconds
IS-IS interface metric	10
Metric style	Narrow
Designated Router priority	64
Circuit Type	Level 1 and Level 2
IS Type	Level 1 and Level 2
Equal Cost Multi Paths	16

Configuration Information

To configure IS-IS, you must enable IS-IS in two modes: ROUTER ISIS and INTERFACE. Commands in ROUTER ISIS mode configure IS-IS globally on the E-Series, while commands executed in the INTERFACE mode enable and configure IS-IS features on that interface only.

Configuration Task List for IS-IS

The following list includes the configuration tasks for IS-IS:

- enable IS-IS on page 492 (mandatory)
- configure IS-IS interface parameters on page 495 (mandatory)
- change LSP attributes on page 496 (optional)

- configure IS-IS metric style and cost on page 497 (optional)
- change the is-type on page 499 (optional)
- control routing updates on page 500 (optional)
- configure authentication passwords on page 502 (optional)
- set the overload bit on page 503 (optional)
- debug IS-IS on page 504 (optional)

For a complete listing of all commands related to IS-IS, refer to .

enable IS-IS

By default, IS-IS is not enabled.

You can create one instance of IS-IS on the E-Series. To enable IS-IS globally on the E-Series, you must create an IS-IS routing process and assign a NET address. To exchange protocol information with neighbors, enable IS-IS on an interface, instead of on a network as with other routing protocols.

In IS-IS, neighbors form adjacencies only when they are same IS type. For example, a Level 1 router never forms an adjacency with a Level 2 router. A Level 1-2 router will form Level 1 adjacencies with a neighboring Level 1 router and will form Level 2 adjacencies with a neighboring Level 2 router.

To configure IS-IS globally on the E-Series, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command	Command Mode	Purpose
1	**router isis** [*tag*]	CONFIGURATION	Create an IS-IS routing process. • *tag* is optional and identifies the name of the IS-IS process.
2	**net** *network-entity-title*	ROUTER ISIS	Configure an IS-IS network entity title (NET) for a routing process. Specify the area address and system ID for an IS-IS routing process. The last byte must be 00. Refer to for more information on configuring a NET.

Step	Command	Command Mode	Purpose
3	**interface** *interface*	CONFIGURATION	Enter the interface configuration mode. Enter the keyword **interface** followed by the type of interface and slot/port information: • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For the Loopback interface on the RPM, enter the keyword **loopback** followed by a number from 0 to 16383. • For a port channel, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN, enter the keyword **vlan** followed by a number from 1 to 4094.
4	**ip address** *ip-address mask*	INTERFACE	Assign an IP address and mask to the interface. The IP address must be on the same subnet as other IS-IS neighbors, but the IP address does not need to relate to the NET address.
5	**ip router isis** [*tag*]	INTERFACE	Enable IS-IS on the interface. If you configure a *tag* variable, it must be the same as the *tag* variable assigned in step 1.

The default IS type is level-1-2. To change the IS type to Level 1 only or Level 2 only, use the **is-type** command in the ROUTER ISIS mode.

To view the IS-IS configuration, enter the **show isis protocol** command in the EXEC privilege mode or the **show config** command in the ROUTER ISIS mode.

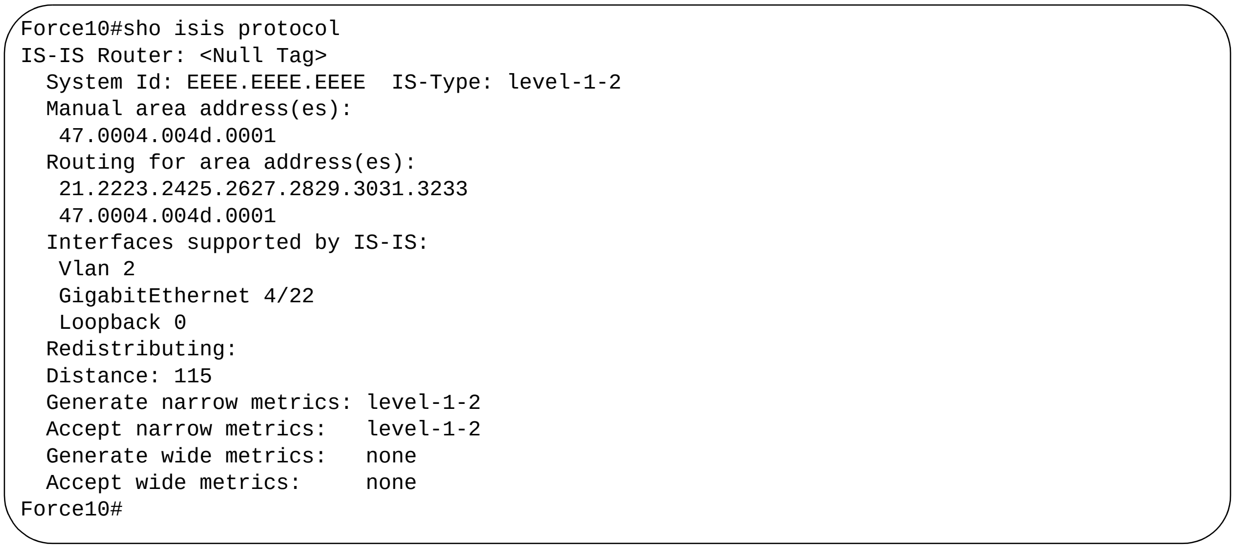

```
Force10#sho isis protocol
IS-IS Router: <Null Tag>
  System Id: EEEE.EEEE.EEEE  IS-Type: level-1-2
  Manual area address(es):
   47.0004.004d.0001
  Routing for area address(es):
   21.2223.2425.2627.2829.3031.3233
   47.0004.004d.0001
  Interfaces supported by IS-IS:
   Vlan 2
   GigabitEthernet 4/22
   Loopback 0
  Redistributing:
  Distance: 115
  Generate narrow metrics: level-1-2
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   none
  Accept wide metrics:     none
Force10#
```

Figure 324 show isis protocol Command Example

To view IS-IS protocol statistics, use the **show isis traffic** command in the EXEC privilege mode (Figure 325).

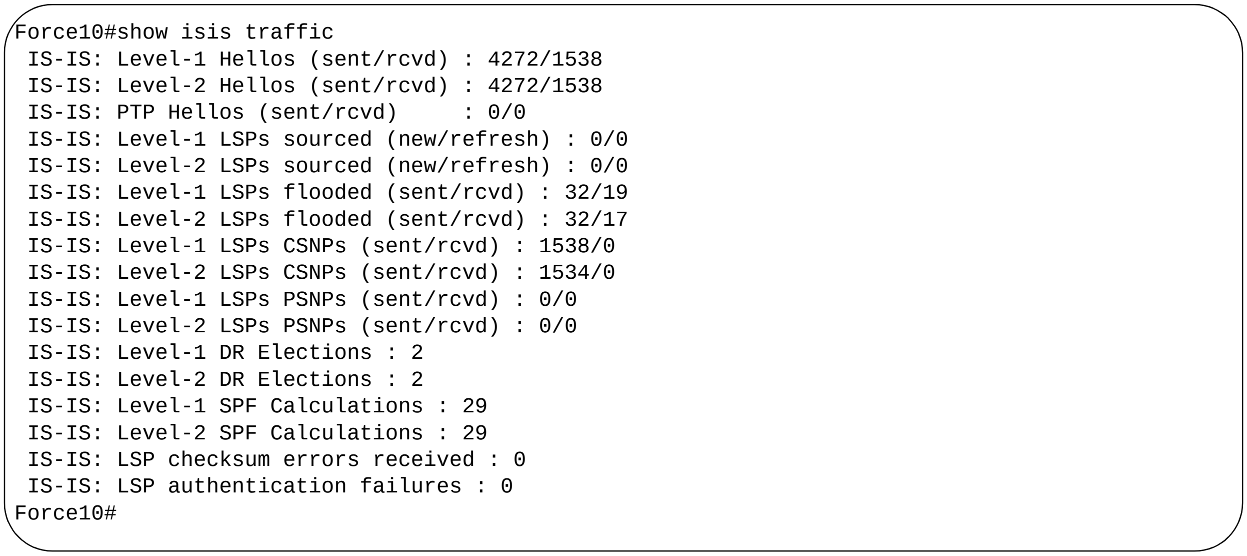

```
Force10#show isis traffic
 IS-IS: Level-1 Hellos (sent/rcvd) : 4272/1538
 IS-IS: Level-2 Hellos (sent/rcvd) : 4272/1538
 IS-IS: PTP Hellos (sent/rcvd)     : 0/0
 IS-IS: Level-1 LSPs sourced (new/refresh) : 0/0
 IS-IS: Level-2 LSPs sourced (new/refresh) : 0/0
 IS-IS: Level-1 LSPs flooded (sent/rcvd) : 32/19
 IS-IS: Level-2 LSPs flooded (sent/rcvd) : 32/17
 IS-IS: Level-1 LSPs CSNPs (sent/rcvd) : 1538/0
 IS-IS: Level-2 LSPs CSNPs (sent/rcvd) : 1534/0
 IS-IS: Level-1 LSPs PSNPs (sent/rcvd) : 0/0
 IS-IS: Level-2 LSPs PSNPs (sent/rcvd) : 0/0
 IS-IS: Level-1 DR Elections : 2
 IS-IS: Level-2 DR Elections : 2
 IS-IS: Level-1 SPF Calculations : 29
 IS-IS: Level-2 SPF Calculations : 29
 IS-IS: LSP checksum errors received : 0
 IS-IS: LSP authentication failures : 0
Force10#
```

Figure 325 show isis traffic Command Example

You can assign additional NET addresses, but the System ID portion of the NET address must remain the same. FTOS supports up to six area addresses.

Some address considerations are:

- In order to be neighbors, Level 1 routers must be configured with at least one common area address.

- A Level 2 router becomes a neighbor with another Level 2 router regardless of the area address configured. However, if the area addresses are different, the link between the Level 2 routers is only at Level 2.

To view the configuration of the interface, use the **show config** command in the INTERFACE mode.

To view all interfaces configured with IS-IS routing and their defaults, use the **show isis interface** command in the EXEC privilege mode (Figure 326).

```
Force10#show isis inter
GigabitEthernet 4/22 is up, line protocol is up
  MTU 1551, Encapsulation SAP
  Routing Protocol: IS-IS
    Circuit Type: Level-1-2
    Interface Index 179929088, Local circuit ID 2
    Level-1 Metric: 10, Priority: 64, Circuit ID: eljefe.02
            Hello Interval: 10, Hello Multiplier: 3, CSNP Interval: 10
    Number of active level-1 adjacencies: 1
    Level-2 Metric: 10, Priority: 64, Circuit ID: eljefe.02
            Hello Interval: 10, Hello Multiplier: 3, CSNP Interval: 10
    Number of active level-2 adjacencies: 1
    Next IS-IS LAN Level-1 Hello in 3 seconds
    Next IS-IS LAN Level-2 Hello in 2 seconds
    LSP Interval: 33
Force10#
```

Figure 326 show isis interface Command Example

configure IS-IS interface parameters

You must enable the IS-IS process on an interface for the IS-IS process to exchange protocol information and form adjacencies. You can modify IS-IS parameters on a per-interface basis, but it is not necessary.

To change IS-IS defaults on an interface, use any or all of the following commands in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
isis circuit-type {**level-1** | **level-1-2** | **level-2-only**}	INTERFACE	Configure the circuit type for the interface. Default is level-1-2.
isis csnp-interval *seconds* [**level-1** | **level-2**]	INTERFACE	Configure the complete sequence number PDU (CSNP) interval. • *seconds* range: 0 to 65535. Default is 10 seconds. Default level is level-1.
isis hello-interval *seconds* [**level-1** | **level-2**]	INTERFACE	Specify the length of time between hello packets sent by FTOS. • *seconds* range: 0 to 65535. Default is 10 seconds. Default level is level-1.

Command Syntax	Command Mode	Purpose
isis hello-multiplier *multiplier* [**level-1** | **level-2**]	INTERFACE	Specify the number of IS-IS hello packets a neighbor must miss before the router declares the adjacency as down. • *multiplier* range: 3 to 1000. Default is 3. Default level is level-1.
isis metric *default-metric* [**level-1** | **level-2**]	INTERFACE	Assign a metric for a link or interface. • *default-metric* range: 0 to 63 for narrow and transition metric styles; 0 to 16777215 for wide metric styles. Default is 10. Default level is level-1. Refer to for more information on this command.
isis password [**hmac-md5**] *password* [**level-1** | **level-2**]	INTERFACE	Configure the password to authenticate between IS-IS neighbors. Simple HMAC-MD5 authentication is supported. • *password:* a text string Default level is level-1. The password must be the same on all neighbors to form adjacencies.
isis priority *value* [**level-1** | **level-2**]	INTERFACE	Set the priority for Designated Router election on the interface. • *value* range: 0 to 127. Default is 64. Default level is level-1.

To view the interface’s non-default configuration, use the **show config** command in the INTERFACE mode.

To view all interfaces routing IS-IS, use the **show isis interface** command in the EXEC privilege mode (Figure 326).

change LSP attributes

IS-IS routers flood Link state PDUs (LSPs) to exchange routing information. LSP attributes include the generation interval, maximum transmission unit (MTU) or size, and the refresh interval. You can modify the LSP attribute defaults, but it is not necessary.

To change the defaults, use any or all of the following commands in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
lsp-gen-interval [**level-1** | **level-2**] *seconds*	ROUTER ISIS	Set interval between LSP generation. • *seconds* range: 0 to 120 Default is 5 seconds. Default level is Level 1.
lsp-mtu *size*	ROUTER ISIS	Set the LSP size. • *size* range: 128 to 9195. Default is 1497.
lsp-refresh-interval *seconds*	ROUTER ISIS	Set the LSP refresh interval. • *seconds* range: 1 to 65535. Default is 900 seconds.
max-lsp-lifetime *seconds*	ROUTER ISIS	Set the maximum time LSPs lifetime. • *seconds* range: 1 to 65535 Default is 1200 seconds.

To view the configuration, use the **show config** command in the ROUTER ISIS mode or the **show running-config isis** command in the EXEC privilege mode (Figure 327).

```
Force10#show running-config isis
!
router isis
 lsp-refresh-interval 902
 net 47.0005.0001.000C.000A.4321.00
 net 51.0005.0001.000C.000A.4321.00
Force10#
```

Figure 327 show running-config isis Command Example

configure IS-IS metric style and cost

All IS-IS links or interfaces are associated with a cost that is used in the SPF calculations. The possible cost varies depending on the metric style supported. If you configure narrow, transition or narrow transition metric style, the cost can be a number between 0 and 63. If you configure wide or wide transition metric style, the cost can be a number between 0 and 16,777,215. FTOS supports five different metric styles: narrow, wide, transition, narrow transition, and wide transition.

By default, FTOS generates and receives narrow metric values. Metrics or costs higher than 63 are not supported. To accept or generate routes with a higher metric, you must change the metric style of the IS-IS process. For example, if metric is configured as narrow, and an LSP with wide metrics is received, the route is not installed.

FTOS supports the following IS-IS metric styles:

Table 44 Metric Styles

Metric Style	Characteristics	Cost Range Supported on IS-IS Interfaces
narrow	Sends and accepts narrow or old TLVs (Type Length Value).	0 to 63
wide	Sends and accepts wide or new TLVs	0 to 16777215
transition	Sends both wide (new) and narrow (old) TLVs.	0 to 63
narrow transition	Sends narrow (old) TLVs and accepts both narrow (old) and wide (new) TLVs	0 to 63
wide transition	Sends wide (new) TLVs and accepts both narrow (old) and wide (new) TLVs.	0 to 16777215

To change the IS-IS metric style of the IS-IS process, use the following command in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
metric-style {**narrow** [**transition**] | **transition** | **wide** [**transition**]} [**level-1** | **level-2**]	ROUTER ISIS	Set the metric style for the IS-IS process. Default: narrow Default: Level 1 and Level 2 (level-1-2)

To view which metric types are generated and received, use the **show isis protocol** command (Figure 324) in the EXEC privilege mode.

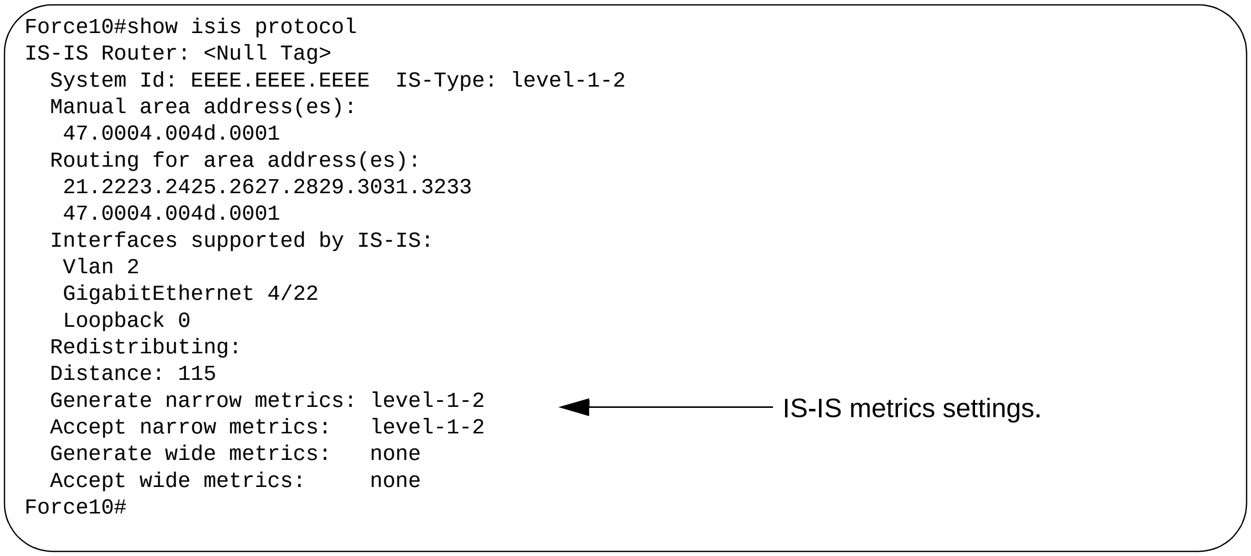

```
Force10#show isis protocol
IS-IS Router: <Null Tag>
  System Id: EEEE.EEEE.EEEE  IS-Type: level-1-2
  Manual area address(es):
   47.0004.004d.0001
  Routing for area address(es):
   21.2223.2425.2627.2829.3031.3233
   47.0004.004d.0001
  Interfaces supported by IS-IS:
   Vlan 2
   GigabitEthernet 4/22
   Loopback 0
  Redistributing:
  Distance: 115
  Generate narrow metrics: level-1-2
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   none
  Accept wide metrics:     none
Force10#
```

IS-IS metrics settings.

Figure 328 show isis protocol Command Example

When you change from one IS-IS metric style to another, the IS-IS metric value could be affected. For each interface with IS-IS enabled, you can assign a cost or metric that is used in the link state calculation. Appendix D, contains details on the behavior of the metric value when you change the metric style.

To change the metric or cost of the interface, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
isis metric *default-metric* [**level-1** | **level-2**]	INTERFACE	*default-value* range: 0 to 63 if the metric-style is narrow, narrow-transition or transition. 0 to 16777215 if the metric style is wide or wide transition. Default: 10.

To view the interface's current metric, use the **show config** command in the INTERFACE mode or the **show isis interface** command in the EXEC privilege mode.

Note: In FTOS, the CLI help always shows the value range (0-16777215) for the metric style. See Table 45 for the correct value range.

Table 45 Correct Value Range for the isis metric command

Metric Style	Correct Value Range
wide	0 to 16777215
narrow	0 to 63
wide transition	0 to 16777215
narrow transition	0 to 63
transition	0 to 63

change the is-type

You can configure the E-Series system to act as one of the following:

- Level 1 router
- Level 1-2 router
- Level 2 router

To change the is-type for the router, use the following command in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
is-type {**level-1** | **level-1-2** | **level-2**}	ROUTER ISIS	Change the is-type for the IS-IS process.

To view which is-type is configured, use the **show isis protocol** command in the EXEC privilege mode (Figure 324). The **show config** command in the ROUTER ISIS mode displays only nondefault information, so if you do not change the is-type, the default value (level-1-2) is not displayed.

The default is Level 1-2 router. When the is-type is Level 1-2, the software maintains two Link State databases, one for each level. Use the **show isis database** command to view the Link State databases (Figure 329).

```
Force10#show isis database
IS-IS Level-1 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
B233.00-00            0x00000003   0x07BF        1088              0/0/0
eljefe.00-00        * 0x00000009   0xF76A        1126              0/0/0
eljefe.01-00        * 0x00000001   0x68DF        1122              0/0/0
eljefe.02-00        * 0x00000001   0x2E7F        1113              0/0/0
Force10.00-00         0x00000002   0xD1A7        1102              0/0/0
IS-IS Level-2 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
B233.00-00            0x00000006   0xC38A        1124              0/0/0
eljefe.00-00        * 0x0000000D   0x51C6        1129              0/0/0
eljefe.01-00        * 0x00000001   0x68DF        1122              0/0/0
eljefe.02-00        * 0x00000001   0x2E7F        1113              0/0/0
Force10.00-00         0x00000004   0xCDA9        1107              0/0/0

Force10#
```

Figure 329 show isis database Command Example

control routing updates

To control the source of IS-IS route information, use the following commands, in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
passive-interface *interface*	ROUTER ISIS	Disable a specific interface from sending or receiving IS-IS routing information. Enter the type of interface and slot/port information: • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For the Loopback interface on the RPM, enter the keyword **loopback** followed by a number from 0 to 16383. • For a port channel, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN, enter the keyword **vlan** followed by a number from 1 to 4094.

Another method of controlling routing information is to filter the information through a prefix list. Prefix lists are applied to incoming or outgoing routes and routes must meet the conditions of the prefix lists or FTOS does not install the route in the routing table. The prefix lists are globally applied on all interfaces running IS-IS. Configure the prefix list in the PREFIX LIST mode prior to assigning it to the IS-IS process.

For configuration information on prefix lists, see Chapter 17, IP Access Control Lists, Prefix Lists, and Route-maps, on page 301.

To apply prefix lists to incoming or outgoing routes, use the following commands in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
distribute-list *prefix-list-name* **in** [*interface*]	ROUTER ISIS	Apply a configured prefix list to all incoming IS-IS routes. Enter the type of interface and slot/port information: • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For the Loopback interface on the RPM, enter the keyword **loopback** followed by a number from 0 to 16383. • For a port channel, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN, enter the keyword **vlan** followed by a number from 1 to 4094.
distribute-list *prefix-list-name* **out** [**connected** | **ospf** *process-id* | **rip** | **static**]	ROUTER ISIS	Apply a configured prefix list to all outgoing IS-IS routes. You can configure one of the optional parameters: • **connected:** for directly connected routes. • **ospf** *process-id:* for OSPF routes only. • **rip:** for RIP routes only. • **static:** for user-configured routes.

In addition to filtering routes, you can add routes from other routing instances or protocols to the IS-IS process. With the **redistribute** command syntax, you can include BGP, OSPF, RIP, static, or directly connected routes in the IS-IS process.

To add routes from other routing instances or protocols, use any of the following commands in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
redistribute {**connected** | **rip** | **static**} [**level-1** **level-1-2** | **level-2**] [**metric** *metric-value*] [**metric-type** {**external** | **internal**}] [**route-map** *map-name*]	ROUTER ISIS	Include directly connected, RIP, or user-configured (static) routes in IS-IS. Configure the following parameters: • **level-1**, **level-1-2**, or **level-2**: Assign all redistributed routes to a level. Default is **level-2**. • *metric* range: 0 to 16777215. Default is 0. • **metric-type:** choose either **external** or **internal**. Default is internal. • *map-name*: name of a configured route map.
redistribute ospf *process-id* [**level-1** | **level-1-2** | **level-2**] [**metric** *value*] [**match external** {**1** | **2**} | **match internal**] [**metric-type** {**external** | **internal**}] [**route-map** *map-name*]	ROUTER ISIS	Include specific OSPF routes in IS-IS. Configure the following parameters: • *process-id* range: 1 to 65535 • **level-1**, **level-1-2**, or **level-2**: Assign all redistributed routes to a level. Default is **level-2**. • *metric* range: 0 to 16777215. Default is 0. • **match external** range: 1 or 2 • **match internal** • **metric-type**: external or internal. • *map-name*: name of a configured route map.

Note: Starting with Release 6.3.1, you can also redistribute BGP routes in IS-IS and IS-IS through BGP. See the IS-IS chapter and BGP chapters in the *FTOS Command Line Interface Reference* for details.

To view the current IS-IS configuration, use the **show running-config isis** command in the EXEC privilege mode or the **show config** command in the ROUTER ISIS mode.

configure authentication passwords

You can assign an authentication password for routers in Level 1 and for routers in Level 2. Since Level 1 and Level 2 routers do not communicate with each other, you can assign different passwords for Level 1 routers and for Level 2 routers. If you want the routers in the level to communicate with each other, though, they must be configured with the same password.

To configure a simple text password, use either or both of the commands in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
area-password [**hmac-md5**] *password*	ROUTER ISIS	Configure authentication password for an area. FTOS supports HMAC-MD5 authentication. This password is inserted in Level 1 LSPs, Complete SNPs, and Partial SNPs.
domain-password [*encryption-type* | **hmac-md5**] *password*	ROUTER ISIS	Set the authentication password for a routing domain. FTOS supports both DES and HMAC-MD5 authentication methods. This password is inserted in Level 2 LSPs, Complete SNPs, and Partial SNPs.

To view the passwords, use the **show config** command in the ROUTER ISIS mode or the **show running-config isis** command in the EXEC privilege mode.

To remove a password, use either **no area-password** or **no domain-password** commands in the ROUTER ISIS mode.

set the overload bit

Another use for the overload bit is to prevent other routers from using the E-Series as an intermediate hop in their shortest path first (SPF) calculations. For example, if the IS-IS routing database is out of memory and cannot accept new LSPs, FTOS sets the overload bit and IS-IS traffic continues to transit the E-Series.

To set the overload bit manually, use this command the following command in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
set-overload-bit	ROUTER ISIS	Set the overload bit in LSPs. This prevents other routers from using it as an intermediate hop in their shortest path first (SPF) calculations.

To remove the overload bit, enter **no set-overload-bit**.

To see if the bit is set, a 1 is placed in the OL column in the show isis database command output. In Figure 330, the overload bit is set in both the Level-1 and Level-2 database because the IS type for the router is Level-1-2.

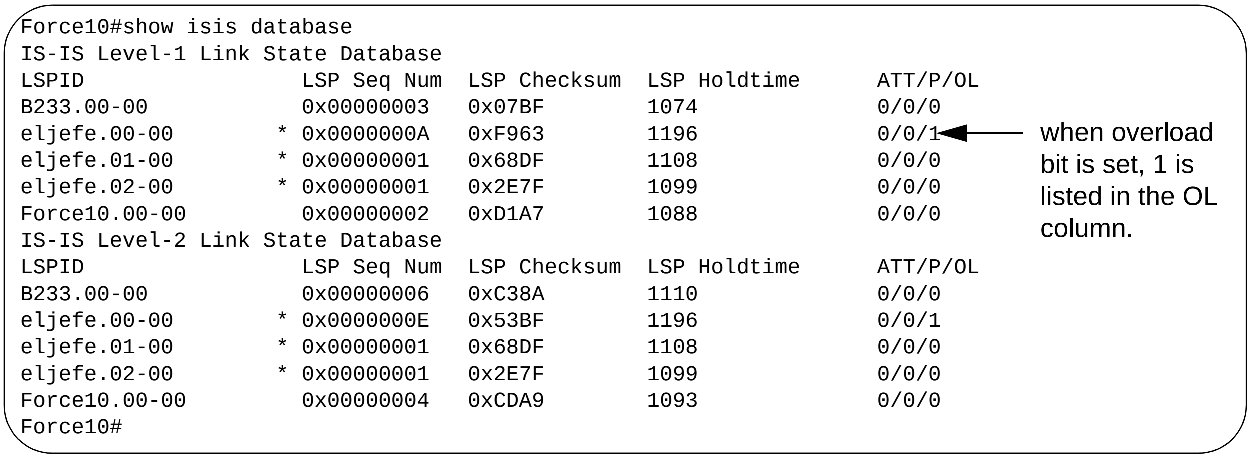

```
Force10#show isis database
IS-IS Level-1 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
B233.00-00            0x00000003   0x07BF        1074              0/0/0
eljefe.00-00        * 0x0000000A   0xF963        1196              0/0/1
eljefe.01-00        * 0x00000001   0x68DF        1108              0/0/0
eljefe.02-00        * 0x00000001   0x2E7F        1099              0/0/0
Force10.00-00         0x00000002   0xD1A7        1088              0/0/0
IS-IS Level-2 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
B233.00-00            0x00000006   0xC38A        1110              0/0/0
eljefe.00-00        * 0x0000000E   0x53BF        1196              0/0/1
eljefe.01-00        * 0x00000001   0x68DF        1108              0/0/0
eljefe.02-00        * 0x00000001   0x2E7F        1099              0/0/0
Force10.00-00         0x00000004   0xCDA9        1093              0/0/0
Force10#
```

when overload bit is set, 1 is listed in the OL column.

Figure 330 show isis database Command Example

debug IS-IS

To debug all IS-IS processes, enter the **debug isis** command in the EXEC privilege mode.

Use the following commands for specific IS-IS debugging:

Command Syntax	Command Mode	Purpose
debug isis adj-packets [*interface*]	EXEC privilege	View information on all adjacency-related activity (for example, hello packets that are sent and received). To view specific information, enter one of the following optional parameters: • *interface:* Enter the type of interface and slot/port information to view IS-IS information on that interface only.
debug isis local-updates [*interface*]	EXEC privilege	View information about IS-IS local update packets. To view specific information, enter one of the following optional parameters: • *interface:* Enter the type of interface and slot/port information to view IS-IS information on that interface only.
debug isis snp-packets [*interface*]	EXEC privilege	View IS-IS SNP packets, include CSNPs and PSNPs. To view specific information, enter one of the following optional parameters: • *interface:* Enter the type of interface and slot/port information to view IS-IS information on that interface only.

Command Syntax	Command Mode	Purpose
debug isis spf-triggers	EXEC privilege	View the events that triggered IS-IS shortest path first (SPF) events for debugging purposes.
debug isis update-packets [*interface*]	EXEC privilege	View sent and received LSPs. To view specific information, enter one of the following optional parameters: • *interface:* Enter the type of interface and slot/port information to view IS-IS information on that interface only.

FTOS displays debug messages on the console. To view which debugging commands are enabled, use the **show debugging** command in the EXEC privilege mode.

To disable a specific debug command, enter the keyword no followed by the debug command. For example, to disable debugging of IS-IS updates, you enter **no debug isis updates-packets** command.

To disable all IS-IS debugging, enter **no debug isis**.

To disable all debugging, enter **undebug all**.

Chapter 31 IPv6 IS-IS

C-Series	NO
E-Series	✓

Platform Specific Feature: IS-IS for IPv6 is supported on E-Series only.

Intermediate System to Intermediate System (IS-IS) protocol is an interior gateway protocol (IGP) that uses a shortest-path-first algorithm. Force10's implementation of the IPv6 IS-IS is detailed in this chapter.

This chapter covers the following topics:

Protocol Overview

The intermediate-system-to-intermediate-system (IS-IS) protocol, developed by the International Organization for Standardization (ISO), is an interior gateway protocol (IGP) that uses a shortest-path-first algorithm.

This protocol supports routers passing both IP and OSI traffic, though the Force10 Networks implementation supports only IP traffic.

IS-IS is organized hierarchally into routing domains, and each router or system resides in at least one area. In IS-IS, routers are designated as Level 1, Level 2 or Level 1-2 systems. Level 1 routers only route traffic within an area, while Level 2 routers route traffic between areas. At its most basic, Level 1 systems route traffic within the area and any traffic destined for outside the area is sent to a Level 1-2 system. Level 2 systems manage destination paths for external routers. Only Level 2 routers can exchange data packets or routing information directly with external routers located outside of the routing domains. Level 1-2 systems manage both inter-area and intra-area traffic by maintaining two separate link databases; one for Level 1 routes and one for Level 2 routes. A Level 1-2 router does not advertise Level 2 routes to a Level 1 router.

To establish adjacencies, each IS-IS router sends different Protocol Data Units (PDU). For IP traffic, the IP addressing information is included in the IS-IS hello PDUs and the Link State PDUs (LSPs).

This brief overview is not intended to provide a complete understanding of IS-IS; for that, consult the documents listed in IS-IS Standards on page 509.

IS-IS Addressing

IS-IS PDUs require ISO-style addressing called Network Entity Title (NET). For those familiar with NSAP addresses, the composition of the NET is identical to an NSAP address, except the last byte is always 0. The NET is composed of IS-IS area address, system ID, and the N-selector. The last byte is the N-selector. All routers within an area have the same area portion. Level 1 routers route based on the system address portion of the address, while the Level 2 routers route based on the area address.

The NET length is variable, with a maximum of 20 bytes and a minimum of 8 bytes. It is composed of the following:

- area address. Within your routing domain or area, each area must have a unique area value. The first byte is called the authority and format indicator (AFI).
- system address. This is usually the router’s MAC address.
- N-selector. This is always 0.

Figure 331 is an example of the ISO-style address to illustrate the address format used by IS-IS. In this example, the first five bytes (47.0005.0001) are the area address. The system portion is 000c.000a.4321 and the last byte is always 0.

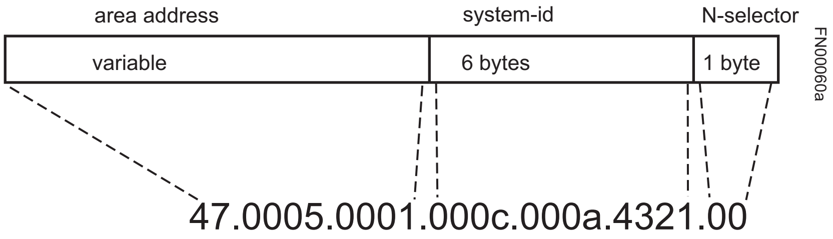

Figure 331 ISO Address Format

IS-IS Standards

The IS-IS protocol is defined in the following documents:

- ISO/IEC 10589, *Information Technology—Telecommunication and information exchange between systems—Intermediate system to Intermediate system intradomain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service* (ISO 8473)
- RFC 1142, *OSI IS-IS Intra-Domain Routing Protocol* (This is an ASCII version of ISO/IEC 10589)
- RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*
- RFC 2763, *Dynamic Hostname Exchange Mechanism for IS-IS*
- RFC 2966, *Domain-wide Prefix Distribution with Two-Level IS-IS*
- RFC 3373, Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies

Implementation Information

The E-Series's implementation of IS-IS is based on RFC 1195 and supports one instance of IS-IS and six areas. The E-Series can be configured as a Level 1 router, a Level 2 router, or a Level 1-2 router. For IPv6, the IPv4 implementation has been expanded to include two new type-length-values (TLV) in the protocol data unit (PDU) that carry information required for IPv6 routing. The new TLVs are *IPv6 Reachability* and *IPv6 Interface Address*. Also, a new IPv6 protocol identifier has also been included in the supported TLVs. The new TLVs use the extended metrics and up/down bit semantics.

By default, FTOS supports dynamic hostname exchange to assist with troubleshooting and configuration. By assigning a name to an IS-IS NET address, you can track IS-IS information on that address easier. FTOS does not support ISO CLNS routing; however, the ISO NET format is supported for addressing.

To support IPv6, the Force10 implementation of IS-IS performs the following tasks:

- Advertise IPv6 information in the PDUs
- Process IPv6 information received in the PDUs
- Compute routes to IPv6 destinations
- Download IPv6 routes to RTM for installing in the FIB
- Accept external IPv6 information and advertise this information in the PDUs

Table 46 displays the default values for IS-IS.

Table 46 E-Series IS-IS Default Values

IS-IS Parameter	Default Value
Complete Sequence Number PDU (CSNP) interval	10 seconds
IS-to-IS hello PDU interval	10 seconds
IS-IS interface metric	10

Table 46 E-Series IS-IS Default Values

IS-IS Parameter	Default Value
Metric style	Narrow
Designated Router priority	64
Circuit Type	Level 1 and Level 2
IS Type	Level 1 and Level 2
Equal Cost Multi Paths	16

Configuration Information

To use IS-IS, you must configure and enable IS-IS in three modes: ROUTER CONFIGURATION, INTERFACE CONFIGURATION, and ADDRESS-FAMILY mode. Commands in ROUTER ISIS mode configure IS-IS globally on the E-Series, while commands executed in the INTERFACE mode enable and configure IS-IS features on that interface only. Commands in the ADDRESS-FAMILY mode are specific to IPv6.

Note that by using the IS-IS routing protocol to exchange IPv6 routing information and to determine destination reachability, you can route IPv6 along with IPv4 while using a single intra-domain routing protocol. The configuration commands allow you to enable and disable IPv6 routing and to configure or remove IPv6 prefixes on links.

The ROUTER CONFIGURATION commands that are applicable to both IPv4 and IPv6 for implementing IS-IS are:

- **area-password**—configures the authentication password for an area
- **clear**—clears IS-IS configuration
- **clns**—assigns a name to a router
- **description**—configuration description
- **domain-password**—sets the authentication password for a routing domain
- **hello**—pads IS-IS hello PDUs to full MTU
- **hostname**—sets the dynamic hostname for IS-IS
- **ignore-lsp-errors**—ignores LSPs with bad checksums
- **is-type**—sets IS Level for this routing process
- **log-adjacency-changes**—logs changes in adjacency states
- **lsp-gen-interval**—sets minimum interval between successive generations of LSP
- **lsp-mtu**—sets maximum LSP size
- **lsp-refresh-interval**—sets LSP refresh interval
- **max-area-addresses**—allows manual configuration of more area addresses
- **max-lsp-lifetime**—sets maximum LSP lifetime
- **net**—sets a network entity title for this process
- **passive-interface**—suppresses routing updates on an interface

- **set-overload-bit**—signals other routers not to use SPF
- **show**—shows the IS-IS configuration
  - **show isis neighbors detail**—for IPv6, Link Local address of neighbor is displayed
  - **show isis database detail**—for IPv6, all IPv6 TLVs are displayed
  - **show isis interface**—for IPv6, IS-IS IPv6 metrics are displayed
- **spf-interval**—sets the minimum interval between SPF calculations

The INTERFACE CONFIGURATION commands that are applicable to both IPv4 and IPv6 for implementing IS-IS are:

- **circuit-type**—configures the circuit type for interface
- **csnp-interval**—sets CSNP interval in seconds
- **hello**—adds padding to the IS-IS hello packets
- **hello-interval**—sets hello interval in seconds
- **hello-multiplier**—sets the multiplier for hello holding time
- **network**—sets the network type
- **password**—configures the authentication password for the interface
- **priority**—sets the priority for designated router election

The following commands are applicable to **IPv6 only** for implementing IS-IS:

ADDRESS-FAMILY context:

- **adjacency-check**—checks IS-IS neighbor protocol support
- **default-information**—controls distribution of default information
- **advertise {level2-into-level1 | level1-into-level2} [prefix-list-name]**—controls which IP routes flow between Layer 1 and Layer 2.

INTERFACE CONFIGURATION mode:

- **ipv6 router isis**—enables IS-IS routing for IPv6
- **isis ipv6 metric**—configures the metric for IPv6

Configuration Task List for IS-IS

The following list includes the configuration tasks for IS-IS:

- enable IS-IS on page 512 (mandatory)
- configure IS-IS interface parameters on page 515 (mandatory)
- change LSP attributes on page 516 (optional)
- configure IS-IS metric style and cost on page 517 (optional)
- change the is-type on page 519 (optional)
- control routing updates on page 520 (optional)
- configure authentication passwords on page 522 (optional)

- set the overload bit on page 523 (optional)
- debug IS-IS on page 524 (optional)

enable IS-IS

By default, IS-IS is not enabled.

You can create one instance of IS-IS on the E-Series. To enable IS-IS globally on the E-Series, you must create an IS-IS routing process and assign a NET address. To exchange protocol information with neighbors, enable IS-IS on an interface, instead of on a network as with other routing protocols.

In IS-IS, neighbors form adjacencies only when they are same IS type. For example, a Level 1 router never forms an adjacency with a Level 2 router. A Level 1-2 router will form Level 1 adjacencies with a neighboring Level 1 router and will form Level 2 adjacencies with a neighboring Level 2 router.

To configure IS-IS globally on the E-Series, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command	Command Mode	Purpose
1	**router isis** [*tag*]	CONFIGURATION	Create an IS-IS routing process. • *tag* is optional and identifies the name of the IS-IS process.
2	**net** *network-entity-title*	ROUTER ISIS	Configure an IS-IS network entity title (NET) for a routing process. Specify the area address and system ID for an IS-IS routing process. The last byte must be 00. Refer to for more information on configuring a NET.
3	**interface** *interface*	CONFIGURATION	Enter the interface configuration mode. Enter the keyword **interface** followed by the type of interface and slot/port information: • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For the Loopback interface on the RPM, enter the keyword **loopback** followed by a number from 0 to 16383. • For a port channel, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN, enter the keyword **vlan** followed by a number from 1 to 4094.

Step	Command	Command Mode	Purpose
4	**ip address** *ip-address mask*	INTERFACE	Assign an IP address and mask to the interface. The IP address must be on the same subnet as other IS-IS neighbors, but the IP address does not need to relate to the NET address.
5	**ip router isis** [*tag*]	INTERFACE	Enable IS-IS on the interface. If you configure a *tag* variable, it must be the same as the *tag* variable assigned in step 1.

The default IS type is level-1-2. To change the IS type to Level 1 only or Level 2 only, use the **is-type** command in the ROUTER ISIS mode.

To view the IS-IS configuration, enter the **show isis protocol** command in the EXEC privilege mode or the **show config** command in the ROUTER ISIS mode.

```
Force10#show isis protocol
IS-IS Router: <Null Tag>
  System Id: EEEE.EEEE.EEEE  IS-Type: level-1-2
  Manual area address(es):
   47.0004.004d.0001
  Routing for area address(es):
   21.2223.2425.2627.2829.3031.3233
   47.0004.004d.0001
  Interfaces supported by IS-IS:
   Vlan 2
   GigabitEthernet 4/22
   Loopback 0
  Redistributing:
  Distance: 115
  Generate narrow metrics: level-1-2
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   none
  Accept wide metrics:     none
Force10#
```

Figure 332 show isis protocol Command Example

To view IS-IS protocol statistics, use the **show isis traffic** command in the EXEC privilege mode (Figure 333).

```
Force10#show isis traffic
 IS-IS: Level-1 Hellos (sent/rcvd) : 4272/1538
 IS-IS: Level-2 Hellos (sent/rcvd) : 4272/1538
 IS-IS: PTP Hellos (sent/rcvd)     : 0/0
 IS-IS: Level-1 LSPs sourced (new/refresh) : 0/0
 IS-IS: Level-2 LSPs sourced (new/refresh) : 0/0
 IS-IS: Level-1 LSPs flooded (sent/rcvd) : 32/19
 IS-IS: Level-2 LSPs flooded (sent/rcvd) : 32/17
 IS-IS: Level-1 LSPs CSNPs (sent/rcvd) : 1538/0
 IS-IS: Level-2 LSPs CSNPs (sent/rcvd) : 1534/0
 IS-IS: Level-1 LSPs PSNPs (sent/rcvd) : 0/0
 IS-IS: Level-2 LSPs PSNPs (sent/rcvd) : 0/0
 IS-IS: Level-1 DR Elections : 2
 IS-IS: Level-2 DR Elections : 2
 IS-IS: Level-1 SPF Calculations : 29
 IS-IS: Level-2 SPF Calculations : 29
 IS-IS: LSP checksum errors received : 0
 IS-IS: LSP authentication failures : 0
Force10#
```

Figure 333 show isis traffic Command Example

You can assign additional NET addresses, but the System ID portion of the NET address must remain the same. FTOS supports up to six area addresses.

Some address considerations are:

- In order to be neighbors, Level 1 routers must be configured with at least one common area address.
- A Level 2 router becomes a neighbor with another Level 2 router regardless of the area address configured. However, if the area addresses are different, the link between the Level 2 routers is only at Level 2.

To view the configuration of the interface, use the **show config** command in the INTERFACE mode.

To view all interfaces configured with IS-IS routing and their defaults, use the **show isis interface** command in the EXEC privilege mode (Figure 334).

```
Force10#show isis interface G1/34
GigabitEthernet 1/34 is up, line protocol is up
  MTU 1497, Encapsulation SAP
  Routing Protocol: IS-IS
    Circuit Type: Level-1-2
    Interface Index 0x48cc03a, Local circuit ID 1
    Level-1 Metric: 10, Priority: 64, Circuit ID: Force10.01
            Hello Interval: 10, Hello Multiplier: 3, CSNP Interval: 10
    Level-1 IPv6 Metric: 10
    Number of active level-1 adjacencies: 1
    Level-2 Metric: 10, Priority: 64, Circuit ID: Force10.01
            Hello Interval: 10, Hello Multiplier: 3, CSNP Interval: 10
    Level-2 IPv6 Metric: 10
    Number of active level-2 adjacencies: 1
    Next IS-IS LAN Level-1 Hello in 2 seconds
    Next IS-IS LAN Level-2 Hello in 1 seconds
    LSP Interval: 33
Force10#
```

Figure 334 show isis interface IPv6 Command Example

configure IS-IS interface parameters

You must enable the IS-IS process on an interface for the IS-IS process to exchange protocol information and form adjacencies. You can modify IS-IS parameters on a per-interface basis, but it is not necessary.

To change IS-IS defaults on an interface, use any or all of the following commands in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
isis circuit-type {**level-1** | **level-1-2** | **level-2-only**}	INTERFACE	Configure the circuit type for the interface. Default is level-1-2.
isis csnp-interval *seconds* [**level-1** | **level-2**]	INTERFACE	Configure the complete sequence number PDU (CSNP) interval. • *seconds* range: 0 to 65535. Default is 10 seconds. Default level is level-1.
isis hello-interval *seconds* [**level-1** | **level-2**]	INTERFACE	Specify the length of time between hello packets sent by FTOS. • *seconds* range: 0 to 65535. Default is 10 seconds. Default level is level-1.

Command Syntax	Command Mode	Purpose
isis hello-multiplier *multiplier* [**level-1** | **level-2**]	INTERFACE	Specify the number of IS-IS hello packets a neighbor must miss before the router declares the adjacency as down. • *multiplier* range: 3 to 1000. Default is 3. Default level is level-1.
isis metric *default-metric* [**level-1** | **level-2**]	INTERFACE	Assign a metric for a link or interface. • *default-metric* range: 0 to 63 for narrow and transition metric styles; 0 to 16777215 for wide metric styles. Default is 10. Default level is level-1. Refer to for more information on this command.
isis password [**hmac-md5**] *password* [**level-1** | **level-2**]	INTERFACE	Configure the password to authenticate between IS-IS neighbors. Simple HMAC-MD5 authentication is supported. • *password:* a text string Default level is level-1. The password must be the same on all neighbors to form adjacencies.
isis priority *value* [**level-1** | **level-2**]	INTERFACE	Set the priority for Designated Router election on the interface. • *value* range: 0 to 127. Default is 64. Default level is level-1.

To view the interface’s non-default configuration, use the **show config** command in the INTERFACE mode.

To view all interfaces routing IS-IS, use the **show isis interface** command in the EXEC privilege mode (Figure 334).

change LSP attributes

IS-IS routers flood Link state PDUs (LSPs) to exchange routing information. LSP attributes include the generation interval, maximum transmission unit (MTU) or size, and the refresh interval. You can modify the LSP attribute defaults, but it is not necessary.

To change the defaults, use any or all of the following commands in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
lsp-gen-interval [**level-1** | **level-2**] *seconds*	ROUTER ISIS	Set interval between LSP generation. • *seconds* range: 0 to 120 Default is 5 seconds. Default level is Level 1.
lsp-mtu *size*	ROUTER ISIS	Set the LSP size. • *size* range: 128 to 9195. Default is 1497.
lsp-refresh-interval *seconds*	ROUTER ISIS	Set the LSP refresh interval. • *seconds* range: 1 to 65535. Default is 900 seconds.
max-lsp-lifetime *seconds*	ROUTER ISIS	Set the maximum time LSPs lifetime. • *seconds* range: 1 to 65535 Default is 1200 seconds.

To view the configuration, use the **show config** command in the ROUTER ISIS mode or the **show running-config isis** command in the EXEC privilege mode (Figure 335).

```
Force10#show running-config isis
!
router isis
 lsp-refresh-interval 902
 net 47.0005.0001.000C.000A.4321.00
 net 51.0005.0001.000C.000A.4321.00
Force10#
```

Figure 335 show running-config isis Command Example

configure IS-IS metric style and cost

All IS-IS links or interfaces are associated with a cost that is used in the SPF calculations. The possible cost varies depending on the metric style supported. If you configure narrow, transition or narrow transition metric style, the cost can be a number between 0 and 63. If you configure wide or wide transition metric style, the cost can be a number between 0 and 16,777,215. FTOS supports five different metric styles: narrow, wide, transition, narrow transition, and wide transition.

By default, FTOS generates and receives narrow metric values. Metrics or costs higher than 63 are not supported. To accept or generate routes with a higher metric, you must change the metric style of the IS-IS process. For example, if metric is configured as narrow, and an LSP with wide metrics is received, the route is not installed.

FTOS supports the following IS-IS metric styles:

Table 47 Metric Styles

Metric Style	Characteristics	Cost Range Supported on IS-IS Interfaces
narrow	Sends and accepts narrow or old TLVs (Type Length Value).	0 to 63
wide	Sends and accepts wide or new TLVs	0 to 16777215
transition	Sends both wide (new) and narrow (old) TLVs.	0 to 63
narrow transition	Sends narrow (old) TLVs and accepts both narrow (old) and wide (new) TLVs	0 to 63
wide transition	Sends wide (new) TLVs and accepts both narrow (old) and wide (new) TLVs.	0 to 16777215

To change the IS-IS metric style of the IS-IS process, use the following command in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
metric-style {**narrow** [**transition**] | **transition** | **wide** [**transition**]} [**level-1** | **level-2**]	ROUTER ISIS	Set the metric style for the IS-IS process. Default: narrow Default: Level 1 and Level 2 (level-1-2)

To view which metric types are generated and received, use the **show isis protocol** command (Figure 332) in the EXEC privilege mode.

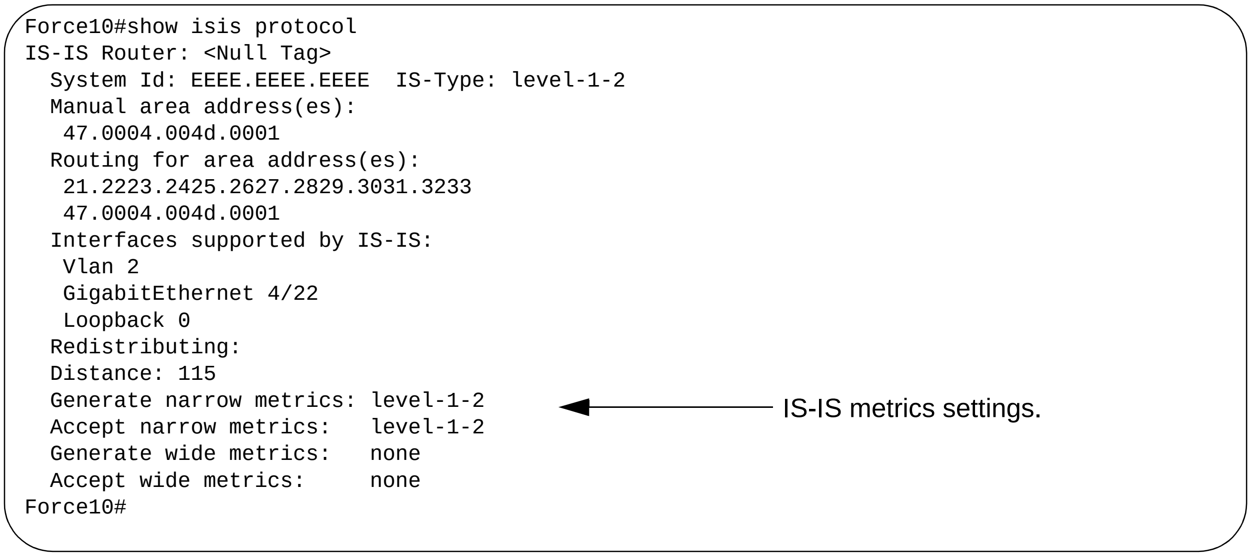

```
Force10#show isis protocol
IS-IS Router: <Null Tag>
  System Id: EEEE.EEEE.EEEE  IS-Type: level-1-2
  Manual area address(es):
   47.0004.004d.0001
  Routing for area address(es):
   21.2223.2425.2627.2829.3031.3233
   47.0004.004d.0001
  Interfaces supported by IS-IS:
   Vlan 2
   GigabitEthernet 4/22
   Loopback 0
  Redistributing:
  Distance: 115
  Generate narrow metrics: level-1-2
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   none
  Accept wide metrics:     none
Force10#
```

IS-IS metrics settings.

Figure 336 show isis protocol Command Example

When you change from one IS-IS metric style to another, the IS-IS metric value could be affected. For each interface with IS-IS enabled, you can assign a cost or metric that is used in the link state calculation. Appendix D, contains details on the behavior of the metric value when you change the metric style.

To change the metric or cost of the interface, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
isis metric *default-metric* [**level-1** | **level-2**]	INTERFACE	*default-value* range: 0 to 63 if the metric-style is narrow, narrow-transition or transition. 0 to 16777215 if the metric style is wide or wide transition. Default: 10.

To view the interface's current metric, use the **show config** command in the INTERFACE mode or the **show isis interface** command in the EXEC privilege mode.

Note: In FTOS, the CLI help always shows the value range (0-16777215) for the metric style. See Table 48 for the correct value range.

Table 48 Correct Value Range for the isis metric command

Metric Style	Correct Value Range
wide	0 to 16777215
narrow	0 to 63
wide transition	0 to 16777215
narrow transition	0 to 63
transition	0 to 63

change the is-type

You can configure the E-Series system to act as one of the following:

- Level 1 router
- Level 1-2 router
- Level 2 router

To change the is-type for the router, use the following command in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
is-type {**level-1** | **level-1-2** | **level-2**}	ROUTER ISIS	Change the is-type for the IS-IS process.

To view which is-type is configured, use the **show isis protocol** command in the EXEC privilege mode (Figure 332). The **show config** command in the ROUTER ISIS mode displays only nondefault information, so if you do not change the is-type, the default value (level-1-2) is not displayed.

The default is Level 1-2 router. When the is-type is Level 1-2, the software maintains two Link State databases, one for each level. Use the **show isis database** command to view the Link State databases (Figure 337).

```
Force10#show isis database
IS-IS Level-1 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
B233.00-00            0x00000003   0x07BF        1088              0/0/0
eljefe.00-00        * 0x00000009   0xF76A        1126              0/0/0
eljefe.01-00        * 0x00000001   0x68DF        1122              0/0/0
eljefe.02-00        * 0x00000001   0x2E7F        1113              0/0/0
Force10.00-00         0x00000002   0xD1A7        1102              0/0/0
IS-IS Level-2 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
B233.00-00            0x00000006   0xC38A        1124              0/0/0
eljefe.00-00        * 0x0000000D   0x51C6        1129              0/0/0
eljefe.01-00        * 0x00000001   0x68DF        1122              0/0/0
eljefe.02-00        * 0x00000001   0x2E7F        1113              0/0/0
Force10.00-00         0x00000004   0xCDA9        1107              0/0/0

Force10#
```

Figure 337 show isis database Command Example

control routing updates

To control the source of IS-IS route information, use the following commands, in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
passive-interface *interface*	ROUTER ISIS	Disable a specific interface from sending or receiving IS-IS routing information. Enter the type of interface and slot/port information: • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For the Loopback interface on the RPM, enter the keyword **loopback** followed by a number from 0 to 16383. • For a port channel, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN, enter the keyword **vlan** followed by a number from 1 to 4094.

Another method of controlling routing information is to filter the information through a prefix list. Prefix lists are applied to incoming or outgoing routes and routes must meet the conditions of the prefix lists or FTOS does not install the route in the routing table. The prefix lists are globally applied on all interfaces running IS-IS. Configure the prefix list in the PREFIX LIST mode prior to assigning it to the IS-IS process.

For configuration information on prefix lists, see Chapter 17, IP Access Control Lists, Prefix Lists, and Route-maps.

To apply prefix lists to incoming or outgoing routes, use the following commands in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
distribute-list *prefix-list-name* **in** [*interface*]	ROUTER ISIS	Apply a configured prefix list to all incoming IS-IS routes. Enter the type of interface and slot/port information: • For a 1-Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For the Loopback interface on the RPM, enter the keyword **loopback** followed by a number from 0 to 16383. • For a port channel, enter the keyword **port-channel** followed by a number from 1 to 255 for TeraScale, 1 to 32 for EtherScale. • For a SONET interface, enter the keyword **sonet** followed by slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • For a VLAN, enter the keyword **vlan** followed by a number from 1 to 4094.
distribute-list *prefix-list-name* **out** [**connected** | **ospf** *process-id* | **rip** | **static**]	ROUTER ISIS	Apply a configured prefix list to all outgoing IS-IS routes. You can configure one of the optional parameters: • **connected:** for directly connected routes. • **ospf** *process-id:* for OSPF routes only. • **rip:** for RIP routes only. • **static:** for user-configured routes.

In addition to filtering routes, you can add routes from other routing instances or protocols to the IS-IS process. With the **redistribute** command syntax, you can include BGP, OSPF, RIP, static, or directly connected routes in the IS-IS process.

To add routes from other routing instances or protocols, use any of the following commands in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
redistribute {**connected** | **rip** | **static**} [**level-1 level-1-2** | **level-2**] [**metric** *metric-value*] [**metric-type** {**external** | **internal**}] [**route-map** *map-name*]	ROUTER ISIS	Include directly connected, RIP, or user-configured (static) routes in IS-IS. Configure the following parameters: • **level-1**, l**evel-1-2**, or **level-2**: Assign all redistributed routes to a level. Default is **level-2**. • *metric* range: 0 to 16777215. Default is 0. • **metric-type:** choose either **external** or **internal**. Default is internal. • *map-name*: name of a configured route map.
redistribute ospf *process-id* [**level-1**| **level-1-2** | **level-2**] [**metric** *value*] [**match external** {**1** | **2**} | **match internal**] [**metric-type** {**external** | **internal**}] [**route-map** *map-name*]	ROUTER ISIS	Include specific OSPF routes in IS-IS. Configure the following parameters: • *process-id* range: 1 to 65535 • **level-1**, l**evel-1-2**, or **level-2**: Assign all redistributed routes to a level. Default is **level-2**. • *metric* range: 0 to 16777215. Default is 0. • **match external** range: 1 or 2 • **match internal** • **metric-type**: external or internal. • *map-name*: name of a configured route map.

Note: Starting with Release 6.3.1, you can also redistribute BGP routes in IS-IS and IS-IS through BGP. See the IS-IS chapter and BGP chapters in the *FTOS Command Line Interface Reference* for details.

To view the current IS-IS configuration, use the **show running-config isis** command in the EXEC privilege mode or the **show config** command in the ROUTER ISIS mode.

configure authentication passwords

You can assign an authentication password for routers in Level 1 and for routers in Level 2. Since Level 1 and Level 2 routers do not communicate with each other, you can assign different passwords for Level 1 routers and for Level 2 routers. If you want the routers in the level to communicate with each other, though, they must be configured with the same password.

To configure a simple text password, use either or both of the commands in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
area-password [**hmac-md5**] *password*	ROUTER ISIS	Configure authentication password for an area. FTOS supports HMAC-MD5 authentication. This password is inserted in Level 1 LSPs, Complete SNPs, and Partial SNPs.
domain-password [*encryption-type* | **hmac-md5**] *password*	ROUTER ISIS	Set the authentication password for a routing domain. FTOS supports both DES and HMAC-MD5 authentication methods. This password is inserted in Level 2 LSPs, Complete SNPs, and Partial SNPs.

To view the passwords, use the **show config** command in the ROUTER ISIS mode or the **show running-config isis** command in the EXEC privilege mode.

To remove a password, use either **no area-password** or **no domain-password** commands in the ROUTER ISIS mode.

set the overload bit

Another use for the overload bit is to prevent other routers from using the E-Series as an intermediate hop in their shortest path first (SPF) calculations. For example, if the IS-IS routing database is out of memory and cannot accept new LSPs, FTOS sets the overload bit and IS-IS traffic continues to transit the E-Series.

To set the overload bit manually, use this command the following command in the ROUTER ISIS mode:

Command Syntax	Command Mode	Purpose
set-overload-bit	ROUTER ISIS	Set the overload bit in LSPs. This prevents other routers from using it as an intermediate hop in their shortest path first (SPF) calculations.

To remove the overload bit, enter **no set-overload-bit**.

To see if the bit is set, a 1 is placed in the OL column in the show isis database command output. In Figure 338, the overload bit is set in both the Level-1 and Level-2 database because the IS type for the router is Level-1-2.

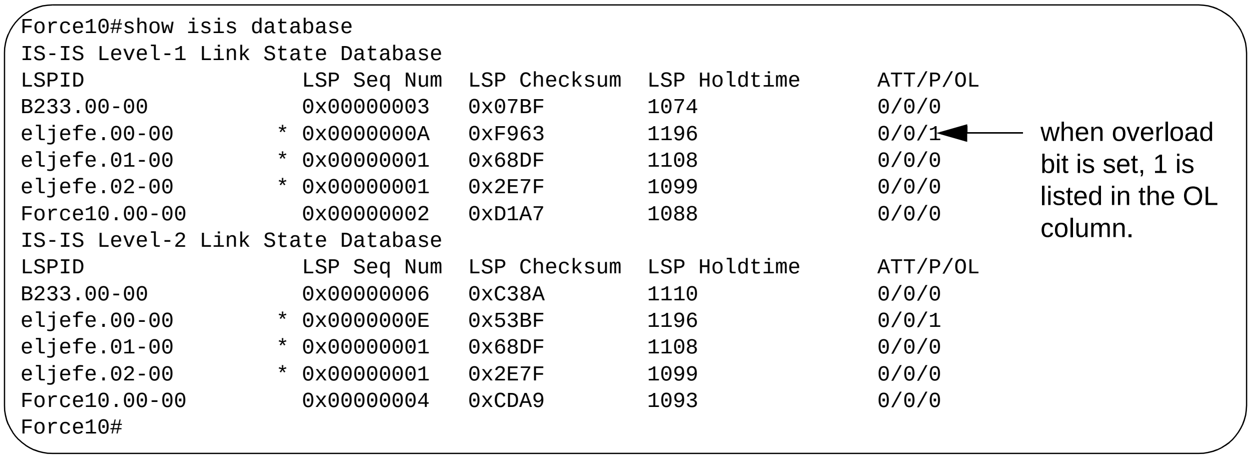

```
Force10#show isis database
IS-IS Level-1 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
B233.00-00             0x00000003   0x07BF        1074              0/0/0
eljefe.00-00         * 0x0000000A   0xF963        1196              0/0/1
eljefe.01-00         * 0x00000001   0x68DF        1108              0/0/0
eljefe.02-00         * 0x00000001   0x2E7F        1099              0/0/0
Force10.00-00          0x00000002   0xD1A7        1088              0/0/0
IS-IS Level-2 Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
B233.00-00             0x00000006   0xC38A        1110              0/0/0
eljefe.00-00         * 0x0000000E   0x53BF        1196              0/0/1
eljefe.01-00         * 0x00000001   0x68DF        1108              0/0/0
eljefe.02-00         * 0x00000001   0x2E7F        1099              0/0/0
Force10.00-00          0x00000004   0xCDA9        1093              0/0/0
Force10#
```

when overload bit is set, 1 is listed in the OL column.

Figure 338 show isis database Command Example

debug IS-IS

To debug all IS-IS processes, enter the **debug isis** command in the EXEC privilege mode.

Use the following commands for specific IS-IS debugging:

Command Syntax	Command Mode	Purpose
debug isis adj-packets [*interface*]	EXEC privilege	View information on all adjacency-related activity (for example, hello packets that are sent and received). To view specific information, enter one of the following optional parameters: • *interface:* Enter the type of interface and slot/port information to view IS-IS information on that interface only.
debug isis local-updates [*interface*]	EXEC privilege	View information about IS-IS local update packets. To view specific information, enter one of the following optional parameters: • *interface:* Enter the type of interface and slot/port information to view IS-IS information on that interface only.
debug isis snp-packets [*interface*]	EXEC privilege	View IS-IS SNP packets, include CSNPs and PSNPs. To view specific information, enter one of the following optional parameters: • *interface:* Enter the type of interface and slot/port information to view IS-IS information on that interface only.

Command Syntax	Command Mode	Purpose
debug isis spf-triggers	EXEC privilege	View the events that triggered IS-IS shortest path first (SPF) events for debugging purposes.
debug isis update-packets [*interface*]	EXEC privilege	View sent and received LSPs. To view specific information, enter one of the following optional parameters: • *interface:* Enter the type of interface and slot/port information to view IS-IS information on that interface only.

FTOS displays debug messages on the console. To view which debugging commands are enabled, use the **show debugging** command in the EXEC privilege mode.

To disable a specific debug command, enter the keyword no followed by the debug command. For example, to disable debugging of IS-IS updates, you enter **no debug isis updates-packets** command.

To disable all IS-IS debugging, enter **no debug isis**.

To disable all debugging, enter **undebug all**.

Chapter 32 BGP

C-Series	NO
E-Series	✓

Platform Specific Feature: Border Gateway Protocol is supported on E-Series only.

FTOS supports Border Gateway Protocol (BGP) version 4. This chapter describes protocol configuration information and contains the following sections:

- Border Gateway Protocol on page 527
- Force10 Implementation of BGP on page 528
- BGP Configuration on page 530
- MBGP Configuration on page 559
- BGP4 MIB on page 560

Border Gateway Protocol

Border Gateway Protocol (BGP) is an external gateway protocol that transmits interdomain routing information within and between Autonomous Systems (AS). Its primary function is to exchange network reachability information with other BGP systems. Internal BGP (IBGP) exchanges routing information between BGP routers within the same AS and External BGP (EBGP) exchanges routing information between BGP routers in different ASs. IBGP provides internal routers with information on reaching external destinations.

BGP version 4 (BGPv4) supports classless interdomain routing and the aggregation of routes and AS paths. Basically, two routers (called neighbors or peers) exchange information including full routing tables and periodically send messages to update those routing tables.

BGP RFCs Supported

The E-Series implementation of BGP is based on the following IETF documents:

- RFC 1771 (BGPv4)

- ID draft-ietf-idr-bgp4-15.txt (revision to BGPv4)
- RFC 1772 (Application of BGP in the Internet)
- RFC 1997 (BGP Communities Attribute)
- RFC 1998 (Application of the BGP Community Attribute in Multi-home Routing)
- RFC 2270 (Using a Dedicated AS for Sites Homed to a Single Provider)
- RFC 2439 (BGP Route Flap Dampening)
- RFC 2519 (A Framework for Inter-Domain Route Aggregation)
- RFC 2796 (BGP Route Reflection - An Alternative to Full Mesh IBGP)
- RFC 2842 (Capabilities Advertisement with BGP-4)
- RFC 3065 (Autonomous System Confederations for BGP)

Force10 Implementation of BGP

The E-Series software supports BGPv4 as well as the following:

- deterministic MED is the default
- a path with a missing MED is treated as worst and assigned a MED value of (`0xffffffff`)
- the community format follows RFC 1998.
- delayed configuration, which means that the software at system boot reads the entire configuration file prior to sending messages to start BGP peer sessions.

In the E-Series software, the following are not yet supported:

- auto-summarization (the default is no auto-summary);
- synchronization (the default is no synchronization).

Best Path Selection Criteria

Paths for active routes are grouped in ascending order according to their neighboring external AS number (BGP best path selection is deterministic by default, which means the **bgp non-deterministic-med** command is NOT applied). The best path in each group is selected based on the criteria listed below.

1. Prefer the path with the largest WEIGHT attribute.
2. Prefer the path with the largest LOCAL_PREF attribute.

Prefer the path that was locally originated via a **network** command, **redistribute** command or **aggregate-address** command. Routes originated via the **network** or **redistribute** commands are preferred over routes originated via the **aggregate-address** command.

3. Prefer the path with the shortest AS_PATH (unless the **bgp bestpath as-path ignore** command is configured, then AS_PATH is not considered). The following criteria apply:
    - An AS_SET has a path length of 1, no matter how many ASs are in the set.

    - A path with no AS_PATH configured has a path length of 0.
    - AS_CONFED_SET is not included in the AS_PATH length.
    - AS_CONFED_SEQUENCE has a path length of 1, no matter how many ASs are in the AS_CONFED_SEQUENCE.
4. Prefer the path with the lowest origin type (IGP is lower than EGP, and EGP is lower than INCOMPLETE).
5. Prefer the path with the lowest multi-exit discriminator (MED) attribute. The following criteria apply:
    - This comparison is only done if the first (neighboring) AS is the same in the two paths. In other words, the MEDs are compared only if the first AS in the AS_SEQUENCE is the same for both paths.
    - If the **bgp always-compare-med** command is entered, MEDs are compared for all paths.
    - Paths with no MED are treated as "worst" and assigned a MED of 4294967295.
6. Prefer external (EBGP) to internal (IBGP) paths or confederation EBGP paths.
7. Prefer the path with the lowest IGP metric to the BGP next-hop.
8. FTOS deems the paths as equal and does not perform steps 10 through 12 listed below, if the following criteria is met:
    - the IBGP multipath or EBGP multipath are configured (**maximum-path** command)
    - the paths being compared were received from the same AS with the same number of ASes in the AS Path but with different NextHops
    - the paths were received from IBGP or EBGP neighbor respectively
9. Prefer the path originated from the BGP router with the lowest router ID. For paths containing a Route Reflector (RR) attribute, the originator ID is substituted for the router ID.
10. If two paths have the same router ID, prefer the path with the lowest cluster ID length. Paths without a cluster ID length are set to a 0 cluster ID length.
11. Prefer the path originated from the neighbor with the lowest address. (The neighbor address is used in the BGP neighbor configuration, and corresponds to the remote peer used in the TCP connection with the local router.)

After a number of best paths is determined, this selection criteria is applied to group's best to determine the ultimate best path.

In non-deterministic mode (the **bgp non-deterministic-med** command is applied), paths are compared in the order in which they arrive. This method can lead to FTOS choosing different best paths from a set of paths, depending on the order in which they were received from the neighbors since MED may or may not get compared between adjacent paths. In deterministic mode, FTOS compares MED between adjacent paths within an AS group since all paths in the AS group are from the same AS.

BGP Enhancements

BGP Configuration

To enable the BGP process and begin exchanging information, you must assign an AS number and use commands in the ROUTER BGP mode to configure a BGP neighbor.

By default, BGP is disabled.

Defaults

By default, FTOS compares the MED attribute on different paths from within the same AS (that is, the **bgp always-compare-med** command is not enabled).

Note: In FTOS, all newly configured neighbors and peer groups are disabled. You must enter the **neighbor** {*ip-address* | *peer-group-name*} **no shutdown** command to enable a neighbor or peer group.

Table 49 displays the default values for BGP on FTOS.

Table 49 FTOS BGP Defaults

Item	Default
BGP Neighbor Adjacency changes	All BGP neighbor changes are logged.
Fast External Fallover feature	Enabled
graceful restart feature	Disabled
Local preference	100
MED	0

Route Flap Damping Parameters	half-life = 15 minutes
	reuse = 750
	suppress = 2000
	max-suppress-time = 60 minutes

Distance	external distance = 20 internal distance = 200 local distance = 200

Timers	keepalive = 60 seconds
	holdtime = 180 seconds

Configuration Task List for BGP

The following list includes the configuration tasks for BGP:

- enable BGP by configuring BGP neighbors on page 531 (required)
- configure peer groups on page 534
- configure passive peering on page 536
- enable graceful restart on page 537

- filter on AS-Path attribute on page 539
- configure IP community lists on page 541
- manipulate the COMMUNITY attribute on page 543
- change MED attribute on page 544
- change LOCAL_PREFERENCE attribute on page 545
- change NEXT_HOP attribute on page 546
- change WEIGHT attribute on page 546
- enable multipath on page 547
- filter BGP routes on page 547
- configure BGP route reflectors on page 550
- aggregate routes on page 550
- configure BGP confederations on page 551
- enable route flap dampening on page 552
- change path selection to non-deterministic on page 554
- change BGP timers on page 555
- Debugging BGP on page 555

For a complete listing of all commands related to BGP, refer to .

enable BGP by configuring BGP neighbors

By default, BGP is not enabled on the E-Series. FTOS supports one Autonomous System (AS) and you must assign an AS number. To establish BGP sessions and route traffic, you must configure at least one BGP neighbor or peer.

In BGP, routers with an established TCP connection are called neighbors or peers. Once a connection is established, the neighbors exchange full BGP routing tables with incremental updates afterwards. In addition, neighbors exchange KEEPALIVE messages to maintain the connection.

In BGP, neighbor routers or peers can be classified as internal or external. External BGP peers must be connected physically to one another (unless you enable the EBGP multihop feature), while internal BGP peers do not need to be directly connected. The IP address of an EBGP neighbor is usually the IP address of the interface directly connected to the router. First, the BGP process determines if all internal BGP peers are reachable, and then it determines which peers outside the AS are reachable.

To establish BGP sessions on the router, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**router bgp** *as-number*	CONFIGURATION	Assign an AS number and enter the ROUTER BGP mode. Only one AS is supported per E-Series system.

Step	Command Syntax	Command Mode	Purpose
2	**neighbor** {*ip-address* | *peer-group-name*} **remote-as** *number*	ROUTER BGP	Add a neighbor by specifying its IP address. (You must first create a peer group before assigning it a remote AS.) To add an external BGP neighbor, configure the *as-number* parameter with a number different from the BGP *as-number* configured in the **router bgp** *as-number* command. To add an internal BGP neighbor, configure the *as-number* parameter with the same BGP *as-number* configured in the **router bgp** *as-number* command.
3	**neighbor** {*ip-address* | *peer-group-name*} **no shutdown**	ROUTER BGP	Enable the BGP neighbor.

Note: When you change the configuration of a BGP neighbor, always reset it by entering the **clear ip bgp** command in the EXEC privilege mode.

To view the BGP configuration, enter **show config** in the ROUTER BGP mode. To view the BGP status, use the **show ip bgp summary** command in the EXEC privilege mode (Figure 339).

```
E1200>show ip bgp summary
BGP router identifier 63.114.8.39, local AS number 65519
BGP table version is 74001, main routing table version 11163
56123 network entrie(s) and 95183 paths using 13742524 bytes of memory
7665 BGP path attribute entrie(s) using 429240 bytes of memory
7127 BGP AS-PATH entrie(s) using 328447 bytes of memory
157 BGP community entrie(s) using 6383 bytes of memory

Neighbor        AS      MsgRcvd  MsgSent     TblVer  InQ  OutQ Up/Down  State/Pfx

192.168.0.0    18508     2153        3          0   30     0 00:00:16      5640
192.168.0.1    18508     2629        3          0   30     0 00:00:16     42154
192.168.0.2    18508     2469        3          0   14     0 00:00:16     11979
192.168.0.3    18508     2236        3          0   30     0 00:00:16     35410
E1200>
```

Figure 339 show ip bgp summary Command Example

For the router's identifier, FTOS uses the highest IP address of the Loopback interfaces configured. Since Loopback interfaces are virtual, they cannot go down, thus preventing changes in the router ID. If no Loopback interfaces are configured, the highest IP address of any interface is used as the router ID.

To view the status of BGP neighbors, use the **show ip bgp neighbors** (Figure 340) command in the EXEC privilege mode. For BGP neighbor configuration information, use the **show running-config bgp** command in the EXEC privilege mode (Figure 341).

Figure 340 displays two neighbors, one is an external and the second one is an internal BGP neighbor. The first line of the output for each neighbor displays the AS number and states whether the link is an external or internal.

The third line of the **show ip bgp neighbors** output contains the BGP State. If anything other than ESTABLISHED is listed, the neighbor is not exchanging information and routes. For more details on using the **show ip bgp neighbors** command, refer to the .

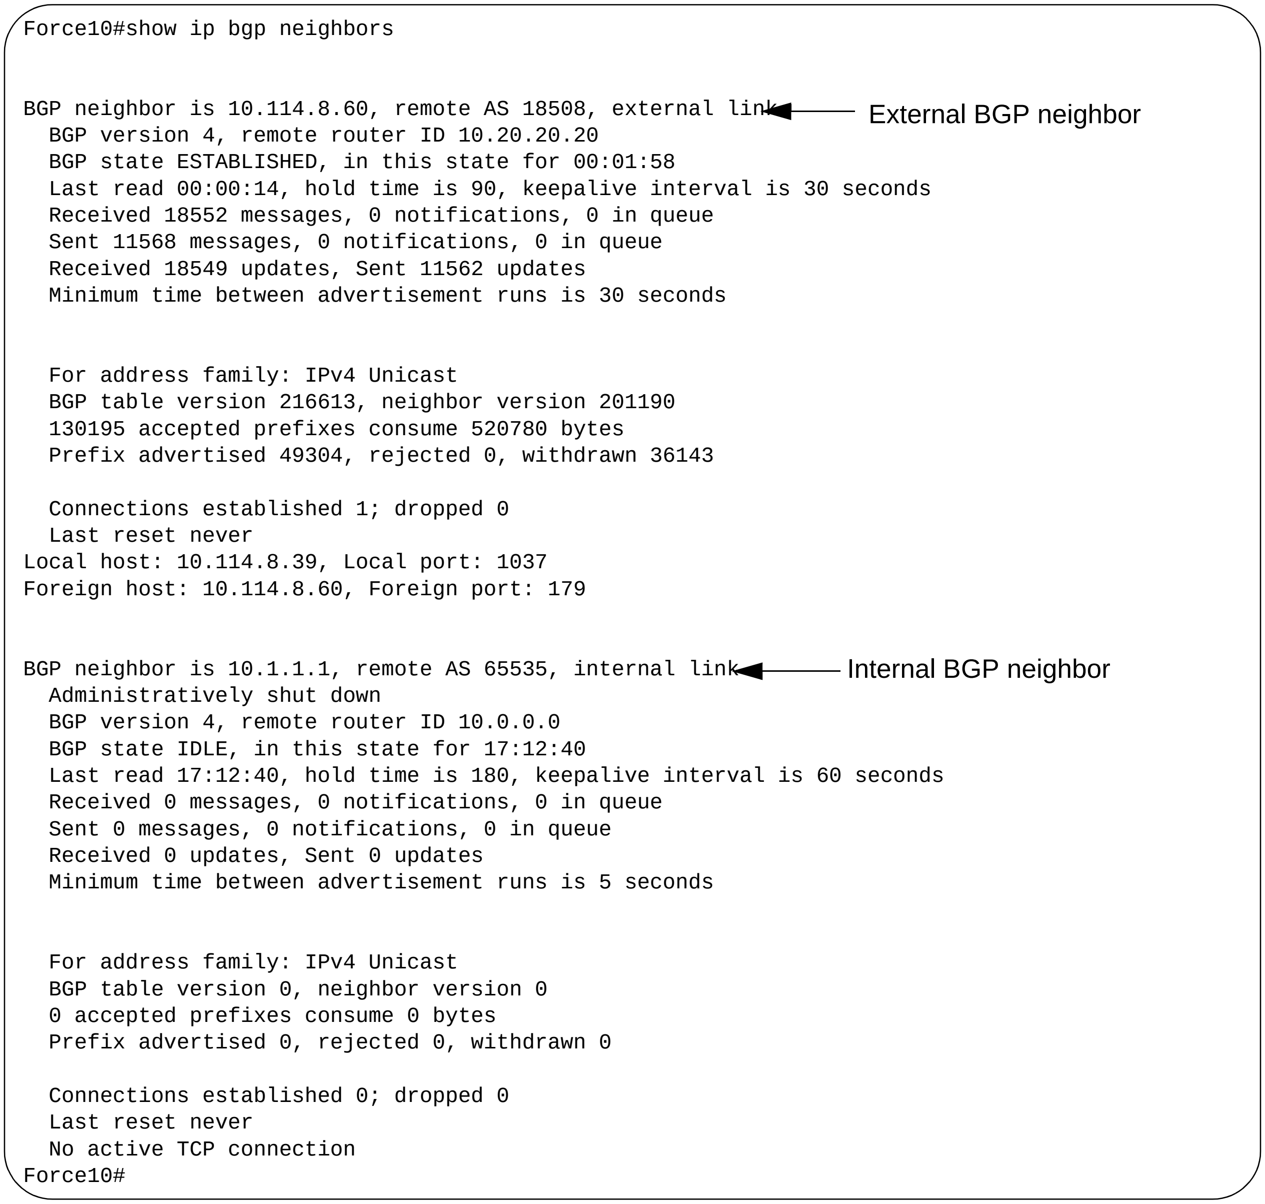

```
Force10#show ip bgp neighbors


BGP neighbor is 10.114.8.60, remote AS 18508, external link
  BGP version 4, remote router ID 10.20.20.20
  BGP state ESTABLISHED, in this state for 00:01:58
  Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
  Received 18552 messages, 0 notifications, 0 in queue
  Sent 11568 messages, 0 notifications, 0 in queue
  Received 18549 updates, Sent 11562 updates
  Minimum time between advertisement runs is 30 seconds


  For address family: IPv4 Unicast
  BGP table version 216613, neighbor version 201190
  130195 accepted prefixes consume 520780 bytes
  Prefix advertised 49304, rejected 0, withdrawn 36143

  Connections established 1; dropped 0
  Last reset never
Local host: 10.114.8.39, Local port: 1037
Foreign host: 10.114.8.60, Foreign port: 179


BGP neighbor is 10.1.1.1, remote AS 65535, internal link
  Administratively shut down
  BGP version 4, remote router ID 10.0.0.0
  BGP state IDLE, in this state for 17:12:40
  Last read 17:12:40, hold time is 180, keepalive interval is 60 seconds
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 messages, 0 notifications, 0 in queue
  Received 0 updates, Sent 0 updates
  Minimum time between advertisement runs is 5 seconds


  For address family: IPv4 Unicast
  BGP table version 0, neighbor version 0
  0 accepted prefixes consume 0 bytes
  Prefix advertised 0, rejected 0, withdrawn 0

  Connections established 0; dropped 0
  Last reset never
  No active TCP connection
Force10#
```

Figure 340 show ip bgp neighbors Command Example

```
Force10(conf-router_bgp)#show conf
!
router bgp 45
 bgp fast-external-fallover
 bgp log-neighbor-changes
 neighbor 10.1.1.1 remote-as 65535
 neighbor 10.1.1.1 shutdown
 neighbor 10.14.8.60 remote-as 18505
 neighbor 10.14.8.60 no shutdown
Force10(conf-router_bgp)#
```

Figure 341 show running-config bgp Command Example

configure peer groups

To configure multiple BGP neighbors at one time, create and populate a BGP peer group. Another advantage of peer groups is that members of a peer groups inherit the configuration properties of the group and share same update policy.

You create a peer group by assigning it a name, then adding members to the peer group. Once a peer group is created, you can configure route policies for it. Refer to filter BGP routes on page 547 for information on configuring route policies for a peer group.

To create a peer group, use these commands in the following sequence starting in the ROUTER BGP mode:

Step	Command Syntax	Command Mode	Purpose
1	**neighbor** *peer-group-name* **peer-group**	ROUTER BGP	Create a peer group by assigning a name to it.
2	**neighbor** *peer-group-name* **no shutdown**	ROUTER BGP	Enable the peer group. By default, all peer groups are disabled
3	**neighbor** *ip-address* **remote-as** *as-number*	ROUTER BGP	Create a BGP neighbor.
4	**neighbor** *ip-address* **no shutdown**	ROUTER BGP	Enable the neighbor.
5	**neighbor** *ip-address* **peer-group** *peer-group-name*	ROUTER BGP	Add an enabled neighbor to the peer group.

After you create a peer group, you can use any of the commands beginning with the keyword **neighbor** to configure that peer group. Refer to for a complete list of all commands beginning with the **neighbor** keyword.

When you add a peer to a peer group, it inherits all the peer group's configured parameters. A peer cannot become part of a peer group if any of the following commands are configured on the peer:

- **neighbor advertisement-interval**

- **neighbor distribute-list out**
- **neighbor filter-list out**
- **neighbor next-hop-self**
- **neighbor route-map out**
- **neighbor route-reflector-client**
- **neighbor send-community**

A neighbor may keep its configuration after it was added to a peer group if the neighbor's configuration is more specific than the peer group's, and the neighbor's configuration does not affect outgoing updates.

Note: When you configure a new set of BGP policies for a peer group, always reset the peer group by entering the **clear ip bgp peer-group** *peer-group-name* command in the EXEC privilege mode.

To view the configuration, use the **show config** command in the ROUTER BGP mode. When you create a peer group, it is disabled (**shutdown**). Figure 342 shows the creation of a peer group (Zanzibar).

```
Force10(conf-router_bgp)#neighbor zanzibar peer-group
Force10(conf-router_bgp)#show conf
!
router bgp 45
 bgp fast-external-fallover
 bgp log-neighbor-changes
 neighbor zanzibar peer-group
 neighbor zanzibar shutdown
 neighbor 10.1.1.1 remote-as 65535
 neighbor 10.1.1.1 shutdown
 neighbor 10.14.8.60 remote-as 18505
 neighbor 10.14.8.60 no shutdown
Force10(conf-router_bgp)#
```

Figure 342 show config Command Example in ROUTER BGP Mode

To enable a peer group, use the **neighbor** *peer-group-name* **no shutdown** command in the ROUTER BGP mode.

```
Force10(conf-router_bgp)#neighbor zanzibar no shutdown
Force10(conf-router_bgp)#show config
!
router bgp 45
 bgp fast-external-fallover
 bgp log-neighbor-changes
 neighbor zanzibar peer-group
 neighbor zanzibar no shutdown
 neighbor 10.1.1.1 remote-as 65535
 neighbor 10.1.1.1 shutdown
 neighbor 10.14.8.60 remote-as 18505
 neighbor 10.14.8.60 no shutdown
Force10(conf-router_bgp)#
```

Figure 343 show config Command Example with Enabled Peer Group

To disable a peer group, use the **neighbor** *peer-group-name* **shutdown** command in the ROUTER BGP mode. The configuration of the peer group is maintained, but it is not applied to the peer group members. When you disable a peer group, all the peers within the peer group that are in ESTABLISHED state are moved to IDLE state.

To view the status of peer groups, use the **show ip bgp peer-group** command in the EXEC privilege mode (Figure 344).

```
Force10>show ip bgp peer-group


Peer-group zanzibar, remote AS 65535
BGP version 4
Minimum time between advertisement runs is 5 seconds

For address family: IPv4 Unicast
BGP neighbor is zanzibar, peer-group internal,
Number of peers in this group 26
Peer-group members (* - outbound optimized):
  10.68.160.1
  10.68.161.1
  10.68.162.1
  10.68.163.1
  10.68.164.1
  10.68.165.1
  10.68.166.1
  10.68.167.1
  10.68.168.1
  10.68.169.1
  10.68.170.1
  10.68.171.1
  10.68.172.1
  10.68.173.1
  10.68.174.1
  10.68.175.1
  10.68.176.1
  10.68.177.1
  10.68.178.1
  10.68.179.1
  10.68.180.1
  10.68.181.1
  10.68.182.1
  10.68.183.1
  10.68.184.1
  10.68.185.1
Force10>
```

Figure 344 show ip bgp peer-group Command Example

configure passive peering

When you enable a peer-group, the software sends an OPEN message to initiate a TCP connection. If you enable passive peering for the peer group, the software does not send an OPEN message, but it will respond to an OPEN message.

When a BGP neighbor connection with authentication configured is rejected by a passive peer-group, FTOS does not allow another passive peer-group on the same subnet to connect with the BGP neighbor. To work around this, change the BGP configuration or change the order of the peer group configuration.

To configure passive peering, use these commands in the following sequence, starting in the ROUTER BGP mode:

Step	Command Syntax	Command Mode	Purpose
1	**neighbor** *peer-group-name* **peer-group passive**	ROUTER BGP	Configure a peer group that does not initiate TCP connections with other peers.
2	**neighbor** *peer-group-name* **subnet** *subnet-number mask*	ROUTER BGP	Assign a subnet to the peer group. The peer group will respond to OPEN messages sent on this subnet.
3	**neighbor** *peer-group-name* **no shutdown**	ROUTER BGP	Enable the peer group.
4	**neighbor** *peer-group-name* **remote-as** *as-number*	ROUTER BGP	Create and specify a remote peer for BGP neighbor.

Only after the peer group responds to an OPEN message sent on the subnet does its BGP state change to ESTABLISHED. Once the peer group is ESTABLISHED, the peer group is the same as any other peer group.

For more information on peer groups, refer to configure peer groups on page 534.

enable graceful restart

Use this feature to lessen the negative effects of a BGP restart. FTOS advertises support for this feature to BGP neighbors through a capability advertisement. You can enable graceful restart by router and/or by peer or peer group.

➡ **Note:** By default, BGP graceful restart is disabled.

The default role for BGP on the E-Series is as a receiving or restarting peer. If you enable BGP, when a peer that supports graceful restart resumes operating, FTOS performs the following tasks:

- Continues saving routes received from the peer if the peer advertised it had graceful restart capability. Continues forwarding traffic to the peer.
- Flags routes from the peer as Stale and sets a timer to delete them if the peer does not perform a graceful restart.
- Deletes all routes from the peer if forwarding state information is not saved.
- Speeds convergence by advertising a special update packet known as an end-of-RIB marker. This marker indicates the peer has been updated with all routes in the local RIB.

If you configure your E-Series to do so, FTOS can perform the following actions during a hot failover:

- Save all FIB and CAM entries on the line card and continue forwarding traffic while the secondary RPM is coming online.
- Advertise to all BGP neighbors and peer-groups that the forwarding state of all routes has been saved. This prompts all peers to continue saving the routes they receive from your E-Series and to continue forwarding traffic.
- Bring the secondary RPM online as the primary and re-open sessions with all peers operating in “no shutdown” mode.
- Defer best path selection for a certain amount of time. This help optimize path selection and results in fewer updates being sent out.

You enable graceful restart using the **configure router bgp graceful-restart** command. The table below shows the command and its available options:

Command Syntax	Command Mode	Usage
bgp graceful-restart	ROUTER BGP	Enable graceful restart for the BGP node.
bgp graceful-restart [**restart-time** *time-in-seconds*]	ROUTER BGP	Set maximum restart time for all peers. Default is 120 seconds.
bgp graceful-restart [**role receiver-only**]	ROUTER BGP	Local router supports graceful restart as a receiver only.
bgp graceful-restart [**stale-path-time** *time-in-seconds*]	ROUTER BGP	Set maximum time to retain the restarting peer’s stale paths. Default is 360 seconds.

With the graceful restart feature, FTOS enables the receiving/restarting mode by default. In receiver-only mode, graceful restart saves the advertised routes of peers that support this capability when they restart. However, the E-Series does not advertise that it saves these forwarding states when it restarts. Essentially, this option provides support for remote peers for their graceful restart without supporting the feature itself.

You can implement BGP graceful restart either by neighbor or by BGP peer-group. For more information, please see the following table or the *FTOS Command Line Interface Reference*.

Command Syntax	Command Mode	Purpose
neighbor {*ip-address* | *peer-group-name*} **graceful-restart**	ROUTER BGP	Add graceful restart to a BGP neighbor or peer-group.
neighbor {*ip-address* | *peer-group-name*} **graceful-restart** [**restart-time** *time-in-seconds*]	ROUTER BGP	Set maximum restart time for the neighbor or peer-group. Default is 120 seconds.
neighbor {*ip-address* | *peer-group-name*} **graceful-restart** [**role receiver-only**]	ROUTER BGP	Local router supports graceful restart for this neighbor or peer-group as a receiver only.
neighbor {*ip-address* | *peer-group-name*} **graceful-restart** [**stale-path-time** *time-in-seconds*]	ROUTER BGP	Set maximum time to retain the restarting neighbor’s or peer-group’s stale paths. Default is 360 seconds.

filter on AS-Path attribute

A BGP attribute, AS_PATH, can be used to manipulate routing policies. The AS_PATH attribute contains a sequence of AS numbers representing the route's path. As the route traverses an Autonomous System, the AS number is prepended to the route. You can manipulate routes based on their AS_PATH to affect interdomain routing. By identifying certain AS numbers in the AS_PATH, you can permit or deny routes based on the number in its AS_PATH.

To view all BGP path attributes in the BGP database, use the **show ip bgp paths** command in the EXEC privilege mode (Figure 345).

```
Force10#show ip bgp paths
Total 30655 Paths
Address       Hash Refcount Metric Path
0x4014154       0        3        18508 701 3549 19421 i
0x4013914       0        3        18508 701 7018 14990 i
0x5166d6c       0        3        18508 209 4637 1221 9249 9249 i
0x5e62df4       0        2        18508 701 17302 i
0x3a1814c       0       26        18508 209 22291 i
0x567ea9c       0       75        18508 209 3356 2529 i
0x6cc1294       0        2        18508 209 1239 19265 i
0x6cc18d4       0        1        18508 701 2914 4713 17935 i
0x5982e44       0      162        18508 209 i
0x67d4a14       0        2        18508 701 19878 ?
0x559972c       0       31        18508 209 18756 i
0x59cd3b4       0        2        18508 209 7018 15227 i
0x7128114       0       10        18508 209 3356 13845 i
0x536a914       0        3        18508 209 701 6347 7781 i
0x2ffe884       0        1        18508 701 3561 9116 21350 i
0x2ff7284       0       99        18508 701 1239 577 855 ?
0x2ff7ec4       0        4        18508 209 3561 4755 17426 i
0x2ff8544       0        3        18508 701 5743 2648 i
0x736c144       0        1        18508 701 209 568 721 1494 i
0x3b8d224       0       10        18508 209 701 2019 i
0x5eb1e44       0        1        18508 701 8584 16158 i
0x5cd891c       0        9        18508 209 6453 4759 i
--More--
```

Figure 345 show ip bgp paths Command Example

AS-PATH ACLs use regular expressions to search AS_PATH values. AS-PATH ACLs have an "implicit deny", that is, routes that do not meet a deny or match filter are dropped.

To configure an AS-PATH ACL to filter a specific AS_PATH value, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip as-path access-list** *as-path-name*	CONFIGURATION	Assign a name to a AS-PATH ACL and enter AS-PATH ACL mode.

Step	Command Syntax	Command Mode	Purpose
2	{**deny** | **permit**} *as-regular-expression*	AS-PATH ACL	Enter a regular expression to match BGP AS-PATH attributes. Use one or a combination of the following: • . = (period) matches on any single character, including white space • * = (asterisk) matches on sequences in a pattern (zero or more sequences) • + = (plus sign) matches on sequences in a pattern (one or more sequences) • ? = (question mark) matches sequences in a pattern (0 or 1 sequences). **You must enter an escape sequence (CNTL+v) prior to entering the ? regular expression.** • [] = (brackets) matches a range of single-character patterns. • ^ = (caret) matches the beginning of the input string. (If the caret is used at the beginning of a sequence or range, it matches on everything BUT the characters specified.) • $ = (dollar sign) matches the end of the output string. • _ = (underscore) matches a comma (,), left brace ({), right brace (}), left parenthesis, right parenthesis, the beginning of the input string, the end of the input string, or a space. • | = (pipe) matches either character.
3	**exit**	AS-PATH ACL	Return to CONFIGURATION mode
4	**router bgp** *as-number*	CONFIGURATION	Enter ROUTER BGP mode.
5	**neighbor** {*ip-address* | *peer-group-name*} **filter-list** *as-path-name* {**in** | **out**}	ROUTER BGP	Use a configured AS-PATH ACL for route filtering and manipulation. If you assign an non-existent or empty AS-PATH ACL, the software allows all routes.

To view the AS-PATH ACL configuration, use the **show config** command in the AS-PATH ACL mode and the **show ip as-path-access-list** command in the EXEC privilege mode (Figure 346).

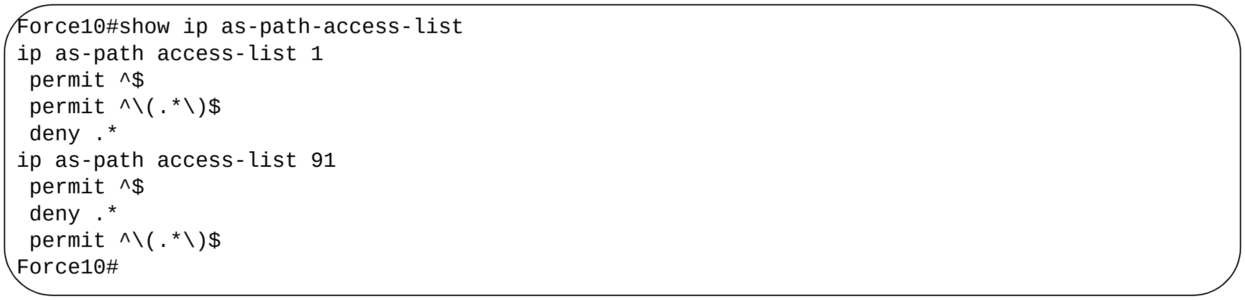

```
Force10#show ip as-path-access-list
ip as-path access-list 1
 permit ^$
 permit ^\(.*\)$
 deny .*
ip as-path access-list 91
 permit ^$
 deny .*
 permit ^\(.*\)$
Force10#
```

Figure 346 show ip as-path-access-list Command Example

For more information on this command and route filtering, refer to filter BGP routes on page 547.

configure IP community lists

Within FTOS, you have multiple methods of manipulating routing attributes. One attribute you can manipulate is the COMMUNITY attribute. This attribute is an optional attribute that is defined for a group of destinations. In FTOS, you can assign a COMMUNITY attribute to BGP routers by using an IP Community list. After you create an IP Community list, you can apply routing decisions to all routers meeting the criteria in the IP Community list.

IETF RFC 1997 defines the COMMUNITY attribute and the pre-defined communities of INTERNET, NO_EXPORT_SUBCONFED, NO_ADVERTISE, and NO_EXPORT. All BGP routes belong to the INTERNET community. In the RFC, the other communities are defined as follows:

- All routes with the NO_EXPORT_SUBCONFED (0xFFFFFF03) community attribute are not sent to CONFED-EBGP or EBGP peers, but are sent to IBGP peers within CONFED-SUB-AS.
- All routes with the NO_ADVERTISE (0xFFFFFF02) community attribute must not be advertised.
- All routes with the NO_EXPORT (0xFFFFFF01) community attribute must not be advertised outside a BGP confederation boundary, but are sent to CONFED-EBGP and IBGP peers.

To configure an IP community list, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip community-list** *community-list-name*	CONFIGURATION	Create a Community list and enter the COMMUNITY-LIST mode.
2	{**deny** | **permit**} {*community-number* | **local-AS** | **no-advertise** | **no-export** | **quote-regexp** *regular-expression-list* | **regexp** *regular-expression*}	COMMUNITY-LIST	Configure a Community list by denying or permitting specific community numbers or types of community • *community-number:* use AA:NN format where AA is the AS number (2 bytes) and NN is a value specific to that autonomous system. • **local-AS**: routes with the COMMUNITY attribute of NO_EXPORT_SUBCONFED. • **no-advertise:** routes with the COMMUNITY attribute of NO_ADVERTISE. • **no-export:** routes with the COMMUNITY attribute of NO_EXPORT. • **quote-regexp:** followed by any number of regular expressions. The software applies all regular expressions in the list. • **regexp:** followed by a regular expression.

To view the configuration, use the **show config** command in the COMMUNITY-LIST mode or the **show ip community-lists** command in the EXEC privilege mode (Figure 347).

```
Force10#show ip community-lists
ip community-list standard 1
 deny 701:20
 deny 702:20
 deny 703:20
 deny 704:20
 deny 705:20
 deny 14551:20
 deny 701:112
 deny 702:112
 deny 703:112
 deny 704:112
 deny 705:112
 deny 14551:112
 deny 701:667
 deny 702:667
 deny 703:667
```

Figure 347 show ip community-lists Command Example

To use an IP Community list to filter routes, you must apply a **match community** filter to a route map and then apply that route map to a BGP neighbor or peer group. Use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**route-map** *map-name* [**permit** | **deny**] [*sequence-number*]	CONFIGURATION	Enter the ROUTE-MAP mode and assign a name to a route map.
2	**match community** *community-list-name* [**exact**]	ROUTE-MAP	Configure a match filter for all routes meeting the criteria in the IP Community list.
3	**exit**	ROUTE-MAP	Return to the CONFIGURATION mode.
4	**router bgp** *as-number*	CONFIGURATION	Enter the ROUTER BGP mode.
5	**neighbor** {*ip-address* | *peer-group-name*} **route-map** *map-name* {**in** | **out**}	ROUTER BGP	Apply the route map to the neighbor or peer group’s incoming or outgoing routes.

To view the BGP configuration, use the **show config** command in the ROUTER BGP mode. To view a route map configuration, use the **show route-map** command in the EXEC privilege mode.

To view which BGP routes meet an IP Community list’s criteria, use the **show ip bgp community-list** command in the EXEC privilege mode.

manipulate the COMMUNITY attribute

In addition to permitting or denying routes based on the values of the COMMUNITY attributes, you can manipulate the COMMUNITY attribute value and send the COMMUNITY attribute with the route information.

By default, FTOS does not send the COMMMUNITY attribute.

To send the COMMUNITY attribute to BGP neighbors, use the following command in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
neighbor {*ip-address* | *peer-group-name*} **send-community**	ROUTER BGP	Enable the software to send the router's COMMUNITY attribute to the BGP neighbor or peer group specified.

To view the BGP configuration, use the **show config** command in the ROUTER BGP mode.

If you want to remove or add a specific COMMUNITY number from a BGP path, you must create a route map with one or both of the following statements in the route map. Then apply that route map to a BGP neighbor or peer group. Use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**route-map** *map-name* [**permit** | **deny**] [*sequence-number*]	CONFIGURATION	Enter the ROUTE-MAP mode and assign a name to a route map.
2	**set comm-list** *community-list-name* **delete**	ROUTE-MAP	Configure a set filter to delete all COMMUNITY numbers in the IP Community list.
	set community {*community-number* | **local-as** | **no-advertise** | **no-export** | **none**}	ROUTE-MAP	Configure a Community list by denying or permitting specific community numbers or types of community • *community-number:* use AA:NN format where AA is the AS number (2 bytes) and NN is a value specific to that autonomous system. • **local-AS**: routes with the COMMUNITY attribute of NO_EXPORT_SUBCONFED and are not sent to EBGP peers. • **no-advertise:** routes with the COMMUNITY attribute of NO_ADVERTISE and are not advertised. • **no-export:** routes with the COMMUNITY attribute of NO_EXPORT. • **none:** remove the COMMUNITY attribute. • **additive:** add the communities to already existing communities.
3	**exit**	ROUTE-MAP	Return to the CONFIGURATION mode.

Step	Command Syntax	Command Mode	Purpose
4	**router bgp** *as-number*	CONFIGURATION	Enter the ROUTER BGP mode.
5	**neighbor** {*ip-address* | *peer-group-name*} **route-map** *map-name* {**in** | **out**}	ROUTER BGP	Apply the route map to the neighbor or peer group's incoming or outgoing routes.

To view the BGP configuration, use the **show config** command in the ROUTER BGP mode. To view a route map configuration, use the **show route-map** command in the EXEC privilege mode.

To view BGP routes matching a certain community number or pre-defined BGP community, use the **show ip bgp community** command in the EXEC privilege mode (Figure 348).

```
Force10>show ip bgp community
BGP table version is 3762622, local router ID is 10.114.8.48
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network             Next Hop            Metric        LocPrf    Weight  Path
* i 3.0.0.0/8           195.171.0.16                         100         0  209 701 80 i
*>i 4.2.49.12/30        195.171.0.16                         100         0  209 i
* i 4.21.132.0/23       195.171.0.16                         100         0  209 6461 16422 i
*>i 4.24.118.16/30      195.171.0.16                         100         0  209 i
*>i 4.24.145.0/30       195.171.0.16                         100         0  209 i
*>i 4.24.187.12/30      195.171.0.16                         100         0  209 i
*>i 4.24.202.0/30       195.171.0.16                         100         0  209 i
*>i 4.25.88.0/30        195.171.0.16                         100         0  209 3561 3908 i
*>i 6.1.0.0/16          195.171.0.16                         100         0  209 7170 1455 i
*>i 6.2.0.0/22          195.171.0.16                         100         0  209 7170 1455 i
*>i 6.3.0.0/18          195.171.0.16                         100         0  209 7170 1455 i
*>i 6.4.0.0/16          195.171.0.16                         100         0  209 7170 1455 i
*>i 6.5.0.0/19          195.171.0.16                         100         0  209 7170 1455 i
*>i 6.8.0.0/20          195.171.0.16                         100         0  209 7170 1455 i
*>i 6.9.0.0/20          195.171.0.16                         100         0  209 7170 1455 i
*>i 6.10.0.0/15         195.171.0.16                         100         0  209 7170 1455 i
```

Figure 348 show ip bgp community Command Example (Partial)

change MED attribute

By default, FTOS uses the MULTI_EXIT_DISC or MED attribute when comparing EBGP paths from the same AS.

To change how the MED attribute is used, use any or all of the following commands in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
bgp always-compare-med	ROUTER BGP	Enable MED comparison in the paths from neighbors with different ASs. By default, this comparison is not performed.

Command Syntax	Command Mode	Purpose
bgp bestpath compare med	ROUTER BGP	Enable MED comparison of paths learned from BGP confederations. By default, this comparison is not performed.

To view the nondefault values, use the **show config** command in the ROUTER BGP mode.

change LOCAL_PREFERENCE attribute

In FTOS, you can change the value of the LOCAL_PREFERENCE attribute.

To change the default values of this attribute for all routes received by the router, use the following command in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
bgp default local-preference *value*	ROUTER BGP	Change the LOCAL_PREF value. • *value* range: 0 to 4294967295 Default is 100.

To view BGP configuration, use the **show config** command in the ROUTER BGP mode or the **show running-config bgp** command in the EXEC privilege mode.

A more flexible method for manipulating the LOCAL_PREF attribute value is to use a route map.

To change the default value of the LOCAL_PREF attribute for specific routes, you must use these commands in the following sequence, starting the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**route-map** *map-name* [**permit** | **deny**] [*sequence-number*]	CONFIGURATION	Enter the ROUTE-MAP mode and assign a name to a route map.
2	**set local-preference** *value*	ROUTE-MAP	Change LOCAL_PREF value for routes meeting the criteria of this route map.
3	**exit**	ROUTE-MAP	Return to the CONFIGURATION mode.
4	**router bgp** *as-number*	CONFIGURATION	Enter the ROUTER BGP mode.
5	**neighbor** {*ip-address* | *peer-group-name*} **route-map** *map-name* {**in** | **out**}	ROUTER BGP	Apply the route map to the neighbor or peer group’s incoming or outgoing routes.

To view the BGP configuration, use the **show config** command in the ROUTER BGP mode. To view a route map configuration, use the **show route-map** command in the EXEC privilege mode.

change NEXT_HOP attribute

You can change how the NEXT_HOP attribute is used.

To change the how the NEXT_HOP attribute is used, use the following command in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
neighbor {*ip-address* | *peer-group-name*} **next-hop-self**	ROUTER BGP	Disable next hop processing and configure the router as the next hop for a BGP neighbor.

To view BGP configuration, use the **show config** command in the ROUTER BGP mode or the **show running-config bgp** command in the EXEC privilege mode.

You can also use route maps to change this and other BGP attributes. For example, you can include the following command in a route map to specify the next hop address:

Command Syntax	Command Mode	Purpose
set next-hop *ip-address*	ROUTE-MAP	Sets the next hop address.

change WEIGHT attribute

To change the how the WEIGHT attribute is used, use the following command in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
neighbor {*ip-address* | *peer-group-name*} **weight** *weight*	ROUTER BGP	Assign a weight to the neighbor connection. • *weight* range: 0 to 65535

To view BGP configuration, use the **show config** command in the ROUTER BGP mode or the **show running-config bgp** command in the EXEC privilege mode.

You can also use route maps to change this and other BGP attributes. For example, you can include the following command in a route map to specify the next hop address:

Command Syntax	Command Mode	Purpose
set weight *weight*	ROUTE-MAP	Sets weight for the route. • *weight* range: 0 to 65535

enable multipath

By default, the software allows one path to a destination. You can enable multipath to allow up to 16 parallel paths to a destination.

To allow more than one path, use the following command in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
maximum-paths {**ebgp** | **ibgp**} *number*	ROUTER BGP	Enable multiple parallel paths. • *number* range: 1 to 16.

The **show ip bgp** *network* command includes multipath information for that network.

filter BGP routes

Filtering routes allows you to implement BGP policies. You can use either IP prefix lists, route maps, AS-PATH ACLs or IP Community lists (via a route map) to control which routes are accepted and advertised by the BGP neighbor or peer group. Prefix lists filter routes based on route and prefix length, while AS-Path ACLs filter routes based on the Autonomous System number. Route maps can filter and set conditions, change attributes, and assign update policies.

With FTOS, you can create inbound and outbound policies. Each of the commands used for filtering, has **in** and **out** parameters that must be applied. In FTOS, the order of preference varies depending on whether the attributes are applied for inbound updates or outbound updates.

For inbound and outbound updates the order of preference is:

- prefix lists (using **neighbor distribute-list** command)
- AS-PATH ACLs (using **neighbor filter-list** command)
- route maps (using **neighbor route-map** command)

Prior to filtering BGP routes, you must create the prefix list, AS-PATH ACL, or route map to be used.

Refer to for configuration information on prefix lists, AS-PATH ACLs, and route maps.

Note: When you configure a new set of BGP policies, always reset the neighbor or peer group by entering the **clear ip bgp** command in the EXEC privilege mode.

To filter routes using prefix lists, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip prefix-list** *prefix-name*	CONFIGURATION	Create a prefix list and assign it a name.

Step	Command Syntax	Command Mode	Purpose
2	**seq** *sequence-number* {**deny** | **permit**} *ip-prefix* [**ge** *max-prefix-length*] [**le** *min-prefix-length*]	PREFIX LIST	Create multiple prefix list filters with a deny or permit action. Refer to for information on configuring prefix lists.
3	**exit**	PREFIX LIST	Return to the CONFIGURATION mode.
4	**router bgp** *as-number*	CONFIGURATION	Enter ROUTER BGP mode.
5	**neighbor** {*ip-address* | *peer-group-name*} **distribute-list** *prefix-list-name* {**in** | **out**}	ROUTER BGP	Filter routes based on the criteria in the configured prefix list. Configure the following parameters: • *ip-address* or *peer-group-name:* enter the neighbor's IP address or the peer group's name. • *prefix-list-name:* enter the name of a configured prefix list. • **in:** apply the prefix list to inbound routes. • **out:** apply the prefix list to outbound routes.

As a reminder, below are some rules concerning prefix lists:

- If the prefix list contains no filters, all routes are permitted.
- If none of the routes match any of the filters in the prefix list, the route is denied. This action is called an implicit deny. (If you want to forward all routes that do not match the prefix list criteria, you must configure a prefix list filter to permit all routes. For example, you could have the following filter as the last filter in your prefix list **permit 0.0.0.0/0 le 32**).
- Once a route matches a filter, the filter's action is applied. No additional filters are applied to the route.

To view the BGP configuration, use the **show config** command in the ROUTER BGP mode. To view a prefix list configuration, use the **show ip prefix-list detail** or **show ip prefix-list summary** commands in the EXEC privilege mode.

To filter routes using a route map, use these commands in the following sequence, starting in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**route-map** *map-name* [**permit** | **deny**] [*sequence-number*]	CONFIGURATION	Create a route map and assign it a name.
2	{**match** | **set**}	ROUTE-MAP	Create multiple route map filters with a match or set action. Refer to for information on configuring route maps.
3	**exit**	ROUTE-MAP	Return to the CONFIGURATION mode.

Step	Command Syntax	Command Mode	Purpose
4	**router bgp** *as-number*	CONFIGURATION	Enter ROUTER BGP mode.
	neighbor {*ip-address* | *peer-group-name*} **route-map** *map-name* {**in** | **out**}	ROUTER BGP	Filter routes based on the criteria in the configured route map. Configure the following parameters: • *ip-address* or *peer-group-name:* enter the neighbor's IP address or the peer group's name. • *map-name:* enter the name of a configured route map. • **in:** apply the route map to inbound routes. • **out:** apply the route map to outbound routes.

To view the BGP configuration, use the **show config** command in the ROUTER BGP mode. To view a route map configuration, use the **show route-map** command in the EXEC privilege mode.

To filter routes based on AS-PATH information, use these commands in the following sequence, beginning in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip as-path access-list** *as-path-name*	CONFIGURATION	Create a AS-PATH ACL and assign it a name.
2	{**deny** | **permit**} *as-regular-expression*	AS-PATH ACL	Create a AS-PATH ACL filter with a deny or permit action. Refer to for information on configuring AS-PATH ACLs.
3	**exit**	AS-PATH ACL	Return to the CONFIGURATION mode.
4	**router bgp** *as-number*	CONFIGURATION	Enter ROUTER BGP mode.
5	**neighbor** {*ip-address* | *peer-group-name*} **filter-list** *as-path-name* {**in** | **out**}	ROUTER BGP	Filter routes based on the criteria in the configured route map. Configure the following parameters: • *ip-address* or *peer-group-name:* enter the neighbor's IP address or the peer group's name. • *as-path-name:* enter the name of a configured AS-PATH ACL. • **in:** apply the AS-PATH ACL map to inbound routes. • **out:** apply the AS-PATH ACL to outbound routes.

To view which commands are configured, use the **show config** command in the ROUTER BGP mode and **show ip as-path-access-list** command in the EXEC privilege mode.

Include this filter **permit .*** in your AS-PATH ACL to forward all routes not meeting the AS-PATH ACL criteria.

configure BGP route reflectors

BGP route reflectors are intended for Autonomous Systems with a large mesh and they reduce the amount of BGP control traffic. With route reflection configured properly, IBGP routers are not fully meshed within a cluster but all receive routing information.

Configure clusters of routers where one router is a concentration router and others are clients who receive their updates from the concentration router.

To configure a route reflector, use the following commands in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
bgp cluster-id *cluster-id*	ROUTER BGP	Assign an ID to a router reflector cluster. You can have multiple clusters in an AS.
neighbor {*ip-address* | *peer-group-name*} **route-reflector-client**	ROUTER BGP	Configure the local router as a route reflector and the neighbor or peer group identified is the route reflector client.

To view a route reflector configuration, use the **show config** command in the ROUTER BGP mode or **show running-config bgp** in the EXEC privilege mode.

When you enable a route reflector, FTOS automatically enables route reflection to all clients. To disable route reflection between all clients in this reflector, use the **no bgp client-to-client reflection** command in the ROUTER BGP mode. All clients should be fully meshed before you disable route reflection.

aggregate routes

FTOS provides multiple ways to aggregate routes in the BGP routing table. At least one more-specific route of the aggregate must be in the routing table for the configured aggregate to become active.

To aggregate routes, use the following command in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
aggregate-address *ip-address mask* [**advertise-map** *map-name*] [**as-set**] [**attribute-map** *map-name*] [**summary-only**] [**suppress-map** *map-name*]	ROUTER BGP	Assign the IP address and mask of the prefix to be aggregated. Optional parameters are: • **advertise-map** *map-name*: to set filters for advertising an aggregate route • **as-set**: to generate path attribute information and include it in the aggregate. • **attribute-map** *map-name:* to modify attributes of the aggregate, except for the AS_PATH and NEXT_HOP attributes • **summary-only**: to advertise only the aggregate address. Specific routes will not be advertised • **suppress-map** *map-name*: to identify which more-specific routes in the aggregate are suppressed

AS_SET includes AS_PATH and community information from the routes included in the aggregated route.

In the **show ip bgp** command, aggregates contain an ‘a’ in the first column and routes suppressed by the aggregate contain an ‘s’ in the first column.

```
Force10#show ip bgp
BGP table version is 0, local router ID is 10.101.15.13
Status codes: s suppressed, d damped, h history, * valid, > best
Path source: I - internal, a - aggregate, c - confed-external, r - redistributed, n - network
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network            Next Hop            Metric     LocPrf Weight Path
*>  7.0.0.0/29         10.114.8.33              0                 0 18508 ?
*>  7.0.0.0/30         10.114.8.33              0                 0 18508 ?
*>a 9.0.0.0/8         192.0.0.0                              32768 18508 701 {7018 2686 3786} ?
*>  9.2.0.0/16         10.114.8.33                                0 18508 701 i
*>  9.141.128.0/24     10.114.8.33                                0 18508 701 7018 2686 ?
```

Figure 349 show ip bgp Command Example with Aggregates

configure BGP confederations

Another way to organize routers within an AS and reduce the mesh for IBGP peers is to configure BGP confederations. As with route reflectors, BGP confederations are recommended only for IBGP peering involving a large number of IBGP peering sessions per router. Basically, when you configure BGP confederations, you break the AS into smaller sub-AS, and to those outside your network, the confederations appear as one AS. Within the confederation sub-AS, the IBGP neighbors are fully meshed and the MED, NEXT_HOP, and LOCAL_PREF attributes are maintained between confederations.

To configure BGP confederations, use the following commands in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
bgp confederation identifier *as-number*	ROUTER BGP	Specifies the confederation ID. Use your public AS number.
bgp confederation peers *as-number* [... *as-number*]	ROUTER BGP	Specifies which confederation sub-AS are peers.

To view the configuration, use the **show config** command in the ROUTER BGP mode.

enable route flap dampening

When EBGP routes become unavailable, they “flap” and the router issues both WITHDRAWN and UPDATE notices. A flap is when a route

- is withdrawn
- is readvertised after being withdrawn
- has an attribute change

The constant router reaction to the WITHDRAWN and UPDATE notices causes instability in the BGP process. To minimize this instability, you may configure penalties, a numeric value, for routes that flap. When that penalty value reaches a configured limit, the route is not advertised, even if the route is up. In FTOS, that penalty value is 1024. As time passes and the route does not flap, the penalty value decrements or is decayed. However, if the route flaps again, it is assigned another penalty.

When dampening is applied to a route, its path is described by one of the following terms:

- history entry—an entry that stores information on a downed route
- dampened path—a path that is no longer advertised
- penalized path—a path that is assigned a penalty

The CLI example below shows configuring values to start reusing or restarting a route, as well as their default values:

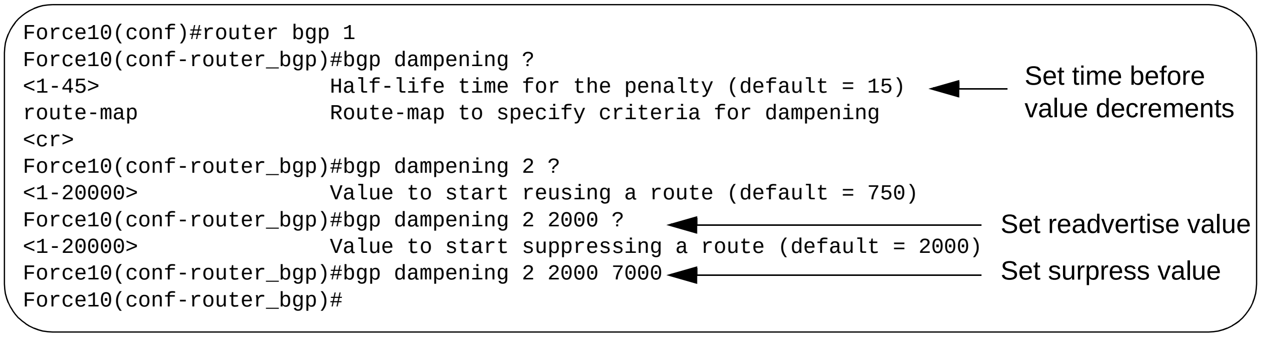

```
Force10(conf)#router bgp 1
Force10(conf-router_bgp)#bgp dampening ?
<1-45>                  Half-life time for the penalty (default = 15)
route-map               Route-map to specify criteria for dampening
<cr>
Force10(conf-router_bgp)#bgp dampening 2 ?
<1-20000>               Value to start reusing a route (default = 750)
Force10(conf-router_bgp)#bgp dampening 2 2000 ?
<1-20000>               Value to start suppressing a route (default = 2000)
Force10(conf-router_bgp)#bgp dampening 2 2000 7000
Force10(conf-router_bgp)#
```

Figure 350 Setting Reuse and Restart Route Values

To configure route flap dampening parameters, use the following command in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
bgp dampening [*half-life reuse suppress max-suppress-time*] [**route-map** *map-name*]	ROUTER BGP	Enable route dampening. Enter the following optional parameters to configure route dampening parameters: • *half-life* range: 1 to 45. Number of minutes after which the Penalty is decreased. After the router assigns a Penalty of 1024 to a route, the Penalty is decreased by half after the half-life period expires. (Default: 15 minutes) • *reuse* range: 1 to 20000. This number is compared to the flapping route's Penalty value. If the Penalty value is less than the reuse value, the flapping route is once again advertised (or no longer suppressed). Withdrawn routes are removed from history state. (Default: 750) • *suppress* range: 1 to 20000. This number is compared to the flapping route's Penalty value. If the Penalty value is greater than the suppress value, the flapping route is no longer advertised (that is, it is suppressed). (Default: 2000.) • *max-suppress-time* range: 1 to 255. The maximum number of minutes a route can be suppressed. The default is four times the half-life value. (Default: 60 minutes.) • **route-map** *map-name:* name of a configured route map. Only match commands in the configured route map are supported. Use this parameter to apply route dampening to selective routes.

To view the BGP configuration, use **show config** in the ROUTER BGP mode or **show running-config bgp** in the EXEC privilege mode.

To view a count of dampened routes, history routes and penalized routes when route dampening is enabled, look at the seventh line of the **show ip bgp summary** command output (Figure 351).

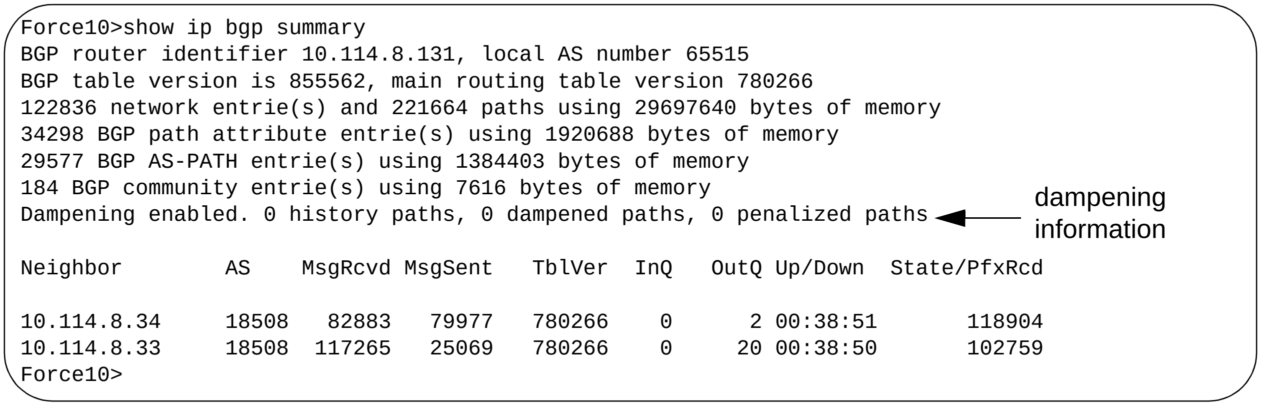

```
Force10>show ip bgp summary
BGP router identifier 10.114.8.131, local AS number 65515
BGP table version is 855562, main routing table version 780266
122836 network entrie(s) and 221664 paths using 29697640 bytes of memory
34298 BGP path attribute entrie(s) using 1920688 bytes of memory
29577 BGP AS-PATH entrie(s) using 1384403 bytes of memory
184 BGP community entrie(s) using 7616 bytes of memory
Dampening enabled. 0 history paths, 0 dampened paths, 0 penalized paths

Neighbor        AS    MsgRcvd MsgSent   TblVer  InQ   OutQ Up/Down  State/PfxRcd

10.114.8.34     18508   82883   79977   780266    0      2 00:38:51       118904
10.114.8.33     18508  117265   25069   780266    0     20 00:38:50       102759
Force10>
```

Figure 351 show ip bgp summary Command Example

To view which routes are dampened (non-active), use the **show ip bgp dampened-routes** command in the EXEC privilege mode.

To clear information on route dampening and return suppressed routes to active state, use the following command in the EXEC privilege mode:

Command Syntax	Command Mode	Purpose
clear ip bgp dampening [*ip-address mask*]	EXEC privilege	Clear all information or only information on a specific route.

To view statistics on route flapping, use the following command in the EXEC and EXEC privilege mode:

Command Syntax	Command Mode	Purpose
show ip bgp flap-statistics [*ip-address* [*mask*]] [**filter-list** *as-path-name*] [**regexp** *regular-expression*]	EXEC EXEC privilege	View all flap statistics or for specific routes meeting the following criteria: • *ip-address* [*mask*]: enter the IP address and mask • **filter-list** *as-path-name:* enter the name of an AS-PATH ACL. • **regexp** *regular-expression:* enter a regular express to match on.

change path selection to non-deterministic

By default, the path selection in FTOS is deterministic, that is, paths are compared irrespective of the order of their arrival. You can change the path selection method to non-deterministic, that is, paths are compared in the order in which they arrived (starting with the most recent). Furthermore, in non-deterministic mode, the software may not compare MED attributes though the paths are from the same AS.

To change the path selection from the default mode (deterministic) to non-deterministic, use the following command in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
bgp non-deterministic-med	ROUTER BGP	Change the best path selection method to non-deterministic.

Note: When you change the best path selection method, path selection for existing paths remains unchanged until you reset it by entering the **clear ip bgp** command in the EXEC privilege mode.

change BGP timers

To configure BGP timers, use either or both of the following commands in the ROUTER BGP mode:

Command Syntax	Command Mode	Purpose
neighbors {*ip-address* | *peer-group-name*} **timers** *keepalive holdtime*	ROUTER BGP	Configure timer values for a BGP neighbor or peer group. • *keepalive* range: 1 to 65535. Time interval, in seconds, between keepalive messages sent to the neighbor routers. (Default: 60 seconds) • *holdtime* range: 3 to 65536. Time interval, in seconds, between the last keepalive message and declaring the router dead. (Default: 180 seconds)
timers bgp *keepalive holdtime*	ROUTER BGP	Configure timer values for all neighbors. • *keepalive* range: 1 to 65535. Time interval, in seconds, between keepalive messages sent to the neighbor routers. (Default: 60 seconds) • *holdtime* range: 3 to 65536. Time interval, in seconds, between the last keepalive message and declaring the router dead. (Default: 180 seconds)

To view non-default values, enter the **show config** command in the ROUTER BGP mode or the **show running-config bgp** command in the EXEC privilege mode.

Timer values configured with the **neighbor timers** command override the timer values configured with the **timers bgp** command.

When two neighbors, configured with different *keepalive* and *holdtime* values, negotiate for new values, the resulting values will be as follows:

- the lower of the *holdtime* values is the new *holdtime* value, and
- whichever is the lower value; one-third of the new *holdtime* value, or the configured *keepalive* value is the new *keepalive* value.

Debugging BGP

To enable BGP debugging, use any of the commands in the EXEC privilege mode:

Command Syntax	Command Mode	Purpose
debug ip bgp [*ip-address* | **peer-group** *peer-group-name*] [**in** | **out**]	EXEC privilege	View all information on BGP, including BGP events, keepalives, notifications, and updates.
debug ip bgp dampening [**in** | **out**]	EXEC privilege	View information on BGP route being dampened.

Command Syntax	Command Mode	Purpose
debug ip bgp [*ip-address* | **peer-group** *peer-group-name*] **events** [**in** | **out**]	EXEC privilege	View information on local BGP state changes and other BGP events.
debug ip bgp [*ip-address* | **peer-group** *peer-group-name*] **keepalive** [**in** | **out**]	EXEC privilege	View information about BGP KEEPALIVE messages.
debug ip bgp [*ip-address* | **peer-group** *peer-group-name*] **notifications** [**in** | **out**]	EXEC privilege	View information about BGP notifications received from or sent to neighbors.
debug ip bgp [*ip-address* | **peer-group** *peer-group-name*] **updates** [**in** | **out**] [**prefix-list name**]	EXEC privilege	View information about BGP updates and filter by prefix name
debug ip bgp {*ip-address* | *peer-group-name*} **soft-reconfiguration**	EXEC privilege	Enable soft-reconfiguration debug.Enable soft-reconfiguration debug.

FTOS displays debug messages on the console. To view which debugging commands are enabled, use the **show debugging** command in the EXEC privilege mode.

To disable a specific debug command, enter the keyword no followed by the debug command. For example, to disable debugging of BGP updates, you enter **no debug ip bgp updates** command.

To disable all BGP debugging, enter **no debug ip bgp**.

To disable all debugging, enter **undebug all**.

Storing Last and Bad PDUs

FTOS stores the last notification sent/received, and the last bad PDU received on per peer basis. The last bad PDU is the one that causes a notification to be issued. Thes PDUs are shown in the output of the command **show ip bgp neighbor**, as shown in Figure 352.

Figure 352 Viewing the Last Bad PDU from BGP Peers

```
Force10(conf-router_bgp)#do show ip bgp neighbors 1.1.1.2

BGP neighbor is 1.1.1.2, remote AS 2, external link
  BGP version 4, remote router ID 2.4.0.1
  BGP state ESTABLISHED, in this state for 00:00:01
  Last read 00:00:00, last write 00:00:01
  Hold time is 90, keepalive interval is 30 seconds
  Received 1404 messages, 0 in queue
    3 opens, 1 notifications, 1394 updates
    6 keepalives, 0 route refresh requests
  Sent 48 messages, 0 in queue
    3 opens, 2 notifications, 0 updates
    43 keepalives, 0 route refresh requests
  Minimum time between advertisement runs is 30 seconds
  Minimum time before advertisements start is 0 seconds

  Capabilities received from neighbor for IPv4 Unicast :
    MULTIPROTO_EXT(1)
    ROUTE_REFRESH(2)
    CISCO_ROUTE_REFRESH(128)

  Capabilities advertised to neighbor for IPv4 Unicast :
    MULTIPROTO_EXT(1)
    ROUTE_REFRESH(2)
    CISCO_ROUTE_REFRESH(128)

  For address family: IPv4 Unicast
  BGP table version 1395, neighbor version 1394
  Prefixes accepted 1 (consume 4 bytes), 0 withdrawn by peer
  Prefixes advertised 0, rejected 0, 0 withdrawn from peer

  Connections established 3; dropped 2
  Last reset 00:00:12, due to Missing well known attribute

  Notification History
   'UPDATE error/Missing well-known attr' Sent : 1  Recv: 0
   'Connection Reset' Sent : 1  Recv: 0

   Last notification (len 21) sent 00:26:02 ago
    ffffffff ffffffff ffffffff ffffffff 00160303 03010000
   Last notification (len 21) received 00:26:20 ago
    ffffffff ffffffff ffffffff ffffffff 00150306 00000000
   Last PDU (len 41) received 00:26:02 ago that caused notification to be issued
    ffffffff ffffffff ffffffff ffffffff 00290200 00000e01 02040201 00024003 04141414
0218c0a8
    01000000
Local host: 1.1.1.1, Local port: 179
Foreign host: 1.1.1.2, Foreign port: 41758
```

Capturing PDUs

Capture incoming and outgoing PDUs on a per-peer basis using the command **capture bgp-pdu neighbor direction.** Disable capturing using the no form of this command.

The buffer size supports a maximum value between 40 MB (the default) and 100 MB. The capture buffers are cyclic, and reaching the limit prompts the system to overwrite the oldest PDUs when new ones are received for a given neighbor or direction. Setting the buffer size to a value lower than the current max, might cause captured PDUs to be freed to set the new limit.

Note: Memory on RP1 is not pre-allocated, but rather is allocated only when a PDU needs to be captured.

To change the maximum buffer size use the command **capture bgp-pdu max-buffer-size** (Figure 353). View the captured PDUs using the command **show capture bgp-pdu neighbor**

Figure 353 Viewing Captured PDUsr

```
Force10#show capture bgp-pdu neighbor 20.20.20.2

Incoming packet capture enabled for BGP neighbor 20.20.20.2
Available buffer size 40958758, 26 packet(s) captured using 680 bytes
  PDU[1] : len 101, captured 00:34:51 ago
    ffffffff ffffffff ffffffff ffffffff 00650100 00000013 00000000 00000000 419ef06c
00000000
    00000000 00000000 00000000 00000000 0181a1e4 0181a25c 41af92c0 00000000 00000000
00000000
    00000000 00000001 0181a1e4 0181a25c 41af9400 00000000
  PDU[2] : len 19, captured 00:34:51 ago
    ffffffff ffffffff ffffffff ffffffff 00130400
  PDU[3] : len 19, captured 00:34:51 ago
    ffffffff ffffffff ffffffff ffffffff 00130400
  PDU[4] : len 19, captured 00:34:22 ago
    ffffffff ffffffff ffffffff ffffffff 00130400
[. . .]

Outgoing packet capture enabled for BGP neighbor 20.20.20.2
Available buffer size 40958758, 27 packet(s) captured using 562 bytes
  PDU[1] : len 41, captured 00:34:52 ago
    ffffffff ffffffff ffffffff ffffffff 00290104 000100b4 14141401 0c020a01 04000100
01020080
    00000000
  PDU[2] : len 19, captured 00:34:51 ago
    ffffffff ffffffff ffffffff ffffffff 00130400
  PDU[3] : len 19, captured 00:34:50 ago
    ffffffff ffffffff ffffffff ffffffff 00130400
  PDU[4] : len 19, captured 00:34:20 ago
    ffffffff ffffffff ffffffff ffffffff 00130400
[. . .]
```

The buffers storing the PDUs free memory when:

- BGP is disabled.

- A neighbor is unconfigured.
- **clear ip bgp** is issued.
- New PDU are captured and there is no more space to store them.
- The max buffer size is reduced. (This may causes PDUs to be cleared depending upon the buffer space consumed and the new limit.)

With full internet feed (205K) captured, approximately 11.8MB is required to store all of the PDUs, as shown in Figure 354.

Figure 354 Required Memory for Captured PDUs

```
Force10(conf-router_bgp)#do show capture bgp-pdu neighbor 172.30.1.250

Incoming packet capture enabled for BGP neighbor 172.30.1.250
Available buffer size 29165743, 192991 packet(s) captured using 11794257 bytes
 [. . .]

Force10(conf-router_bgp)#do sho ip bg s
BGP router identifier 172.30.1.56, local AS number 65056
BGP table version is 313511, main routing table version 313511
207896 network entrie(s) and 207896 paths using 42364576 bytes of memory
59913 BGP path attribute entrie(s) using 2875872 bytes of memory
59910 BGP AS-PATH entrie(s) using 2679698 bytes of memory
3 BGP community entrie(s) using 81 bytes of memory

Neighbor        AS     MsgRcvd  MsgSent     TblVer  InQ  OutQ Up/Down  State/Pfx

1.1.1.2         2           17    18966          0    0     0 00:08:19 Active
172.30.1.250    18508   243295       25     313511    0     0 00:12:46    207896
```

PDU Counters

FTOS version 7.5.1.0 introduces additional counters for various types of PDUs sent and received from neighbors. These are seen in the output of the command **show ip bgp neighbor**.

MBGP Configuration

Support for different address families is advertised to BGP neighbors through a capability advertisement. When BGP is configured, the support for IPv4 Multicast is turned off by default. IPv4 Multicast is enabled using the commands outlined below.

FTOS MBGP is implemented as per IETF RFC 1858. The MBGP feature can be enabled per router and/or per peer/peer-group. Default is IPv4 Unicast routes.

Command Syntax	Command Mode	Purpose
address family ipv4 multicast	ROUTER BGP (conf-router_bgp)	Enables support for the IPv4 Multicast family on the BGP node

Command Syntax	Command Mode	Purpose
neighbor [*ip-address* | *peer-group-name*] **activate**	ROUTER BGP Address Family (conf-router_bgp_af)	Enable IPv4 Multicast support on a BGP neighbor/peer group

When a peer is configured to support IPv4 Multicast, FTOS takes the following actions:

- Send a capacity advertisement to the peer in the BGP Open message specifying IPv4 Multicast as a supported AFI/SAFI (Subsequent Address Family Identifier).
- If the corresponding capability is received in the peer's Open message, BGP will mark the peer as supporting the AFI/SAFI.
- When exchanging updates with the peer, BGP sends and receives IPv4 Multicast routes if the peer is marked as supporting that AFI/SAFI.
- Exchange of IPv4 Multicast route information occurs through the use of two new attributes called MP_REACH_NLRI and MP_UNREACH_NLRI, for feasible and withdrawn routes, respectively.
- If the peer has not been activated in any AFI/SAFI, the peer remains in Idle state.

Most FTOS BGP IPv4 Unicast commands are extended to support the IPv4 Multicast RIB using extra options to the command. See the *FTOS Command Line Interface Reference* for a detailed description of the MBGP commands.

BGP4 MIB

The FORCE10-BGP4-V2-MIB enhances FTOS BGP MIB support with many new SNMP objects and notifications (traps) defined in the *draft-ietf-idr-bgp4-mibv2-05*. To see these enhancements, download the MIB from the Force10 website, www.force10networks.com.

Note: See the Force10 iSupport webpage for the *Force10-BGP4-V2-MIB* and othe MIB documentation.

Important Points to Remember

- In f10BgpM2AsPathTableEntry table, f10BgpM2AsPathSegmentIndex, and f10BgpM2AsPathElementIndex are used to retrieve a particular ASN from the AS path. These indices are assigned to the AS segments and individual ASN in each segment starting from 0. For example, an AS path list of {200 300 400} 500 consists of two segments: {200 300 400} with segment index 0 and 500 with segment index 1. ASN 200, 300, and 400 are be assigned 0, 1, and 2 element indices in that order.
- Unknown optional transitive attributes within a given path attribute (PA) are assigned indices in order. These indices correspond to f10BgpM2PathAttrUnknownIndex field in the f10BgpM2PathAttrUnknownEntry table.
- Negotiation of multiple instances of the same capability is not supported. F10BgpM2PeerCapAnnouncedIndex and f10BgpM2PeerCapReceivedIndex are ignored in the peer capability lookup.

- Inbound BGP soft-reconfiguration must be configured on a peer for f10BgpM2PrefixInPrefixesRejected to display the number of prefixes filtered due to a policy. If BGP soft-reconfig is not enabled, the denied prefixes are not accounted for.
- F10BgpM2AdjRibsOutRoute stores the pointer to the NLRI in the peer's Adj-Rib-Out.
- PA Index (f10BgpM2PathAttrIndex field in various tables) is used to retrieve specific attributes from the PA table. The Next-Hop, RR Cluster-list, Originator ID attributes are not stored in the PA Table and cannot be retrieved using the index passed in. These fields are not populated in f10BgpM2PathAttrEntry, f10BgpM2PathAttrClusterEntry, f10BgpM2PathAttrOriginatorIdEntry.
- F10BgpM2PathAttrUnknownEntry contains the optional-transitive attribute details.
- Query for f10BgpM2LinkLocalNextHopEntry returns default value for Link-local Next-hop.
- RFC 2545 and the f10BgpM2Rfc2545Group are not supported.
- An SNMP query will display up to 89 AS paths. A query for a larger AS path count will display as "…" at the end of the output.
- SNMP set for BGP is not supported. For all peer configuration tables (f10BgpM2PeerConfigurationGroup, f10BgpM2PeerRouteReflectorCfgGroup, and f10BgpM2PeerAsConfederationCfgGroup), an SNMP set operation will return an error. Only SNMP queries are supported. In addition, the f10BgpM2CfgPeerError, f10BgpM2CfgPeerBgpPeerEntry, and f10BgpM2CfgPeerRowEntryStatus fields are to hold the SNMP set status and are ignored in SNMP query.
- The AFI/SAFI is not used as an index to the f10BgpM2PeerCountersEntry table. The BGP peer's AFI/SAFI (IPv4 Unicast or IPv6 Multicast) is used for various outbound counters. Counters corresponding to IPv4 Multicast cannot be queried.
- The f10BgpM2[Cfg]PeerReflectorClient field is populated based on the assumption that route-reflector clients are not in a full mesh if BGP client-2-client reflection is enabled and that the BGP speaker acting as reflector will advertise routes learned from one client to another client. If disabled, it is assumed that clients are in a full mesh, and there is no need to advertise prefixes to the other clients.
- High CPU utilization may be observed during an SNMP walk of a large BGP Loc-RIB.
- To avoid SNMP timeouts with a large-scale configuration (large number of BGP neighbors and a large BGP Loc-RIB), Force10 recommends setting the timeout and retry count values to a relatively higher number. e.g. t = 60 or r = 5.
- To return all values on an snmpwalk for the f10BgpM2Peer sub-OID, use the -C c option, such as snmpwalk -v 2c -C c -c public <IP_address> <OID>.
- An SNMP walk may terminate pre-maturely if the index does not increment lexicographically. Force10 recommends using options to ignore such errors.
- Multiple BPG process instances are not supported. Thus, the F10BgpM2PeerInstance field in various tables is not used to locate a peer.
- Multiple instances of the same NLRI in the BGP RIB are not supported and are set to zero in the SNMP query response.
- F10BgpM2NlriIndex and f10BgpM2AdjRibsOutIndex fields are not used.
- Carrying MPLS labels in BGP is not supported. F10BgpM2NlriOpaqueType and f10BgpM2NlriOpaquePointer fields are set to zero.
- 4-byte ASN is not supported. The f10BgpM2AsPath4byteEntry table contains default values.
- Extended-Communities are not supported. The f10BgpM2PathAttrExtCommEntry is populated with default values.

- Traps (notifications) specified in the BGP4 MIB draft <draft-ietf-idr-bgp4-mibv2-05.txt> are not supported. Such traps (bgpM2Established and bgpM2BackwardTransition) are supported as part of RFC 1657.

Chapter 33 Multicast Protocols

C-Series	NO
E-Series	✓

Platform Specific Feature: Multicast Protocols are supported on E-Series only.

FTOS supports IP multicast and Internet Group Management Protocol (IGMP) and Protocol Independent Multicast—Sparse Mode (PIM-SM) protocols. This chapter includes the following sections:

- IP Multicast on page 563
- IGMP version 2 on page 564
- IGMP version 3 on page 568
- IGMP Snooping on page 568
- PIM-Sparse Mode (Version 2) on page 573
- Multicast Traceroute on page 579
- PIM Source Specific Multicast on page 579

IP Multicast

Prior to enabling the multicast protocols on an interface or configuring a static Rendezvous Point (RP), you must enable IP multicasting on the E-Series.

Note: Multicast is not supported on secondary IP addresses.

To enable IP multicast routing on an E-Series, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ip multicast-routing	CONFIGURATION	Enable multicast routing.

➡ **Note:** Egress L3 ACL is not applied to multicast data traffic if multicast routing is turned on.

Protocol Control Traffic Redirected Through MAC

Protocol control traffic in FTOS is redirected through the MAC address. In an InterVLAN scenario, certain types of multicast traffic may hit the CPU in addition to normal Layer 2 flooding, since multiple Multicast IP addresses and Layer 2 traffic both map to the same MAC address. For example, 224.0.0.5 is a well known IP address for OSPF that maps to the multicast MAC address 01:00:5e:00:00:05. The Layer 2 FIB alone can not differentiate multicast control traffic, such as OSPF or RIPv2, from certain multicast data traffic. Since addresses such as 224.0.0.5, 225.0.0.5, 226.0.0.5, etc. all map to this same multicast MAC address, the data traffic and OSPF traffic hit the same entry and are forwarded to the CPU. Therefore, Force10 recommends to avoid using those multicast IP address that map to well-known MAC addresses for data transmission.

As the upper five bits of an IP Multicast address are dropped in the translation, 32 different multicast group IDs all map to the same Ethernet address. For example, when the user uses IP address 225.0.0.5 that maps to the same multicast MAC address 01:00:5e:00:00:05, the traffic is treated as an OSPF multicast entry and is also sent to the CPU.

Here are well known MAC addresses that are used in the system:

- OSPF 01:00:5e:00:00:05
- OSPF 01:00:5e:00:00:06
- RIP 01:00:5e:00:00:09
- NTP 01:00:5e:00:01:01
- VRRP 01:00:5e:00:00:12
- PIMSM 01:00:5e:00:00:0d

IGMP version 2

Multicast routers use Internet Group Management Protocol (IGMP) to learn which groups have members on each of their attached physical networks. A multicast router keeps a list of multicast group memberships (that is, the presence of at least one member of a multicast group) for each attached network, and a timer for each membership.

In IGMP, a multicast router is either a Querier or not. Queriers are routers with the lowest IP address of the multicast routers on an attached network. When a router receives an IGMP Membership Report for a group, it adds that group to the list of multicast group memberships on the network on which it received the report.

For more information on the protocol, refer to RFC 2236 *Internet Group Management Protocol*.

Implementation Information

The E-Series cannot serve as an IGMP host, but may support IGMP version 1 hosts. The E-Series does not support IGMP version 1 router.

The FTOS implementation of IGMP is based on IETF RFC 2236.

Configuration Tasks for IGMP

The following list includes the configuration tasks for IGMP:

- enable IGMP on an interface on page 565 (mandatory)
- configure static IGMP-group on page 566 (optional)
- adjust timers on page 567 (optional)

For a complete listing of all commands related to IGMP, refer to .

enable IGMP on an interface

When you enter the **ip multicast-routing** command, you enable IP Multicast on the E-Series, however, PIM and IGMP are not enabled on any interfaces.

To enable IGMP and PIM on an interface, use these commands in the following sequence, beginning in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip multicast-routing**	CONFIGURATION	Enable multicast routing on the system If you have previously entered this command, skip this step.
2	**interface** *interface*	CONFIGURATION	Specify the physical interface type, slot, and number. • For a Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a Loopback interface, enter the keyword **loopback** followed by a number from 1 to 16383. • For a SONET interface, enter the keyword **sonet** followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information.
3	**ip address** *ip-address mask*	INTERFACE	Assign an IP address to the interface. • *ip-address:* enter an address and mask.
4	**no shutdown**	INTERFACE	Enable the interface.

Step	Command Syntax	Command Mode	Purpose
5	**ip pim sparse-mode**	INTERFACE	Enable IGMP and PIM on an interface.

To view which interfaces are IGMP-enabled, enter the **show ip igmp interface** command in the EXEC privilege mode.

```
Force10#show ip igmp interface
GigabitEthernet 7/16 is up, line protocol is up
  Internet address is 10.87.3.2/24
  IGMP is enabled on interface
  IGMP query interval is 60 seconds
  IGMP querier timeout is 300 seconds
  IGMP max query response time is 10 seconds
  Last member query response interval is 199 ms
  IGMP activity: 0 joins, 0 leaves
  IGMP querying router is 10.87.3.2 (this system)
  IGMP version is 2
Force10#
```

Figure 355 show ip igmp interface Command Example

configure static IGMP-group

To configure a static IGMP group, use the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip igmp static-group *group-address*	INTERFACE	Assign a multicast group address to an interface. • *group-address:* enter a multicast group address.

Note: A static IGMP group never expires.

To view both learned and statically configured IGMP groups, use the **show ip igmp groups** command in the EXEC privilege mode.

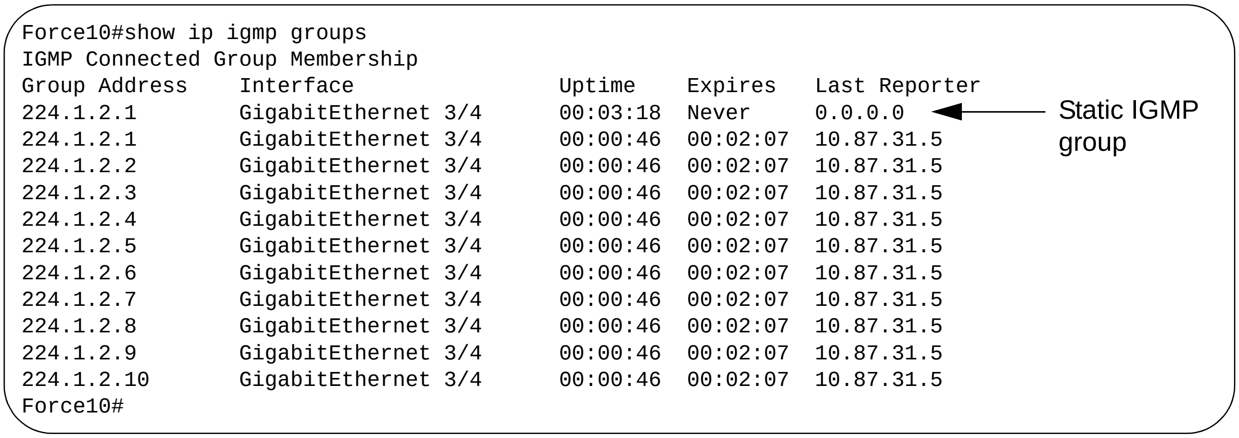

```
Force10#show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface              Uptime    Expires   Last Reporter
224.1.2.1        GigabitEthernet 3/4    00:03:18  Never     0.0.0.0
224.1.2.1        GigabitEthernet 3/4    00:00:46  00:02:07  10.87.31.5
224.1.2.2        GigabitEthernet 3/4    00:00:46  00:02:07  10.87.31.5
224.1.2.3        GigabitEthernet 3/4    00:00:46  00:02:07  10.87.31.5
224.1.2.4        GigabitEthernet 3/4    00:00:46  00:02:07  10.87.31.5
224.1.2.5        GigabitEthernet 3/4    00:00:46  00:02:07  10.87.31.5
224.1.2.6        GigabitEthernet 3/4    00:00:46  00:02:07  10.87.31.5
224.1.2.7        GigabitEthernet 3/4    00:00:46  00:02:07  10.87.31.5
224.1.2.8        GigabitEthernet 3/4    00:00:46  00:02:07  10.87.31.5
224.1.2.9        GigabitEthernet 3/4    00:00:46  00:02:07  10.87.31.5
224.1.2.10       GigabitEthernet 3/4    00:00:46  00:02:07  10.87.31.5
Force10#
```

Static IGMP group

Figure 356 show ip igmp groups Command Example

adjust timers

Routers periodically send a General Query on each attached network for which this router is the Querier. A General Query is addressed to the all-system multicast group (224.0.0.1). In FTOS, you can adjust the frequency of the General Query and other messages.

To adjust IGMP timers, use any of the following commands in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip igmp last-member-query-interval *milliseconds*	INTERFACE	Change the last member query interval. • *milliseconds:* enter a number from 100 to 65535. Default: 1000 milliseconds
ip igmp querier-timeout *seconds*	INTERFACE	Change the interval that must pass before a multicast router decides there is no longer another Querier. • *seconds:* enter a number from 60 to 300. Default: 120 seconds
ip igmp query-interval *seconds*	INTERFACE	Change the transmission frequency of IGMP general queries. • *seconds:* enter a number from 1 to 18000. Default: 60 seconds.
ip igmp query-max-resp-time *seconds*	INTERFACE	Change the maximum query response time advertised in the General Query. • *seconds:* enter a number from 1 to 25. Default: 10 seconds

To view the current IGMP timers settings, use the **show ip igmp interface** command in the EXEC privilege mode.

IGMP version 3

The E-Series supports IGMP v3 as specified in RFC 3376. The FTOS implementation of IGMPv3 does not support interoperability with IGMPv1 and v2 routers on the same subnet. There is no limit on the number of supported groups. Up to 512 interfaces can support IGMPv3.

To enable IGMPv3 on an interface, use the same **configuration**, **show**, and **debug** commands as IGMPv1 and v2.

The E-Series supports include and exclude modes for PIM-SM and PIM-SSM. Use the following command beginning in CONFIGURATION node to configure the mode:

Command Syntax	Command Mode	Purpose
ip igmp static-group [include | exclude]	INTERFACE	Assign a multicast group address to an interface.

➡ **Note:** A static IGMP group never expires.

IGMP Snooping

IGMP snooping enables the switch to constrain IP multicast traffic at Layer-2 by forwarding traffic only to interested receivers. IGMP snooping optimizes the usage of network bandwidth. The switch does not reduce any IP multicast traffic in the local multicast domain (224.0.0.0/24).

An IGMP-enabled switch listens in on the IGMP frames between hosts and routers. When a switch receives a Membership Report from a host for a given multicast group, it adds the host's interface to the OIF (outgoing interface) list for that multicast group in the VLAN. Similarly, when a host sends a "Leave Group" message the switch removes the interface number from the OIF list for that multicast group.

The following list includes the following sections on IGMP snooping:

- Joining a Multicast Group on page 569
- Leaving Multicast Group on page 569
- IGMP Snooping Querier Functionality on page 570
- Fast Convergence after MSTP-Triggered Topology Changes on page 570
- Multicast Router Interface on page 570
- Important Things to Remember for IGMP Snooping on page 570
- Important Things to Remember for IGMP Querier on page 571
- Configuration Task for IGMP Snooping on page 571

Joining a Multicast Group

A host joins a multicast group by sending either an unsolicited JOIN message or a JOIN message in response to the General Queries sent by the router or the IGMP snooping Querier. An IGMP switch floods general queries to all members of the VLAN interface.

When the IGMP switch receives a JOIN request from the host, it creates a multicast group entry for that VLAN and adds the ingress interface in the OIF (outgoing interface) list. The switch creates only one multicast group entry for a VLAN. If another host in the same VLAN sends a JOIN message for the same multicast group, that interface is also added to the OIF list of the previous Layer 3 flow entry. IGMP snooping switch sends only the first JOIN message for a multicast group to the multicast router. Subsequent JOIN messages for the same multicast group from other hosts in the VLAN are suppressed.

Leaving Multicast Group

normal leave process

In order to maintain the multicast group membership with the IGMP snooping switch, a host must either continue to respond to the general queries, or in the absence of a Querier, send an unsolicited membership report once during the "IGMP query interval." As long as there is one interested host in the VLAN the multicast router continues to forward multicast traffic to the VLAN interface.

When a host is no longer interested in receiving multicast traffic for that multicast group, either they may stop responding to the general queries (as done by IGMP version 1 host), or send a group leave message. This message is a IGMP version 2 group specific message.

If an IGMP snooping switch does not receive a membership report for two IGMP query intervals, it waits another ten seconds and then expires the host's multicast membership by removing the host interface from the OIF list.

When a group specific IGMP v2 leave message is received by the IGMP snooping switch, it removes the host interface from the OIF list and forwards this leave message to the multicast router interface only if the host was the only member of the multicast group in this VLAN. In other words, a group leave message is forwarded to the multicast router interface only when it is the LAST leave message for the multicast group in the VLAN.

When an IGMP snooping switch is also acting as a Querier, it sends out two Group Specific queries, separated by last-member-query-interval, on the interface where it received the group leave message. This ensures uninterrupted multicast data forwarding when there is another host on the same Ethernet segment interested in receiving traffic for that specific multicast group. If there is no JOIN request in response to the group specific query, the interface is removed from the OIF list.

fast leave process

Fast leave processing configuration in a VLAN enables the switch to remove the interface from the OIF list, after receiving the group leave message, without sending the two group-specific queries. Fast leave processing is supported on IGMP version 2 hosts only.

IGMP Snooping Querier Functionality

An IGMP switch, when configured, can act as an IGMP Querier for the VLAN without IGMP and PIM configured. This is typically done when IP multicast data traffic does not need to be routed.

When enabled, an IGMP snooping switch periodically sends out general queries to all members of the VLAN interface. These general queries causes the hosts to respond with membership report messages for the groups that wants IP multicast traffic. The switch then listens in on these frames and establishes group memberships for IP multicast data forwarding.

Querier functionality is enabled or disabled per VLAN basis.

When the Querier functionality is enabled on the VLAN of more than one switch, an election based on the IP address takes place. Switches with the lowest source IP address, in the general query frames, is elected as the Querier. Other switches maintain a timer; if they do not receive general query from the Querier for two query intervals, they send out a general query with the IP address assigned to their VLAN interface.

An IP address must be assigned to the VLAN address for the Querier processing to work on that interface.

Fast Convergence after MSTP-Triggered Topology Changes

When a port transitions to the Forwarding state as a result of an STP or MSTP topology change, FTOS sends a general query out of all ports except the multicast router ports. The host sends a response to the general query and the forwarding database is updated without having to wait for the query interval to expire.

When an IGMP snooping switch is not acting as a Querier it sends out the general query, in response to the MSTP triggered link-layer topology change, with the source IP address of 0.0.0.0 to avoid triggering Querier election.

Multicast Router Interface

You can designate an interface in the VLAN as a multicast router interface with the **ip igmp snooping mrouter interface** command. FTOS also has the capability of listening in on the incoming IGMP General Queries and designate those interfaces as the multicast router interface when the frames have a non-zero IP source address. All IGMP control packets and IP multicast data traffic is forwarded to the interfaces designated as multicast router interface.

Important Things to Remember for IGMP Snooping

- FTOS supports version 1 and version 2 hosts.
- FTOS IGMP snooping implementation is based on *draft-ietf-magma-snoop-10*.
- FTOS supports IGMP snooping.
- IGMP snooping is supported on VLANs on EtherScale and TeraScale systems.
- IGMP snooping is not enabled by default on the switch.

- A maximum of 1800 groups and 600 VLANs are supported.
- IGMP snooping is not supported on the default VLAN interface.
- IGMP snooping is not supported over VLAN-STACK enabled VLAN interfaces (you must disable IGMP snooping on a VLAN interface before configuring VLAN-STACK related commands).
- IGMP snooping does not react to Layer-2 topology changes triggered by STP.
- IGMP snooping reacts to Layer-2 topology changes triggered by MSTP by sending a general query on the interface that comes in FWD state.
- PIM and IGMP protocols run on L3 physical interfaces on EtherScale.
- PIM and IGMP protocols run on L3 physical interfaces and VLANs on TeraScale.

Important Things to Remember for IGMP Querier

- The IGMP snooping Querier supports version 2.
- You must configure an IP address to the VLAN interface for IGMP snooping Querier to begin. The IGMP snooping Querier disables itself when a VLAN IP address is cleared, and then it restarts itself when an IP address is re-assigned to the VLAN interface.
- When enabled, IGMP snooping Querier does not start if there is a statically configured multicast router interface in the VLAN.
- When enabled, IGMP snooping Querier starts after one query interval in case no IGMP general query (with IP SA lower than its VLAN IP address) is received on any of its VLAN members.
- When enabled, IGMP snooping Querier periodically sends general queries with an IP source address of the VLAN interface. If it receives a general query on any of its VLAN member, it checks the IP source address of the incoming frame. If the IP SA in the incoming IGMP general query frame is lower than the IP address of the VLAN interface, then the switch disables its IGMP snooping Querier functionality. If the IP SA of the incoming IGMP general query is higher than the VLAN IP address, the switch continues to work as an IGMP Querier.

Configuration Task for IGMP Snooping

The following list includes the configuration tasks for IGMP snooping:

- enable IGMP snooping globally on page 572
- enable IGMP snooping on the VLAN interface on page 572
- enable IGMP snooping Querier functionality on page 572

For a complete list of all IGMP snooping related commands, refer to *FTOS Command Line Interface Reference*.

enable IGMP snooping globally

By default, IGMP snooping is not enabled in FTOS. To enable IGMP snooping globally in FTOS use the following command in CONFIGURATION mode.

Step	Command Syntax	Command Mode	Purpose
1	**ip igmp snooping enable**	CONFIGURATION	Enable IP IGMP snooping globally
2	**show running-config igmp**	EXECUTIVE	View IP IGMP running configuration.

```
Force10#show running-config igmp
!
ip igmp snooping enable
```

Figure 357 enable IGMP snooping Command Example

enable IGMP snooping on the VLAN interface

By default, IGMP snooping is enabled on the VLAN interface when IGMP snooping is enabled globally. Execute the **no-shut** command on the VLAN interface for IGMP snooping to become operational. To disable IGMP snooping on a particular VLAN, use the command **no ip igmp snooping**.

enable IGMP snooping Querier functionality

IGMP snooping Querier functionality is not enabled in FTOS by default. To enable IGMP snooping Querier functionality on a VLAN use the following command in VLAN mode.

Step	Command Syntax	Command Mode	Purpose
1	**ip igmp snooping querier**	INTERFACE VLAN	Enable Querier functionality on a VLAN
2	**show ip igmp interface**	EXECUTIVE	View IP IGMP snooping Querier related information on the VLAN interface

Note: For IGMP Querier to work on VLAN, you must:

- Apply no-shut command on the VLAN
- Assign an IP address to the VLAN interface

```
Force10#show running-config igmp
!
ip igmp snooping enable
Force10#show ip igmp interface
Vlan 2 is up, line protocol is up
  IGMP Snooping query interval is 60 seconds
  IGMP Snooping querier timeout is 120 seconds
  IGMP Snooping last member query response interval is 1000 ms
  IGMP snooping fast-leave is disabled on this interface
  IGMP snooping querier is disabled on this interface
```

Figure 358 show ip igmp interface Command Example

PIM-Sparse Mode (Version 2)

Protocol-Independent Multicast-Sparse Mode (PIM-SM) is a multicast protocol in which multicast receivers explicitly join to receive multicast traffic. The protocol uses a router as the root or Rendezvous Point (RP) of the share tree distribution tree to distribute multicast traffic to a multicast group. Messages to join the multicast group (Join messages) are sent towards the RP and data is sent from senders to the RP so receivers can discover who are the senders and begin receiving traffic destined to the multicast group.

For more information, refer to *Internet Draft draft-ietf-pim-sm-v2-new-05.txt*.

Note: PIM-SM on VLAN interface is supported only on TeraScale platforms.

PIM-SM Implementation

FTOS implementation of PIM-SM is based on the IETF *Internet Draft draft-ietf-pim-sm-v2-new-05.txt*.

If the interface is the last hop router (directly connected), FTOS switches to Shortest Path Tree (SPT) as soon as it receives the first packet.

Configuration Tasks for PIM-SM

By default, IP multicast and all multicast protocols are disabled.

- configuring a router to be an RP on page 574
- enable PIM on an interface on page 574 (mandatory)
- override BSR updates on page 576
- display PIM-SM register messages on page 577
- creating multicast boundries and domains on page 577
- configure a static RP on page 578 (optional)
- modify PIM parameters on page 578 (optional)

configuring a router to be an RP

For a router to be used as a RP, the PIM-SM has to be enabled for the interface whose IP address is used as RP address. In the example below, "ip pim sparse-mode" is configured for interface loopback 0 because its address is the RP address.

```
Force10#sh run int loop0
!
interface Loopback 0
 ip address 1.1.1.1/32
 ip pim sparse-mode
 no shutdown
Force10#sh run pim
!
ip pim rp-address 1.1.1.1 group-address 224.0.0.0/4
```

IP PIM SM configured for Loopback 0 to allow it to be used as Rendezvous Point (RP).

enable PIM on an interface

When you enter **ip multicast-routing** command, IP multicast is enabled on the E-Series, but PIM and IGMP are not enabled on any interfaces. A PIM-enabled interface is added to the PIM routing table when the interface receives a join message from a downstream router or a directly connected host.

To enable IGMP and PIM on an interface, use these commands in the following sequence, beginning in the CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	**ip multicast-routing**	CONFIGURATION	Enable multicast routing on the system. If you have already entered this command, skip this step.
2	**interface** *interface*	CONFIGURATION	Specify the physical interface type, slot, and number. • For a Gigabit Ethernet interface, enter the keyword **GigabitEthernet** followed by the slot/port information. • For a Loopback interface, enter the keyword **loopback** followed by a number from 1 to 16383. • For a SONET interface, enter the keyword **sonet** followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword **TenGigabitEthernet** followed by the slot/port information. • **port-channel** <1 - 255> • **vlan** <1-4094>
3	**ip address** *ip-address mask*	INTERFACE	Assign an IP address to the interface. • *ip-address:* enter an address and mask
4	**no shutdown**	INTERFACE	Enable the interface.

Step	Command Syntax	Command Mode	Purpose
5	**ip pim sparse-mode**	INTERFACE	Enable IGMP and PIM on an interface.

To view which interfaces are enabled for PIM and IGMP, enter the **show ip pim interface** command in the EXEC privilege mode.

```
Force10#show ip pim interface
Address          Interface VIFindex Ver/   Nbr    Query  DR    DR
                                    Mode   Count  Intvl  Prio
189.87.5.6       Gi 4/11   0x2      v2/S   1      30     1     127.87.5.6
189.87.3.2       Gi 4/12   0x3      v2/S   1      30     1     127.87.3.5
189.87.31.6      Gi 7/11   0x0      v2/S   0      30     1     127.87.31.6
189.87.50.6      Gi 7/13   0x4      v2/S   1      30     1     127.87.50.6
Force10#
```

Figure 359 show ip pim interface Command Example

To view the PIM neighbors for each interface, use the **show ip pim neighbor** command in the EXEC privilege mode.

```
Force10#show ip pim neighbor
Neighbor          Interface      Uptime/Expires     Ver  DR
Address                                                  Prio/Mode
127.87.5.5        Gi 4/11        01:44:59/00:01:16  v2   1  / S
127.87.3.5        Gi 4/12        01:45:00/00:01:16  v2   1  / DR
127.87.50.5       Gi 7/13        00:03:08/00:01:37  v2   1  / S
Force10#
```

Figure 360 show ip pim neighbor Command Example

To view the PIM routing table, use the **show ip pim tib** command in the EXEC privilege mode.

```
Force10#show ip pim tib

PIM Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned,
       R - RP-bit set, F - Register flag, T - SPT-bit set, J - Join SPT,
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(*, 192.1.2.1), uptime 00:29:36, expires 00:03:26, RP 10.87.2.6, flags: SCJ
  Incoming interface: GigabitEthernet 4/12, RPF neighbor 10.87.3.5
  Outgoing interface list:
    GigabitEthernet 4/11
    GigabitEthernet 7/13

(10.87.31.5, 192.1.2.1), uptime 00:01:24, expires 00:02:26, flags: FT
  Incoming interface: GigabitEthernet 7/11, RPF neighbor 0.0.0.0
  Outgoing interface list:
    GigabitEthernet 4/11
    GigabitEthernet 4/12
    GigabitEthernet 7/13
--More--
```

Figure 361 show ip pim tib Command Example (Partial)

override BSR updates

To override bootstrap router (BSR) updates with the static RP, use the **ip pim rp-address <addr> group** <*addr/mask*> **override** command. When using this command, configuration changes are applied to the static RP configuration. When this command option is not applied, the RPs advertised by the BSR updates takes precedence over the statically configured RPs. This command is applied to a multicast group range.

Command Syntax	Command Mode	Purpose
ip pim rp-address *<addr>* **group override** *<addr/mask>*	CONFIGURATION	To override bootstrap router updates with the static RP.

The following configuration shows the RP for all multicast groups with the address 165.87.50.5, and shows the address 224.0.0.0/4 as representing all multicast groups:

To view the RP for a multicast group and the group address for all multicast groups, use the show running-configuration pim command in EXEC privilege mode:

```
Force10#show running-configuration pim
!
ip pim rp-address 165.87.50.5 group-address 224.0.0.0/4
```

Figure 362 ip pim rp-address <addr> group <addr/mask> override Command Example

To view the addresses within a group and their assigned RP, use the **show ip pim rp** command in EXEC privilege mode:

```
Force10#show ip pim rp
Group            RP
225.0.1.40       165.87.50.5
226.1.1.1        165.87.50.5
```

Figure 363 show ip pim rp Command Example

To view the group-to-RP mapping, use the **show ip pim rp mapping** command in EXEC privilege mode:

```
Force10#show ip pim rp mapping
PIM Group-to-RP Mappings
Group(s): 224.0.0.0/4, Static
  RP: 165.87.50.5, v2
```

Figure 364 show ip pim rp mapping Command Example

display PIM-SM register messages

To display PIM-SM register messages, use the **debug ip pim register** [**group**] command. The group option allows the user to filter register messages by a specified group.

➡ **Note:** The messages that come under this debug command include register encapsulation, null-register and register-stop.

Command Syntax	Command Mode	Purpose
debug ip pim register [**group**]	CONFIGURATION	To display PIM-SM register messages.

creating multicast boundries and domains

To create multicast boundries and domains by filtering inbound and outbound BSR messages per interface, use the ip pim bsr-border command:

Command Syntax	Command Mode	Purpose
[**no**] **ip pim bsr-border**	INTERFACE	To create multicast boundries and domains.

➡ **Note:** This command gets applied to the subsequent inbound and outbound updates. Already existing BSR advertisements are cleaned up by time out. Candidate RP advertisements can be cleaned up using clear ip pim rp-mapping.

To create multicast boundries and domains by filtering inbound and outbound Bootstrap Router (BSR) messages per interface, use the **ip pim bsr-border** command. This command is applied to the subsequent inbound and outbound updates. Already existing BSR advertisements are cleaned up by timeout.

Candidate RP advertisements can be cleaned up by using the **clear ip pim rp-mapping** command.

configure a static RP

In FTOS, at least one Rendezvous Point (RP) must be either learned or statically configured for multicast packets to flow.

To assign the group address to an RP, use the following command in the CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ip pim rp-address *address* **group-address** *group-address mask* **override**	CONFIGURATION	Assign an address to a group. Configure the following: • *address*: enter the IP address of the RP • *group-address mask:* enter the multicast group address and mask. You can configure multiple RP mappings for different group ranges by entering this command multiple times.

To view the RP mappings, use the **show ip pim rp mapping** command in the EXEC privilege mode.

```
Force10#show ip pim rp mapping
PIM Group-to-RP Mappings
Group(s): 224.0.0.0/4, Static
  RP: 10.87.2.6, v2

Force10#
```

Figure 365 show ip pim rp mapping Command Example

modify PIM parameters

In FTOS, you can change the following PIM timers or priority values:

- designated router (DR) priority
- frequency of PIM Hellos

To change these values, use any of these commands in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip pim dr-priority *priority-value*	INTERFACE	Change the Designated Router priority. • *priority-value:* enter a number from 0 to 4294967294. Default: 1.
ip pim query-interval *seconds*	INTERFACE	Change the frequency of PIM Hellos. • *seconds:* Enter a number from 0 to 65535. Default: 30 seconds.

To view the settings for each PIM-enabled interface, use the **show ip pim interfaces** command in the EXEC privilege mode.

Multicast Traceroute

MTRACE is an IGMP protocol based on the multicast traceroute facility which is implemented as per the IETF draft *draft-fenner-traceroute-ipm*. FTOS supports mtrace client and mtrace transmit functionality.

Mtrace Client

When acting as an mtrace client, FTOS transmits mtrace queries, and receives, parses, and prints out the details in the recieved response packets.

Mtrace Transit

When acting as an mtrace transit or intermediate router, FTOS returns the response to mtrace queries. Upon receiving an mtrace request, FTOS computes the RPF neighbor for the source, fills in the request, and forwards the request to the RPF neighbor. While computing the RPF neighbor, static mroutes and mBGP routes are preferred over unicast routes.

To trace a multicast route use the command **mtrace**.

Figure 366 Tracing a Multicast Route

```
Force10#mtrace 165.87.33.6 165.87.30.5 225.1.2.1
Type Ctrl-C to abort.
Mtrace from 165.87.33.6 to 165.87.30.5 via group 225.1.2.1
From source (?) to destination (?)
Querying full reverse path...
 0  165.87.30.5
-1  165.87.3.5  PIM [165.87.33.0/24]
-2  165.87.3.2  PIM [165.87.33.0/24]
-3  165.87.5.5  PIM [165.87.33.0/24]
-4  165.87.33.6
```

PIM Source Specific Multicast

The PIM-SSM protocol for IPv4 and IPv6 is based on the source specific model for forwarding multicast traffic across multiple domains in the Internet. It is restricted to shortest path trees to specific sources described by hosts using IGMPv3.

PIM-SSM is a subset of PIM-SM and shares with it the ability to join SPTs. However, unlike PIM-SM, PIM-SSM's register states and shared tree states for multicast groups in the SSM range are not maintained.

End-hosts use IGMPv3 to register their interest in a particular source-group (S,G) pair. PIM-SSM interacts with IGMPv3 to construct the multicast forwarding tree rooted at the source S.

Important Points to Remember

- The default SSM range is 232/8 and ff3x/32.
- Only standard ACLs are supported. Extended ACLs cannot be used for configuring SSM range.
- Ensure you configure the ACL first and then apply it to the SSM range.
- Applying an SSM range with an ACL overwrites the default range. To have both the default range and the SSM range be effective, add the default range to the SSM ACL.
- The default range is always supported. The range can never be smaller than the default.

Configuring and Verifying PIM-SSM

PIM-SSM supports a configurable SSM group range using an ACL. To configure an SSM group, use the command [**no**] {**ip** | **ipv6**} **pim ssm range** *access-list*

To display which non-default groups have been added to the PIM SSM, use the command **show ip pim smm-range**.

Chapter 34 Service Agent (FTSA)

C-Series	NO
E-Series	✓

Platform Specific Feature: FTSA is supported on E-Series only.

The Force10 Service Agent (FTSA), commonly called a call-home service, resides on the switch and collects information from the chassis manager, constructs email messages, and sends the messages to Force10 Support (Technical Assistance Center — TAC) and/or other designated servers.

FTSA currently collects only hardware and software inventory information, but it will eventually allow you to designate how much detail to send—error conditions, counters, statistical information, and so forth.

This chapter covers the following topics:

- Configuring FTSA on page 581
- Using FTSA on page 586

Configuring FTSA

FTSA requires minimal configuration to get it running at a basic level. After starting FTSA, you must designate an SMTP server that receives and forwards the email messages from the switch, and you must designate an Administrator email address. All other FTSA configuration tasks are optional, but you would typically want to perform them so that the full FTSA messaging functionality is enabled.

Configuration Task List

The following sections in this chapter detail the configuration tasks:

- starting and stopping FTSA on page 582 (mandatory)
- designate an SMTP server on page 582 (mandatory)
- designate an administrator email address on page 583 (mandatory)
- designate recipients of FTSA email messages on page 583 (optional)

- specify the frequency of FTSA email messages on page 586 (optional)
- enable the sending of emails to all recipients on page 586 (optional)

Note: For details on command syntax, see the Service Agent chapter in the *FTOS Command Line Interface Reference*.

starting and stopping FTSA

FTSA is stopped by default. In the CONFIGURATION mode, start FTSA with the **call-home** command. After you execute the command, you are in the CONFIGURATION (conf-callhome) mode, where you enter the FTSA configuration commands.

Use the **no call-home** command if you need to stop FTSA. You can start and stop FTSA at will, but stopping FTSA removes any FTSA configuration changes from the running configuration that you have made.

Command Syntax	Command Mode	Purpose
call-home	CONFIGURATION	This command has two functions: • Start FTSA. • Enter the CONFIGURATION (conf-callhome) mode. If FTSA is not started when you enter this command, executing the command performs both functions. If FTSA is already started when you enter this command, executing the command performs just the second function.

designate an SMTP server

After you start FTSA, you are in the CONFIGURATION (conf-callhome) mode, where you must designate an SMTP (Simple Mail Transfer Protocol) server that will receive email messages sent from FTSA on the switch and then forward the messages to the addressed recipients.

Command Syntax	Command Mode	Purpose
smtp server-address *server-address* [**smtp-port** *port number*]	CONFIGURATION (conf-callhome)	Enter the keyword server-address followed by the SMTP server address, such as smtp.yourco.com. Optionally, enter the keyword **smtp-port** followed by the SMTP port number, such as smtp-port 40. Range: 0 to 65535 Default: 25

designate an administrator email address

After you start FTSA, you are in the CONFIGURATION (conf-callhome) mode, where you must use the **admin-email** command to designate an Administrator email address. You can enter any email address that you want, but it would logically be for the FTSA administrator.

You can also use the **domain-name** command to specify the domain name for the Administrator email address, which would take priority over any domain name that you enter through the **admin-email** command.

Step	Command Syntax	Command Mode	Purpose
1	**admin-email** *email address*	CONFIGURATION (conf-callhome)	When FTSA is unable to send a message to an email address specified by a **recipient** command (unreachable destination), the SMTP server forwards the message to this Administrator email address. For *email address,* you have two choices: • Enter the administrator's full email address, such as: *admin@domain_name.com* • Enter just the username component, for example: *admin*
2	**domain-name** *domain_name*	CONFIGURATION (conf-callhome)	If you did not enter the domain name part of the Administrator's email address in the previous step, you must do so with this command. In any case, a domain name specified with this command will be used instead of a domain name you enter with the **admin-email** command. Enter the complete domain name of the Administrator's email address, such as *yourco.com.*

You now have a functioning FTSA service that will send periodic emails to Force10 Support. Use the following procedure to modify that recipient or add more.

designate recipients of FTSA email messages

FTSA is pre-configured to send email messages containing basic switch inventory information to Force10 Support. You can change that address, and you can add up to four more recipients. You can enable the sending of email messages to specific recipients or to all recipients.

After you start FTSA with the **call-home** command, as described above, you can add or edit recipients. The following steps describe that optional task. In summary, after using the **server** command to create a server name, you are placed at the server-specific prompt, where you use the **recipient** command to enter the email address of the recipient associated with the server. Still at that server-specific prompt, you have the option of enabling message encryption and using the **enable** command to start the sending of FTSA email messages to the selected recipient.

You can use this procedure for five recipients (including modifying the recipient associated with the Force10 server):

Step	Command Syntax	Command Mode	Purpose
1	**server** *name*	CONFIGURATION (conf-callhome)	For the *name* value, create a name for the recipient's server (if there is a name specified in an encryption file that you want to use for this server, you must enter that name—TeraScale only) in alphanumeric format, up to 25 characters long. The Force10 server is already added as the server name associated with ftsa@force10networks.com as the default recipient. If you want to change that address, enter **server Force10**. You will be placed at that server-specific prompt, where you can proceed to the next step.
2	**recipient** *email address*	CONFIGURATION Server (conf-callhome-*server_name*)	Enter the email address of the recipient to be associated with the selected server, for example: *name@domain_name.com*
3	**keyadd** *public key*	CONFIGURATION Server (conf-callhome-*server_name*)	(TeraScale only) Optionally, enter the filename of the public key (must be PGP 5.0 (Pretty Good Privacy) compatible) created for the selected server.
4	**encrypt**	CONFIGURATION Server (conf-callhome-*server_name*)	(TeraScale only) Encryption is disabled by default. Execute this command if you have entered a key (previous step), and if you want to encrypt messages to this recipient.
5	**enable**	CONFIGURATION Server (conf-callhome-*server_name*)	The default condition is disabled (no FTSA email messages to the selected recipient). Use this command or the **enable-all** command at the global level to enable messages to the recipient.

Step	Command Syntax	Command Mode	Purpose
6	**exit**	CONFIGURATION Server (conf-callhome-*server_name*)	This command returns you to the (conf-callhome) prompt, where you can repeat the sequence starting at Step 3 to add up to five recipients, including the recipient associated with the “Force10” server. Alternatively, you can use the **end** command to back out to the basic CONFIGURATION mode.
7	**exit**	CONFIGURATION (conf-callhome)	This command returns you to the basic CONFIGURATION mode. Before you do so, you might want to use the optional global configuration commands and **show** commands described below.

The use of the enable or enable-all commands causes FTSA to send a status email message to the affected recipients, and that activity is tracked by the CLI, as shown in the following sample:

```
Force10(conf-callhome)#enable-all
Force10(conf-callhome)#Aug 3 13:06:36: %RPM1-P:CP %CALL-HOME-3-CALLHOME: Callhome
service sent a message to Force10 at ftsa@force10networks.com
Aug 3 13:06:37: %RPM1-P:CP %CALL-HOME-3-CALLHOME: Callhome service sent a message
to helen at helenh@yourco.com
```

Figure 367 enable-all Command Example

modifying server settings

Each command in the procedure listed above has a “no” form of the command. Therefore, you have several options for disabling email to a particular recipient, although each has its own ramifications:

- Using the **no enable** command stops messaging to the selected recipient, but does not remove the rest of the server configuration.
- Using the **no recipient** *email address* command removes the recipient from the server configuration, but the rest of the configuration remains so that you can enter a new recipient for the server.
- Using the **server** *name* from the (conf-callhome) prompt removes the server configuration, including the recipient. The one exception is the Force10 server, which you cannot remove, but you can remove or modify the recipient and other settings.

Optional Global FTSA Settings

In addition to the required configuration tasks that you perform at the (conf-callhome) prompt, you have the option of selecting a frequency of status email messages from FTSA and of enabling or disabling messages to all recipients.

specify the frequency of FTSA email messages

Optionally, use the **frequency** command at the (conf-callhome) prompt to select the time interval with which email messages will be sent to all designated recipients. If you do not select a time interval, FTSA will send status messages every 24 hours after FTSA is started.

If you set the frequency to an interval other than the default, you can use the **no frequency** command to return to the default, or you can enter the command again to select another interval.

There is no interval setting override for individual recipients.

Command Syntax	Command Mode	Purpose
frequency *minutes*	CONFIGURATION (conf-callhome)	(OPTIONAL STEP) Enter the time interval, in minutes, that you want between FTSA status emails. Range: 2 to 10080 minutes Default: 1440 minutes (24 hours)

enable the sending of emails to all recipients

Optionally, use the **enable-all** command at the (conf-callhome) prompt to enable the sending of FTSA messages to all recipients. Alternatively, at a server-specific prompt, you can use the **enable** command to enable messages for that specific recipient.

To disable (end) the sending of FTSA email messages to all designated recipients, use the **no enable** command. Alternatively, if you used the **enable-all** command and then want to disable messages to a specific recipient, you can use the **no enable** command at that server-specific prompt.

Command Syntax	Command Mode	Purpose
enable-all	CONFIGURATION (conf-callhome)	The default condition is disabled (no FTSA email messages to any recipients). Any recipient email that is in the default disabled state can be enabled with this command.

Using FTSA

The purpose of FTSA is to automatically send information about the switch on a periodic basis to troubleshooters, so that, if a problem arises, that information can be used to help assess the problem.

After you start FTSA, FTSA runs in the background on the switch, periodically collecting status information and sending it by email to Force10 Support or any recipients that you designate, as described above. After you start and configure FTSA, you should not need to do anything unless you want to modify the configuration.

Interpreting Email Messages from FTSA

FTSA sends email to all enabled recipients in the following situations:

- On startup of FTSA
- On shutdown of FTSA
- On the execution of the **enable-all** command
- On the execution of the **no enable-all** command
- At the designated time interval (**frequency**)

FTSA sends an email to a particular recipient when email to that recipient is enabled or disabled by using the **enable** or **no enable** commands, respectively.

FTSA currently sends three kinds of messages:

- On startup / enable-all / enable (message type 0)
- On shutdown / no enable-all / no enable (message type 1)
- At the designated time interval (message type 3)

The email body always contains information about the type of message, the chassis name, and the time when the message was sent. Encrypted messages (TeraScale only) are encrypted with the recipient's public key and signed with the public key of the chassis.

All message types contain an attachment with the same information as that provided by the **show inventory** command. In addition, when you enable messaging for a recipient with a new encryption key (TeraScale only), the ensuing message to that recipient contains the public key of the chassis encrypted with the recipient's public key.

Using Show Commands

The FTSA service has three show commands. All are executed from the (conf-callhome) prompt:

- **show configuration**: Displays the FTSA (call-home) configuration
- **show debugging**: Displays the status of FTSA debugging
- **show keys**: Displays the email encryption (PGP) keys (TeraScale only)

Example of Using the show configuration Command

```
Force10(conf-callhome)#show configuration
!
call-home
   admin-email traza
   domain-name yourco.com
    frequency 480
   smtp server-address 10.0.2.6
   no enable-all
   server Force10
     recipient ftsa@force10networks.com
     keyadd Force10DefaultPublicKey
     no encrypt
     enable
    server someserver
         recipient joe@yourco.com
         no encrypt
         enable
Force10(conf-callhome)#
```

Figure 368 show configuration Command Example

example of using the show debugging command

```
Force10(conf-callhome)#show debugging

CALLHOME:
    Callhome service debugging is on

Force10(conf-callhome)#
```

Figure 369 show debugging Command Example

example of using the show keys command

```
Force10(conf-callhome)#show keys

Type Bits KeyID      Created    Expires    Algorithm       Use

sec+  768 0x64CE09D9 2005-06-27 ---------- RSA            Sign & Encrypt
uid  E000000003209
pub  1024 0xA8E48C2F 2004-12-08 ---------- DSS            Sign & Encrypt
sub  1024 0xD832BB91 2004-12-08 ---------- Diffie-Hellman
uid  Force10

2 matching keys found
Force10(conf-callhome)#
```

Figure 370 show keys Command Example

Debugging FTSA

Use the **debug call-home** command, from either the EXEC mode or the EXEC Privilege mode, to monitor FTSA messages. If no recipient is enabled, no email is sent, so no messages are displayed on the CLI.

To turn the message monitoring off, use the **no debug call-home** command.

Chapter 35 MSDP

C-Series	NO
E-Series	✓

Platform Specific Feature: MSDP is supported on E-Series only.

Multicast Source Discovery Protocol (MSDP) is a mechanism designed to connect multiple PIM Sparse-Mode (PIM-SM) domains. Each PIM-SM domain uses its own independent Rendezvous Points (RPs) and does not depend on RPs in other domains. The RPs are only aware of the sources and receivers within their domain. MSDP enables RPs of different domains to share information on sources. This shared knowlege enables RPs to send information to receivers to enable multicast traffic to be forwarded between those domains. MSDP can also be used within a domain to provide RP redundancy (Anycast-RP).

MSDP interconnects RPs with TCP connections to pass source active (SA) messages. RPs send SA messages for internal sources to MSDP peers. SAs are checked for peer-RPF (Reverse Path Forwarding) before being accepted or forwarded. RPs learn of existing external sources through the SA messages and may trigger (S,G) joins on behalf of local receivers.

MSDP peers connected using TCP port 639. Peers send KeepAlives every 60 seconds. A peer connection is reset after 75 seconds if no MSDP packets are received.

MSDP connections are typically parallel with multiprotocol BGP (MBGP) connections.

SA packets are periodically sent to MSDP peers indicating:

- Source address of active stream
- Group address of active stream
- IP address of RP originating the SA
- Only originate SAs for sources within local domain

Initial SA messages sent when source first registers to RP and subsequent SA messages periodically refreshed every 60 seconds as long as source still active by originating RP. Interested parties (RPs) can send PIM JOINs towards source to create inter-domain source tree.

This chapter includes the following topics for MSDP:

Peer-RPF Forwarding

The MSDP Peer-RPF Forwarding rules are used for forwarding SA messages throughout an MSDP enabled internet. Unlike the RPF check used when forwarding data packets, which generally compares the packet's source address against the interface upon which the packet was received, the Peer-RPF check compares the RP address carried in the SA message against the MSDP peer from which the message is received.

MSDP Mesh-group Semantics

MSDP mesh-group is an operational mechanism for reducing SA flooding, typically in an intra-domain setting. In particular, when some subset of a domain's MSDP speakers are fully meshed, they can be configured into a mesh-group.

If member R of a mesh-group receives an SA message from an MSDP peer that is also a member of mesh-group, member R accepts the SA message and forwards it to all of its peers that are not part of the mesh-group. Member R must not forward the SA message to other members of the mesh-group.

Timers

The timers for MSDP are:

- SA advertisement timer
- SA cache entry timer
- Peer hold timer
- Keepalive timer
- Connect retry timer

SA-Advertisement Timer

There is one SA-Advertisement Timer covering the sources that an RP may advertise. SA-Advertise-Period must be 60 seconds. An RP must not send more than one periodic SA message for a given (S,G) within an SA Advertisement interval. An originating RP should trigger the transmission of a SA message as soon as it receives data from an internal source for the first time.

An RP packs its active sources into a SA message until the largest MSDP packet that can be sent is built or there are no more sources, and then sends the message.

SA-Cache Timeout (SA-State Timer)

An (S,G) SA-State Timer is started when an (S,G) SA message is initially received by an MSDP peer. The timer is reset to SG-State-Period if another (S,G) SA message is received before the (S,G) SA-State Timer expires. SG-State-Period must not be less than SA-Advertisement-Period + SA-Hold-Down-Period.

Peer Hold Timer

The Peer Hold Timer is initialized to HoldTime-Period when the peer’s transport connection is established, and is reset to HoldTime-Period when any MSDP message is received. If the timer expires, the transport connection is restarted. The HoldTime-Period must be at least three seconds, the recommended value is 75 seconds.

KeepAlive timer

Once an MSDP transport connection is established, each side of the connection sends a KeepAlive message and resets the KeepAlive timer. The KeepAlive timer is set to KeepAlive-Period each time an MSDP message is sent to the peer. KeepAlive-Period must be less than HoldTime-Period, and at least one second. The recommended value is 60 seconds.

ConnectRetry Timer

The ConnectRetry timer is used by the MSDP peer with the lower IP address to transit from INACTIVE to CONNECTING states. ConnectRetry-Period should be set to 30 seconds.

Anycast RP

Anycast RP allows two or more RPs to load share or source registration and to act as running backup routers for each other (RP redundancy).

MSDP may be used to implement the concept of Anycast RP within a PIM-SM domain to provide:

- RP redundancy
- Rapid RP fail-over
- RP load-balancing

To support Anycast RP, MSDP peers the Anycast RP to distribute the active-registered sources.

The Anycast RP mechanism works as follows:

- Each RP is assigned the same RP address (Anycast RP) and advertises its Anycast RP address as a host route in the Unicast routing domain.
- Source Designated Routers (DRs) and group receivers are based on their Unicast routing table to Register or Join the closed Anycast RP.
- Anycast RP are all connected via MSDP. This allows each Anycast RP to learn which sources have been registered with other Anycast RP in the domain.

MSDP Commands and Configuration Examples

The following list includes the configuration tasks for MSDP:

- Configuring MSDP Peer on page 594
- Defining Default Peer on page 595
- Defining Originator ID on page 596
- Configuring a Member of a Mesh Group on page 596
- Showing MSDP on page 597
- Clearing MSDP on page 599
- Shutting Down an MSDP Peer on page 599
- Using IP Address of Interface as RP on page 600
- MSDP Debugging on page 600
- Example: Two Routers with MSDP on page 601

For a complete listing of all commands related to Multiple Spanning Tree Protocol, see the *FTOS Command Line Interface Reference*.

Configuring MSDP Peer

The following command is used to configure an MSDP peer:

Command Syntax	Command Mode	Purpose
[**no**] **ip msdp peer** {*peer-address* } [**description** *description*]	CONFIGURATION	Configures an MSDP peer. Specify the address, and optionally provide a description for the peer.

Figure 371 Configuring an MSDP Peer

```
Force10(conf)# ip msdp peer 192.168.196.1 description peer1
Force10(conf)#sh ip msdp summary
Peer Addr Local Addr State Source SA Up/Down Description
72.30.1.2 72.30.1.1 Established none 0 00:00:03 peer1
72.30.2.2 72.30.2.1 Established none 0 00:00:03 peer2
72.30.3.2 72.30.3.1 Established none 0 00:00:02 test-peer-3
```

Defining Default Peer

The following command is used to configure an MDSP default peer:

Command Syntax	Command Mode	Purpose
[**no**] **ip msdp default-peer** *peer-address* [**prefix-list** {*name*}]	CONFIGURATION	Defines a default peer from which to accept all SA messages. If the prefix-list is not specified, all SA messages received from the default peer are accepted.You can enter multiple default-peer commands. (Optional)

The following example shows how to configure an MSDP default peer and naming it “xyz”:

```
Force10(conf)#ip msdp default-peer xyz
```

Figure 372 Configuring a Defaul Peer

The following output shows two MSDP peers, and two MSDP default peers with their addresses from the command **show running-configuration**:

```
Force10#show run msdp
!
ip multicast-msdp
ip msdp peer 172.21.3.254
ip msdp peer 172.21.5.254
ip msdp default-peer 172.21.3.254
ip msdp default-peer 172.21.5.254
```

Figure 373 Output for MSDP Peer and Defaul Peer

Defining Originator ID

The following command is used to configure an MDSP originator ID:

Command Syntax	Command Mode	Purpose
[**no**] **ip msdp originator-id [fastethernet | gigabitethernet | loopback | port-channel | sonet | tengigabitethernet | vlan]**	CONFIGURATION	Enter the following keywords and slot/port or number information: **fastethernet**—Fast Ethernet interface **gigabitethernet**—Gigabit Ethernet interface **loopback**—Loopback interface **port-channel**—Port Channel interface **sonet**—SONET interface **tengigabitethernet**—TenGigabit Ethernet interface **vlan**—VLAN interface

The following example shows how to configure an MSDP originator ID and naming it “xyz”:

```
Force10(conf)#ip msdp originator-id xyz
```

Figure 374 Configuring an Originator ID

Configuring a Member of a Mesh Group

The following command is used to configure a mesh group:

Command Syntax	Command Mode	Purpose
[**no**] **ip msdp mesh-group** *name peer-address*	CONFIGURATION	Configures a peer to be a member of a mesh group.

The following example shows how to configure an MSDP mesh-group and naming it “xyz”:

```
Force10(conf)#ip msdp mesh-group ?

WORD                    Name of mesh-group (max 16 chars)
Force10(conf)#ip msdp mesh-group xyz
```

Figure 375 Configuring MSDP Mesh-group

The following output shows three MSDP default peers in a mesh-group called “GRP1” and their addresses from the command **show running-configuration**:

```
Force10#show run msdp
!
ip multicast-msdp
ip msdp peer 172.21.3.254
ip msdp peer 172.21.5.254
ip msdp peer 172.21.6.254
ip msdp mesh-group GRP1 172.21.3.254
ip msdp mesh-group GRP1 172.21.5.254
ip msdp mesh-group GRP1 172.21.6.254
```

Figure 376 Showing MSDP Peer and and Mesh-groups Configuration

Showing MSDP

The following commands are used to show MSDP information:

Command Syntax	Command Mode	Purpose
show ip msdp [peer *peer-address* | **sa-cache** | **summary**]	CONFIGURATION	**peer** *peer-address*—Displays status on all or one MSDP peer. **sa-cache**—Displays Source Active Cache database or summary information. **summary**—Displays MSDP peer summary.

peer

The following output shows sample MSDP peer information from the command **show ip msdp peer**:

```
Force10#show ip msdp peer 100.1.1.1

Peer Addr: 100.1.1.1
    Local Addr: 100.1.1.2(639)  Connect Source: none
    State: Established  Up/Down Time: 00:00:08
    Timers: KeepAlive  60 sec, Hold time  75 sec
    SourceActive packet count (in/out): 0/0
    SAs learned from this peer: 0
    SA Filtering:
    Input (S,G) filter: none
    Output (S,G) filter: none
```

showing source active cache information

The following output shows sample MSDP SA cache information from the command **show ip msdp sa-cache**:

```
Force10#show ip msdp sa-cache
MSDP Source-Active Cache - 1 entries
GroupAddr       SourceAddr      RPAddr          LearnedFrom   Expire UpTime
224.1.1.1       172.21.220.10   172.21.3.254    172.21.3.254    102 00:02:52
```

showing MSDP summary

The following output shows sample MSDP summary from the command show:

```
Force10#show ip msdp summary

Peer Addr       Local Addr      State       Source    SA       Up/Down
100.1.1.1       100.1.1.2       Established  none      0       00:00:02
100.1.2.1       100.1.2.2       Established  none      0       00:00:02
100.1.3.1       100.1.3.2       Established  none      0       00:00:02
172.21.10.254   0.0.0.0         Inactive     Lo 0      0       00:00:11
```

Figure 377 show ip msdp summary Command Example

The following output shows sample MSDP information from the command **show running-configuration msdp**:

```
Force10#show run msdp
!
ip multicast-msdp
ip msdp log-adjacency-changes
ip msdp peer 100.1.1.1
ip msdp peer 100.1.2.1
ip msdp peer 100.1.3.1
ip msdp peer 172.21.10.254 connect-source Loopback 0
```

Figure 378 show running-configuration msdp Command Example

Clearing MSDP

The following commands are used to clear peer statistics or source-active entries from cache:

peer statistics

Command Syntax	Command Mode	Purpose
clear ip msdp peer peer-address	CONFIGURATION	Resets the TCP connection to the peer, clears all the peer statistics, and the transmission FIFO.

source active entries

Command Syntax	Command Mode	Purpose
clear ip msdp sa-cache *group-address*	CONFIGURATION	Clears from a Multicast group address all the source-active entries from the source-active cache. (Optional)

Shutting Down an MSDP Peer

The following command is used to shut down an MSDP peer:

Command Syntax	Command Mode	Purpose
[**no**] **ip msdp shutdown** *peer-address*	CONFIGURATION	Administratively shuts down a configured MSDP peer.

The following example shows how to shut down an MSDP peer using the command **ip msdp shutdown**:

```
Force10(conf-t)#ip msdpt shutdown 100.1.1.1
```

Figure 379 Shutting Down an MSDP Peer

Using IP Address of Interface as RP

The following command is used to allows an MSDP speaker to use the IP address of an interface as the rendezvous point:

Command Syntax	Command Mode	Purpose
[**no**] **ip msdp originated-id** *interface*	CONFIGURATION	Allows an MSDP speaker that originates an SA message to use the IP address of the interface as the rendezvous point (RP) address in the SA message.

MSDP Debugging

The following is MSDP is the debugging command **debug ip msdp**:

Command Syntax	Command Mode	Purpose
[**no**] **debug ip msdp [event** *peer-address* | **packet** *peer-address* | **pim]**	CONFIGURATION	Turns on all MSDP debug options: **event** *peer-address*—To debug MSDP protocol event. **packet** *peer-address*—To debug protocol packet. **pim**—To debug advertisement from PIM.

Example: Two Routers with MSDP

The following examples shows two routers configured with MSDP to run Anycast RP. (All PIM-SM routers are configured with 1.1.1.1 as their RP address).

first chassis

```
Force10#show run msdp
...
interface Loopback 0
 ip address 1.1.1.1/32
 ip pim sparse-mode
 no shutdown
!
interface Loopback 1
 ip address 172.21.10.254/32
 no shutdown
!
ip multicast-msdp
ip msdp peer 172.21.6.254 connect-source Loopback 1
ip msdp originator-id Loopback 1
!
ip pim rp-address 1.1.1.1 group-address 224.0.0.0/4
```

second chassis

```
Force10#show run msdp
...
interface Loopback 0
 ip address 1.1.1.1/32
 ip pim sparse-mode
 no shutdown
!
interface Loopback 1
 ip address 172.21.6.254/32
 no shutdown
!
ip multicast-msdp
ip msdp peer 172.21.10.254 connect-source Loopback 1
ip msdp originator-id Loopback 1
!
ip pim rp-address 1.1.1.1 group-address 224.0.0.0/4
```

Source Active Message Limiting

FTOS caches the SA messages it created and those received from MSDP peers. You can limit the number of SA messages that the Force10 system keeps, and limit the number of SA messages that can be received from an MSDP peer.

Limiting the Source Active Message Cache

Set the upper limit of the number of Source Active messages that FTOS keeps. The default SA limit is 500K messages. When the total number of SA messages reaches the specified limit, subsequent SA messages are dropped even if they pass RPF checking and policy checking.

Set an SA message limit using the command **ip msdp sa-limit** *number* from CONFIGURGATION mode.

If the total number of SA messages is already larger than the limit when message limiting is applied, the messages that are already in FTOS are not discarded. To enforce the limit in such a situation, use the command **clear ip msdp sa-cache** to clear all SA messages.

Limiting the Source Active Messages Received from a Peer

Set the upper limit of the number of SA messages allowed from an MSDP peer using the command **ip msdp peer** *peer-address* **sa-limit** *number* from CONFIGURATION mode. The default limit is 100K.

If the total number of SA messages received from the peer is already larger than the limit when this configuration is applies, those SA messages are not discarded. To enforce the limit in such a situation, use the command **clear ip msdp peer** to clear all SA messages received from the peer.

Viewing Source Active Message Information

- View the number of per peer SAs learned and the SA limits of each using the command **sh ip msdp summary**.
- Verify the global SA limit using the command **sh ip msdp sa-cache**.

Chapter 36 FTOS XML Feature

C-Series	NO
E-Series	✓

Platform Specific Feature: XML is supported on E-Series only.

This chapter describes the FTOS XML Feature in the following major sections:

- XML Functionality on page 603
- The Form of XML Requests and Responses on page 604
- The Configuration Request and Response on page 605
- The “Show” Request and Response on page 606
- Configuration Task List on page 606
- XML Error Conditions and Reporting on page 610
- Using display xml as a Pipe Option on page 614

XML Functionality

Through SSH/Telnet client sessions, FTOS XML provides a way of interfacing with the E-Series by entering XML-formatted requests and retrieving XML output. See The Form of XML Requests and Responses on page 604.

FTOS XML supports the following functionality:

- Configure both physical and logical interfaces
- Layer 2 and Layer 3 Standard ACLs
- Layer 2 and Layer 3 Extended ACLs
- Supported show commands and their output. Some show command options supported by FTOS are not supported in XML, so each option that is supported in XML is listed separately here for clarity:

  Protocol commands:

  — **show ip bgp neighbors** (no parameters accepted)

  — **show qos statistics**

  — **show qos statistics wred-profile**

  System commands:

- **show chassis**
- **show rpm** *slot ID*
- **show rpm all**
- **show linecard** *slot ID*
- **show linecard all**
- **show sfm** *slot ID*
- **show logging** *1-65535*
- **show logging reverse**
- **show sfm**
- **show sfm all**
- **show version**
- **show running-config**—Only the full report is supported, no options.
- **show interfaces**—All the options are supported except **rate**:

The Form of XML Requests and Responses

To send an XML-formatted command through a Telnet or SSH client session, you first use the **terminal xml** command to inform FTOS that you wish to switch to XML mode. See run an FTOS XML session on page 607.

Note: FTOS accepts well-formed XML requests, except that it does not currently support XML Namespaces.

Request Format

You can then enter XML-formatted requests that conform to the following schema. Every XML request begins with an XML declaration, followed by a "Method" type tag, followed by an "Operation" type tag, as shown in this shell schema:

```
<?xml version="1.0" encoding="UTF-8"?>
<Request MajorVersion="1" MinorVersion="0">
   <Method>
        <Operation>
             <command>
        :: ! The number of allowed <command> tag sets depends on the type of request. !
        ::
             </command>
        </Operation>
   </Method>
</Request>
```

Currently, for “Method”, you must enter “cli”. In place of “Operation”, you enter either “configuration” or “action”, depending on the CLI mode that you want to invoke:

Namespace	Description
<configuration>	This tag tells the CLI to invoke the CONFIGURATION mode. These requests encapsulate configuration modification commands.
<action>	This tag tells the CLI to invoke the EXEC PRIVILEGE mode. These requests encapsulate “show” commands.

Response Format

Similarly, every response from FTOS begins with the XML declaration, followed by a “Response” tag:

```
<?xml version="1.0" encoding="UTF-8"?>
<Response MajorVersion="1" MinorVersion="0">
    ::
    ::
 </Response>
```

What goes between the Response tags depends on the type of response, as discussed next.

The Configuration Request and Response

To create a configuration request, you know from the introduction above that you put “<cli>” in place of the “<Method>” tag in the schema, and you put “<configuration>” in place of “<Operation>”. The number of configuration commands in one request is not restricted.

Just as you enter commands in the CLI, you have the option of entering abbreviated commands in XML messages. For example, instead of using the full **show running-config** statement, you can enter **show run**. Also, spaces before or after the command are allowed, as shown in the following example.

The following sequence of XML tags shows the structure of a configuration request containing several commands:

```
<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<cli>
<configuration>
<command>ip access standard test2 </command>
<command> seq 10 deny any</command>
<command> seq 20 permit host 10.1.1.1 count </command>
<command>seq 30 deny 10.2.0.0 /16</command>
</configuration>
</cli>
</request>
```

The response from FTOS, if the command executes successfully, is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<response MajorVersion="1" MinorVersion="0">
<responseType>NO_ERROR</responseType>
<responseSeverity>SEVERITY_INFO</responseSeverity>
<responseMsg>Xml request successfully processed.</responseMsg>
</response>
```

For details on responses to error conditions, see XML Error Conditions and Reporting on page 610.

The “Show” Request and Response

To generate an XML request that encapsulates a “show” command (to request a report), you use the <action> tag instead of the <configuration> tag as the Operation type. The schema of a show request allows only one <command>, as shown here for the **show linecard** command. (Note that “<command>**show line all**</command>” demonstrates that you can use both an abbreviated form of the command and options, just as in the standard CLI):

```
<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<cli>
<action>
<command>show line all</command>
</action>
</cli>
</request>
```

The response from FTOS, if the command executes successfully, presents all of the content that you would get in the equivalent CLI report. Note that the data are encapsulated in self-explanatory XML tags. The following is an example of a **show linecard** report embedded in XML tags:

```
<?xml version="1.0" encoding="UTF-8" ?>
<response MajorVersion="1" MinorVersion="0">
<action>
<linecard>
<slotId>3</slotId>
<status>online</status>
<nextBoot>online</nextBoot>
<reqType>E48TF3</reqType>
<numPorts>0</numPorts>
<swVer>6.5.1.1</swVer>
</linecard>
</action>
</response>
```

Configuration Task List

In addition to supporting show commands, FTOS XML currently also supports ACL configuration:

- run an FTOS XML session on page 607
- configure a standard ACL on page 609
- configure an extended ACL on page 609
- apply an IP ACL on page 609
- create an egress ACL and apply rules to the ACL on page 610

run an FTOS XML session

Use the following procedure to start, run, and close an FTOS XML session:

Step	Command Syntax	Command Mode	Purpose
1	**terminal xml**	EXEC PRIVILEGE	Enable XML mode in Telnet and SSH client sessions.
2	[Construct input to the CLI by following the XML request schema, as described in The Form of XML Requests and Responses on page 604.]	FTOS XML	Cut and paste your XML request from a text editor or other type of XML presentation tool, or type your XML request line by line.
3	Press **Ctrl-Y** (or press **Enter** twice, creating an empty line).	FTOS XML	Execute the request. Alternatively, to cancel the request (only possible before sending) and get a fresh XML prompt, press **Ctrl-C**.
4	Press **Ctrl-Z** (or enter **terminal no xml** as the <command> string in the XML request <action> schema).	FTOS XML	Exit from FTOS XML mode.

Figure 380, below, illustrates entering FTOS XML mode. Figure 381, below, illustrates the full sequence of invoking an XML session, entering a command, receiving a success response, and leaving the session with the **terminal no xml** command in XML:

```
Force10# terminal xml
Force10(xml)#
Enter XML request with CTRL-Y or empty line
Clear XML request with CTRL-C
Exit XML mode with CTRL-Z:
```

Figure 380 Example of Entering FTOS XML mode from the CLI

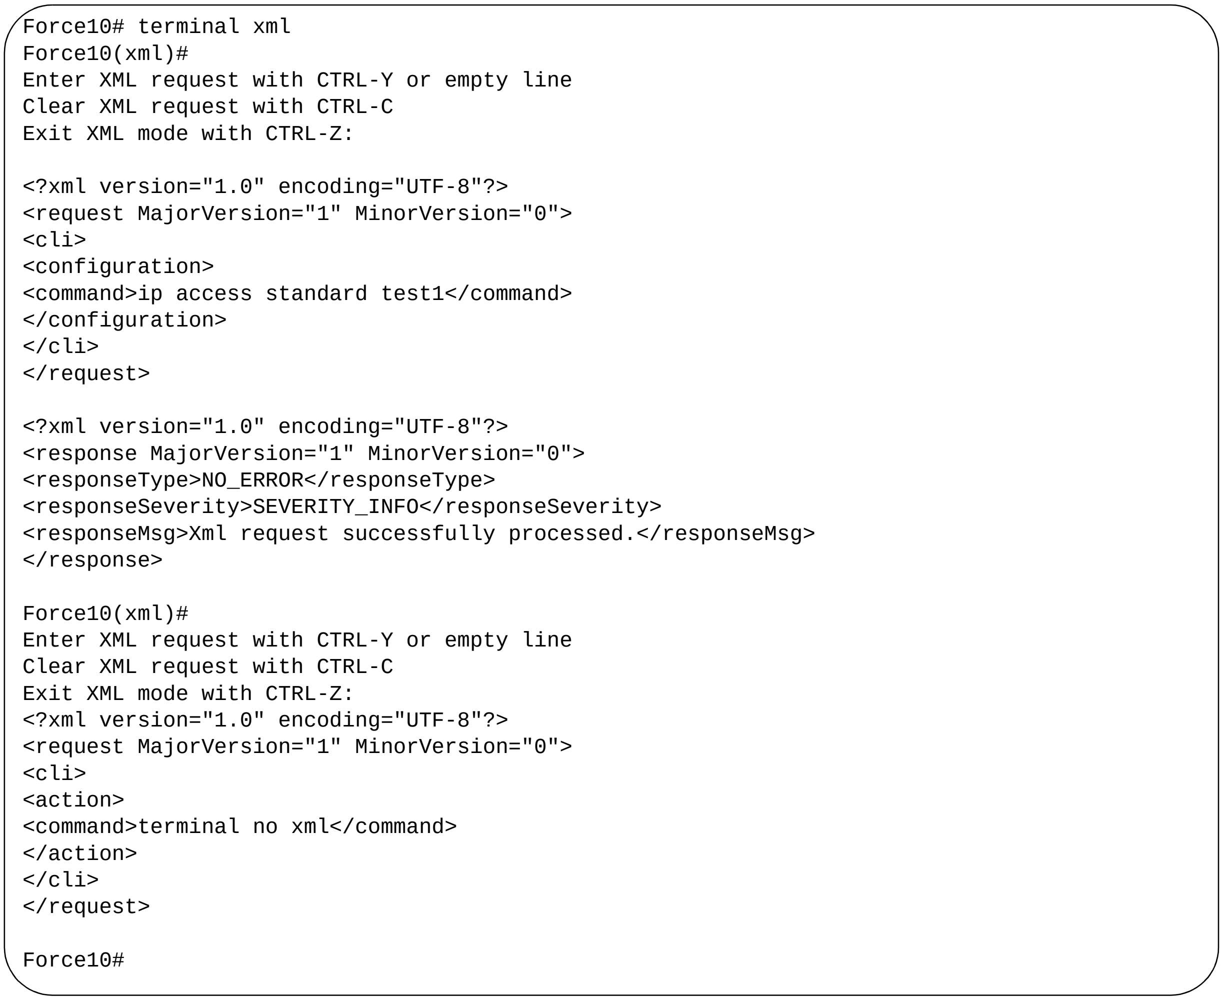

```
Force10# terminal xml
Force10(xml)#
Enter XML request with CTRL-Y or empty line
Clear XML request with CTRL-C
Exit XML mode with CTRL-Z:

<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<cli>
<configuration>
<command>ip access standard test1</command>
</configuration>
</cli>
</request>

<?xml version="1.0" encoding="UTF-8"?>
<response MajorVersion="1" MinorVersion="0">
<responseType>NO_ERROR</responseType>
<responseSeverity>SEVERITY_INFO</responseSeverity>
<responseMsg>Xml request successfully processed.</responseMsg>
</response>

Force10(xml)#
Enter XML request with CTRL-Y or empty line
Clear XML request with CTRL-C
Exit XML mode with CTRL-Z:
<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<cli>
<action>
<command>terminal no xml</command>
</action>
</cli>
</request>

Force10#
```

Figure 381 Example of a Successful XML Session

configure a standard ACL

To configure a standard ACL with XML, first enter FTOS XML mode, and then construct a configuration request, as described above. An example of a complete standard ACL configuration request message is:

```
<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<cli>
<configuration>
<command> ip access list standard ToOspf</command>
<command> seq 5 deny any</command>
<command> seq 10 deny 10.2.0.0 /16</command>
<command> seq 15 deny 10.3.0.0 /16</command>
<command> seq 20 deny 10.4.0.0 /16</command>
<command> seq 25 deny 10.5.0.0 /16</command>
<command> seq 30 deny 10.6.0.0 /16</command>
<command> seq 35 deny 10.7.0.0 /16</command>
<command> seq 40 deny 10.8.0.0 /16</command>
<command> seq 45 deny 10.9.0.0 /16</command>
<command> seq 50 deny 10.10.0.0 /16</command>
</configuration>
</cli>
</request>
```

configure an extended ACL

To configure an extended ACL through XML, enter FTOS XML mode and construct an XML configuration request (see run an FTOS XML session on page 607). An example of a complete request message is:

```
<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<cli>
<configuration>
<command> interface GigabitEthernet 0/0</command>
<command> ip address 10.2.1.100 255.255.255.0 </command>
<command> ip access-group nimule in no shutdown</command>
</configuration>
</cli>
</request>
```

apply an IP ACL

To apply the IP ACL (standard or extended) that you created, above, to a physical or port channel interface, construct an XML configuration request (see **run an FTOS XML session on page 607**) that encapsulates the appropriate CLI commands, as exemplified here:

```
<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<cli>
<configuration>
<command> interface GigabitEthernet 0/0</command>
<command> ip address 10.2.1.100 255.255.255.0 </command>
<command> ip access-group nimule in no shutdown</command>
</configuration>
</cli>
</request>
```

create an egress ACL and apply rules to the ACL

To create an egress ACL and apply rules to the ACL in one single XML request, first enter FTOS XML mode, and then construct the configuration request (see run an FTOS XML session on page 607). The following example shows a configuration request message that accomplishes this task:

```
<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<cli>
<configuration>
<command> interface GigabitEthernet 0/0</command>
<command> ip access-group abcd out</command>
<command> ip access-list extended abcd</command>
<command> seq 5 permit tcp any any</command>
<command> seq 10 deny icmp any any</command>
<command> permit 1.1.1.2</command>
</configuration>
</cli>
</request>
```

XML Error Conditions and Reporting

This section contains examples of various error conditions that might occur in an XML transaction, and the associated responses that the XML generates. Note also, as shown below by the “NO_ERROR” message, that the same response message format is used for a successful configuration request.

The general form of the response is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<response MajorVersion="1" MinorVersion="0">
<responseType></responseType>
<responseSeverity></responseSeverity>
<responseMsg></responseMsg>
</response>
```

Summary of XML Limitations

- The XML response to a **show running-configuration** request is encoded in one single XML tag, instead of the standard XML-encoded response.
- A **show** command, in an XML request, requires <action> for the operation tag; the request is not supported if <configuration> is used for the operation tag.
- Only allowed one show command is supported within a single XML request.
- XML namespace is not supported.

Error Messages

The following strings can appear after the <responseType> tag:

- XML_PARSE_ERROR
- CLI_PARSE_ERROR—This error is caused by:

  - Malformed XML or mismatched XML tags
  - Invalid CLI commands or keywords
  - Invalid range of data specified in the CLI command
- XML_SCHEMA_ERROR—This error is caused by:
  - Invalid XML method or operation tags
  - Invalid object hierarchy or value out of range
- APPLICATION_ERROR—This error is caused by a failure to process the request, or a problem on the FTOS task.
- NO_ERROR—The XML request processed successfully.

The following strings can appear after the <responseSeverity> tag:

- SEVERITY_INFO—This string indicates no error, and is paired with NO_ERROR after the <responseType> tag.
- SEVERITY_ERROR—This string is paired with one of the other four possible <responseType> strings besides NO_ERROR.

The following strings can appear after the <responseMsg> tag:

- “Xml request successfully processed” (paired with NO_ERROR)
- “% Error: Parsing error is detected in the XML request” (paired with XML_PARSE_ERROR)
- “% Error: Schema error is detected in the XML request” (paired with XML_SCHEMA_ERROR)
- “% Error: CLI Parsing error is detected in the XML request” (paired with CLI_PARSE_ERROR)
- “% Error: [content varies, depending on the error]” (paired with APPLICATION_ERROR, indicating an application error from a backend task)

Examples of Error Conditions

XML parsing error

The following XML request is missing the XML declaration (the first line in the schema):

```
<request MajorVersion="1" MinorVersion="0">
<cli>
<configuration>
<command>ip access standard test2</command>
</configuration>
</cli>
</request>
```

The XML response to that malformed request is:

```
<?xml version="1.0" encoding="UTF-8"?>
<response MajorVersion="1" MinorVersion="0">
<responseType>XML_PARSE_ERROR</responseType>
<responseSeverity>SEVERITY_ERROR</responseSeverity>
<responseMsg>% Error: Parsing error detected in the XML request.</responseMsg>
</response>
```

XML schema error

This following XML request has transposed the <configuration> and <cli> tag sets:

```
<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<configuration>
<cli>
<command>ip access standard test2</command>
</cli>
</configuration>
</request>
```

The XML response to that malformed request is:

```
<?xml version="1.0" encoding="UTF-8"?>
<response MajorVersion="1" MinorVersion="0">
<responseType>XML_SCHEMA_ERROR</responseType>
<responseSeverity>SEVERITY_ERROR</responseSeverity>
<responseMsg>% Error: Schema error detected in the XML request.</responseMsg>
</response>
```

XML command error

The following XML request contains an invalid CLI command:

```
<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<cli>
<configuration>
<command>ip access test test1</command>
</configuration>
</cli>
</request>
```

The XML response to that invalid request is:

```
<?xml version="1.0" encoding="UTF-8"?>
<response MajorVersion="1" MinorVersion="0">
<responseType>CLI_PARSE_ERROR</responseType>
<responseSeverity>SEVERITY_ERROR</responseSeverity>
<responseMsg><command>ip access test test1</command></responseMsg>
</response>
```

XML application error

The command in this XML request makes an invalid request:

```
<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<cli>
<configuration>
<command>ip access standard test1</command>
<command>seq 10 permit host 1.2.3.4 log count bytes</command>
</configuration>
</cli>
</request>
```

The error response contains a <responseSeverity> of “APPLICATION_ERROR”, <responseSeverity>SEVERITY_ERROR, and a <responseMsg> of “% Error: Seq number does not exist.”

The second command in this XML request also makes an invalid request:

```
<?xml version="1.0" encoding="UTF-8"?>
<request MajorVersion="1" MinorVersion="0">
<cli>
<configuration>
<command>ip access standard test1</command>
<command>no permit host 2.2.3.4 log count bytes</command>
</configuration>
</cli>
</request>
```

The error response contains a <responseSeverity> of “APPLICATION_ERROR”, <responseSeverity> of “APPLICATION_ERROR” and a <responseMsg> of “% Error: Access-list entry does not exist.”

Using display xml as a Pipe Option

Also, at a CLI prompt in EXEC privilege mode (“enable mode”), you can retrieve XML-formatted responses to the show commands supported by XML (see the list of supported show commands in the section XML Functionality on page 603). The following table describes how to format a show command with a pipe option that will request that the show command report be presented with XML formatting.

Command Syntax	Command Mode	Purpose
show *keyword* **| display xml**	EXEC privilege	FTOS treats “ | **display xml**” as a request to format the show command report in XML format.

As shown in the following Figure 382, FTOS formats the response with the XML tags from the same response schema used by the XML response, discussed in The “Show” Request and Response on page 606. For more on pipe options, see Filtering show Command Outputs on page 55.

```
Force10>#show linecard 0 | display xml
<?xml version="1.0" encoding="UTF-8" ?>
<response MajorVersion="1" MinorVersion="0">
<action>
<linecard>
<slotId>0</slotId>
<status>online</status>
<nextBoot>online</nextBoot>
<reqType>EXW2PF3 - 2-port 10GE LAN/WAN PHY line card with XFP optics (EF3)</reqType>
<curType>EXW2PF3 - 2-port 10GE LAN/WAN PHY line card with XFP optics (EF3)</curType>
<hwRevBase>1.1</hwRevBase>
<hwRevPortPipe0>1.1</hwRevPortPipe0>
<hwRevPortPipe1>n/a</hwRevPortPipe1>
<numPorts>2</numPorts>
<upTime>1 hr, 32 min</upTime>
<swVer>4.4.3.243</swVer>
<lcJumboCapable>yes</lcJumboCapable>
<lcBootFlashA>2.3.0.6 [booted]</lcBootFlashA>
<lcBootFlashB>2.3.0.6  </lcBootFlashB>
<totMemSize>268435456</totMemSize>
<lcTemperature>37</lcTemperature>
<powerStatus>AC</powerStatus>
<voltage>ok</voltage>
<serialNum>0039034</serialNum>
<partNum>7520017400</partNum>
<productRev>08</productRev>
<vendorId>04</vendorId>
<dateCode>01332005</dateCode>
<countryCode>01</countryCode>
</linecard>
</action>
</response>
Force10>
```

Figure 382 Example: show linecard 0 | display xml

Chapter 37

Bidirectional Forwarding Detection

C-Series	NO
E-Series	✓

Platform Specific Feature: Bidirectional Forwarding Detection (BFD) is supported on E-Series only.

Protocol Overview

Bidirectional Forwarding Detection (BFD) is a protocol that is used to rapidly detect communication failures between two adjacent systems. It is a simple and lightweight replacement for existing routing protocol link state detection mechanisms. It also provides a failure detection solution for links on which no routing protocol is used.

BFD is a simple hello mechanism. Two neighboring systems running BFD establish a session using a three-way handshake. After the session has been established, the systems exchange periodic control packets at sub-second intervals. If a system does not receive a hello packet within a specified amount of time, routing protocols are notified that the forwarding path is down.

BFD provides forwarding path failure detection times on the order of milliseconds rather than seconds as with conventional routing protocol hellos. It is independent of routing protocols, and as such provides a consistent method of failure detection when used across a network. Networks converge faster because BFD triggers link state changes in the routing protocol sooner and more consistently, because BFD can eliminate the use of multiple protocol-dependent timers and methods.

BFD also carries less overhead than routing protocol hello mechanisms. Control packets can be encapsulated in any form that is convenient, and, on Force10 routers, sessions are maintained by BFD Agents that reside on the line card, which frees resources on the RPM. Only session state changes are reported to the BFD Manager (on the RPM), which in turn notifies the routing protocols that are registered with it.

BFD is an independent protocol and generic protocol, which all media, topologies, and routing protocols can support using any encapsulation. Force10 has implemented BFD at Layer 3 and uses UDP encapsulation. BFD functionality will be implemented in phases. As of FTOS version 7.5.1.0, OSPF, IS-IS, VRRP, VLANs, LAGs, static routes, and physical ports support BFD, based on the IETF internet draft *draft-ietf-bfd-base-03*.

How BFD Works

Two neighboring systems running BFD establish a session using a three-way handshake. After the session has been established, the systems exchange control packets at agreed upon intervals. In addition, systems send a control packet anytime there is a state change or change in a session parameter; these control packets are sent without regard to transmit and receive intervals.

Note: FTOS does not support multi-hop BFD sessions.

If a system does not receive a control packet within an agreed-upon amount of time, the BFD Agent changes the session state to Down. It then notifies the BFD Manager of the change, and sends a control packet to the neighbor that indicates the state change (though it might not be received if the link or receiving interface is faulty). The BFD Manager notifies the routing protocols that are registered with it (clients) that the forwarding path is down, and a link state change is triggered in all protocols.

Note: A session state change from Up to Down is the only state change that triggers a link state change in the routing protocol client.

BFD Packet Format

Control packets are encapsulated in UDP packets. Figure 383 shows the complete encapsulation of a BFD control packet inside an IPv4 packet.

Figure 383 BFD in IPv4 Packet Format

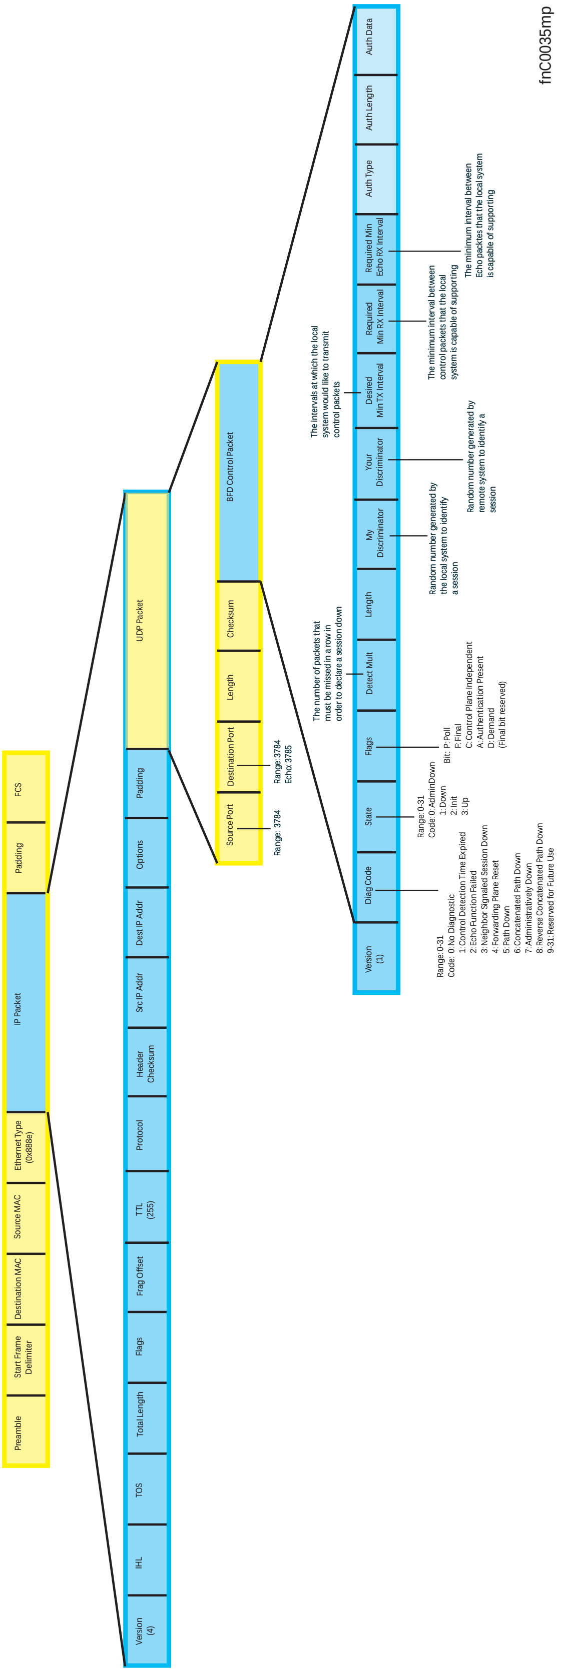
Preamble
Start Frame Delimiter
Destination MAC
Source MAC
Ethernet Type (0x888e)
IP Packet
Padding
FCS
Version (4)
IHL
TOS
Total Length
Flags
Frag Offset
TTL (255)
Protocol
Header Checksum
Src IP Addr
Dest IP Addr
Options
Padding
UDP Packet
Source Port
Destination Port
Length
Checksum
BFD Control Packet
Range: 3784
Range: 3784
Echo: 3785
Version (1)
Diag Code
State
Flags
Detect Mult
Length
My Discriminator
Your Discriminator
Desired Min TX Interval
Required Min RX Interval
Required Min Echo RX Interval
Auth Type
Auth Length
Auth Data
Range: 0-31
Code: 0: No Diagnostic
1: Control Detection Time Expired
2: Echo Function Failed
3: Neighbor Signaled Session Down
4: Forwarding Plane Reset
5: Path Down
6: Concatenated Path Down
7: Administratively Down
8: Reverse Concatenated Path Down
9-31: Reserved for Future Use
Range: 0-31
Code: 0: AdminDown
1: Down
2: Init
3: Up
Bit: P: Poll
F: Final
C: Control Plane Independent
A: Authentication Present
D: Demand
(Final bit reserved)
The number of packets that must be missed in a row in order to declare a session down
Random number generated by the local system to identify a session
Random number generated by remote system to identify a session
The intervals at which the local system would like to transmit control packets
The minimum interval between control packets that the local system is capable of supporting
The minimum interval between Echo packtes that the local system is capable of supporting
fnC0035mp

Table 50 BFD Packet Fields

Field	Description
Diagnostic Code	The reason that the last session failed.
State	The current local session state. See BFD Sessions.
Flag	A bit that indicates packet function. If the poll bit is set, the receiving system must respond as soon as possible, without regard to its transmit interval. The responding system clears the poll bit and sets the final bit in its response. The poll and final bits are used during the handshake and Demand mode (see BFD Sessions). **Note:** FTOS does not currently support multi-point sessions, Demand mode, authentication, or control plane independence; these bits are always clear.
Detection Multiplier	The number of packets that must be missed in order to declare a session down.
Length	The entire length of the BFD packet.
My Discriminator	A random number generated by the local system to identify the session.
Your Discriminator	A random number generated by the remote system to identify the session. Discriminator values are necessary to identify the session to which a control packet belongs since there can be many sessions running on a single interface.
Desired Min TX Interval	The minimum rate at which the local system would like to send control packets to the remote system.
Required Min RX Interval	The minimum rate at which the local system would like to receive control packets from the remote system.
Required Min Echo RX	The minimum rate at which the local system would like to receive echo packets. **Note:** FTOS does not currently support the echo function.
Authentication Type Authentication Length Authentication Data	An optional method for authenticating control packets. **Note:** FTOS does not currently support the BFD authentication function.

Two important parameters are calculated using the values contained in the control packet.

- **Transmit interval —** Transmit interval is the agreed-upon rate at which a system sends control packets. Each system has its own transmit interval, which is the greater of the last received remote Desired TX Interval and the local Required Min RX Interval.
- **Detection time —** Detection time is the amount of time that a system does not receive a control packet, after which the system determines that the session has failed. Each system has its own detection time.
    - In Asynchronous mode: Detection time is the remote Detection Multiplier multiplied by greater of the remote Desired TX Interval and the local Required Min RX Interval.
    - In Demand mode: Detection time is the local Detection Multiplier multiplied by the greater of the local Desired Min TX and the remote Required Min RX Interval.

BFD Sessions

BFD must be enabled on both sides of a link in order to establish a session. The two participating systems can assume either of two roles:

- **Active**—The active system initiates the BFD session. Both systems can be active for the same session.
- **Passive**—The passive system does not initiate a session. It only responds to a request for session initialization from the active system.

A BFD session has two modes:

- **Asynchronous mode**—In Asynchronous mode, both systems send periodic control messages at an agreed upon interval to indicate that their session status is Up.
- **Demand mode**—If one system requests Demand mode, the other system stops sending periodic control packets; it only sends a response to status inquiries from the Demand mode initator. Either system (but not both) can request Demand mode at any time.

Note: FTOS supports currently supports asychronous mode only.

A session can have four states: Adminstratively Down, Down, Init, and Up.

- **Administratively Down**—The local system will not participate in a particular session.
- **Down**—The remote system is not sending any control packets or at least not within the detection time for a particular session.
- **Init**—The local system is communicating.
- **Up**—The both systems are exchanging control packets.

The session is declared down if:

- A control packet is not received within the detection time.
- Sufficient echo packets are lost.
- Demand mode is active and a control packet is not received in response to a poll packet.

BFD Three-way Handshake

A three-way handshake must take place between the systems that will participate in the BFD session. The handshake shown in Figure 384 assumes that there is one active and one passive system, and that this is the first session established on this link. The default session state on both ports is Down.

1. The active system sends a steady stream of control packets that indicates that its session state is Down, until the passive system responds. These packets are sent at the desired transmit interval of the Active system, and the Your Discriminator field is set to zero.
2. When the passive system receives any of these control packets, it changes its session state to Init, and sends a response that indicates its state change. The response includes its session ID in the My Discriminator field, and the session ID of the remote system in the Your Discriminator field.

3. The active system receives the response from the passive system, and changes its session state to Up. It then sends a control packet indicating this state change. This is the third and final part of of the handshake. At this point, the discriminator values have been exchanged, and the transmit intervals have been negotiated.
4. The passive system receives the control packet, changes its state to Up. Both systems agree that a session has been established. However, since both members must send a control packet—that requires a response—anytime there is a state change or change in a session parameter, the passive system sends a final response indicating the state change. After this, periodic control packets are exchanged.

Figure 384 BFD Three-way Handshake

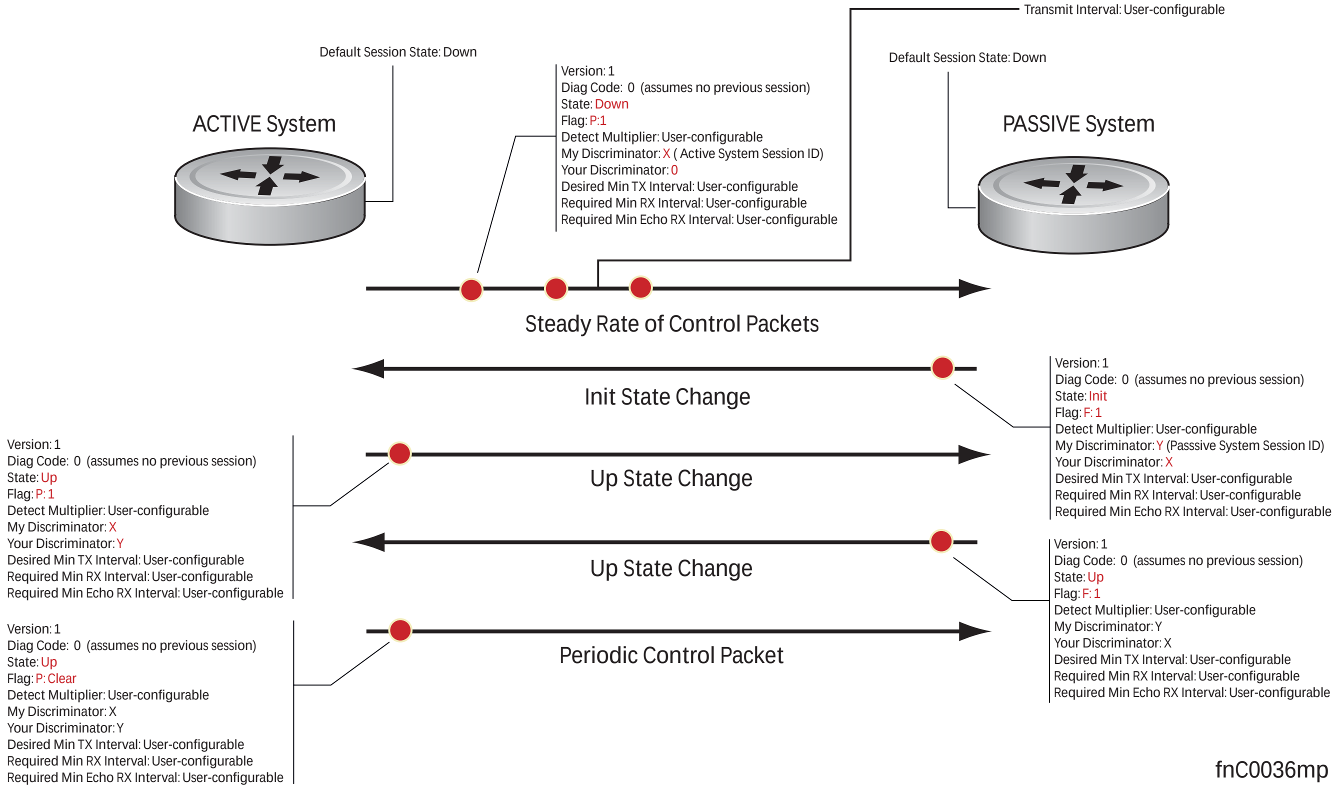
Transmit Interval: User-configurable
Default Session State: Down
Default Session State: Down
ACTIVE System
PASSIVE System
Steady Rate of Control Packets
Init State Change
Up State Change
Up State Change
Periodic Control Packet
fnC0036mp

Version: 1
Diag Code: 0 (assumes no previous session)
State: Down
Flag: P:1
Detect Multiplier: User-configurable
My Discriminator: X (Active System Session ID)
Your Discriminator: 0
Desired Min TX Interval: User-configurable
Required Min RX Interval: User-configurable
Required Min Echo RX Interval: User-configurable

Version: 1
Diag Code: 0 (assumes no previous session)
State: Init
Flag: F: 1
Detect Multiplier: User-configurable
My Discriminator: Y (Passive System Session ID)
Your Discriminator: X
Desired Min TX Interval: User-configurable
Required Min RX Interval: User-configurable
Required Min Echo RX Interval: User-configurable

Version: 1
Diag Code: 0 (assumes no previous session)
State: Up
Flag: P: 1
Detect Multiplier: User-configurable
My Discriminator: X
Your Discriminator: Y
Desired Min TX Interval: User-configurable
Required Min RX Interval: User-configurable
Required Min Echo RX Interval: User-configurable

Version: 1
Diag Code: 0 (assumes no previous session)
State: Up
Flag: F: 1
Detect Multiplier: User-configurable
My Discriminator: Y
Your Discriminator: X
Desired Min TX Interval: User-configurable
Required Min RX Interval: User-configurable
Required Min Echo RX Interval: User-configurable

Version: 1
Diag Code: 0 (assumes no previous session)
State: Up
Flag: P: Clear
Detect Multiplier: User-configurable
My Discriminator: X
Your Discriminator: Y
Desired Min TX Interval: User-configurable
Required Min RX Interval: User-configurable
Required Min Echo RX Interval: User-configurable

Session State Changes

Figure 385 shows how the session state on a system changes based on the status notification it receives from the remote system. For example, if a session on a system is down, and it receives a Down status notification from the remote system, the session state on the local system changes to Init.

Figure 385 BFD State Machine

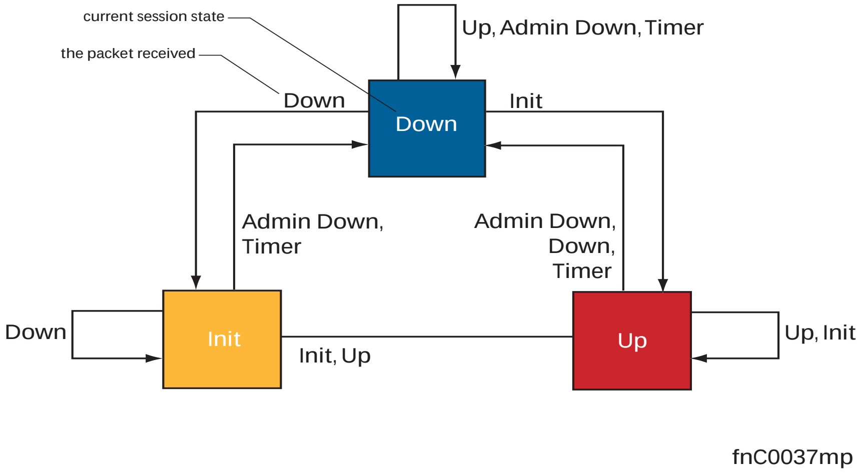

Important Points to Remember

- BFD is supported on E-Series TeraScale only.
- FTOS supports a maximum of 100 sessions per BFD agent.
- BFD must be enabled on both ends of a link.
- Demand mode, authentication, and the Echo function are not supported.
- BFD is not supported on multi-hop and virtual links.
- Protocol Liveness is supported for routing protocols only.
- FTOS supports only OSPF, ISIS, and VRRP as BFD clients.

Configuring Bidirectional Forwarding Detection

The remainder of this chapter is divided into the following sections:

- Configuring BFD for Physical Ports on page 622
- Configuring BFD for Static Routes on page 626
- Configuring BFD for OSPF on page 628
- Configuring BFD for IS-IS on page 631
- Configuring BFD for VRRP on page 634
- Configuring BFD for VLANs on page 637
- Configuring BFD for Port-Channels on page 640
- Configuring Protocol Liveness on page 642
- TroubleShooting BFD on page 643

Configuring BFD for Physical Ports

BFD on physical ports is useful when no routing protocol is enabled. Without BFD, if the remote system fails, the local system does not remove the connected route until the first failed attempt to send a packet. When BFD is enabled, the local system removes the route as soon as it stops receiving periodic control packets from the remote system.

Configuring BFD for a physical port is a two-step process:

1. Enable BFD globally. See page 622.
2. Establish a session with a next-hop neighbor. See page 622.

Related Configuration Tasks

- Change session parameters. See page 624.
- Disable or re-enable BFD on an interface. See page 625.

Enabling BFD Globally

BFD must be enabled globally on both routers, as shown in Figure 387.

To enable BFD globally:

Step	Task	Command Syntax	Command Mode
1	Enable BFD globally.	**bfd enable**	CONFIGURATION

Verify that BFD is enabled globally using the command **show running bfd**, as shown in Figure 386.

Figure 386 Enabling BFD Globally

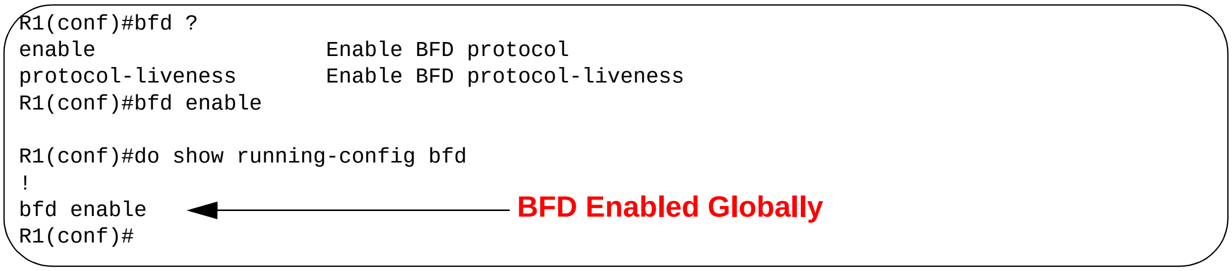

```
R1(conf)#bfd ?
enable                    Enable BFD protocol
protocol-liveness         Enable BFD protocol-liveness
R1(conf)#bfd enable

R1(conf)#do show running-config bfd
!
bfd enable
R1(conf)#
```

Establishing a Session on Physical Ports

To establish a session, BFD must be enabled at interface level on both ends of the link, as shown in Figure 387. The configuration parameters do not need to match.

Figure 387 Establishing a BFD Session for Physical Ports

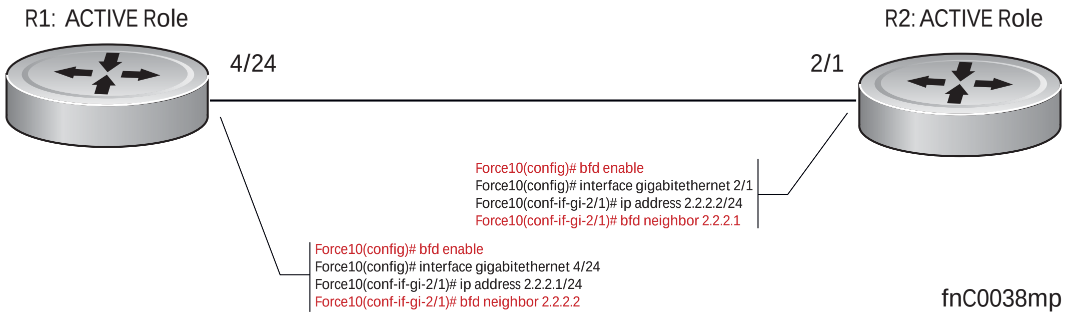

```
Force10(config)# bfd enable
Force10(config)# interface gigabitethernet 2/1
Force10(conf-if-gi-2/1)# ip address 2.2.2.2/24
Force10(conf-if-gi-2/1)# bfd neighbor 2.2.2.1
```

```
Force10(config)# bfd enable
Force10(config)# interface gigabitethernet 4/24
Force10(conf-if-gi-2/1)# ip address 2.2.2.1/24
Force10(conf-if-gi-2/1)# bfd neighbor 2.2.2.2
```

To establish a session:

Step	Task	Command Syntax	Command Mode
1	Enter interface mode	**interface**	CONFIGURATION
2	Assign an IP address to the interface if one is not already assigned.	**ip address** *ip-address*	INTERFACE
3	Identify the neighbor with which the interface will participate in the BFD session.	**bfd neighbor** *ip-address*	INTERFACE

Verify that the session is established using the command **show bfd neighbors**, as shown in Figure 388.

Figure 388 Viewing Established Sessions for Physical Ports

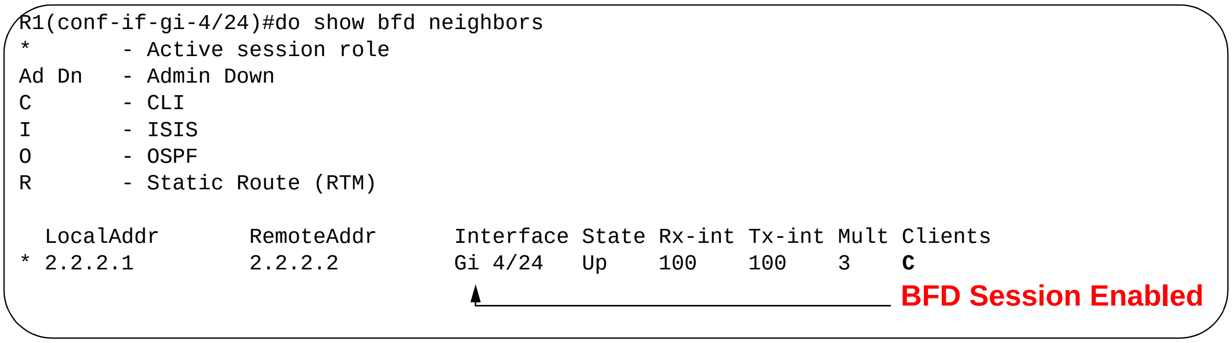

```
R1(conf-if-gi-4/24)#do show bfd neighbors
*        - Active session role
Ad Dn    - Admin Down
C        - CLI
I        - ISIS
O        - OSPF
R        - Static Route (RTM)

  LocalAddr        RemoteAddr        Interface State Rx-int Tx-int Mult Clients
* 2.2.2.1          2.2.2.2           Gi 4/24   Up    100    100    3    C
```

The command **show bfd neighbors detail** shows more specific information about BFD sessions (Figure 389).

Figure 389 Viewing Session Details

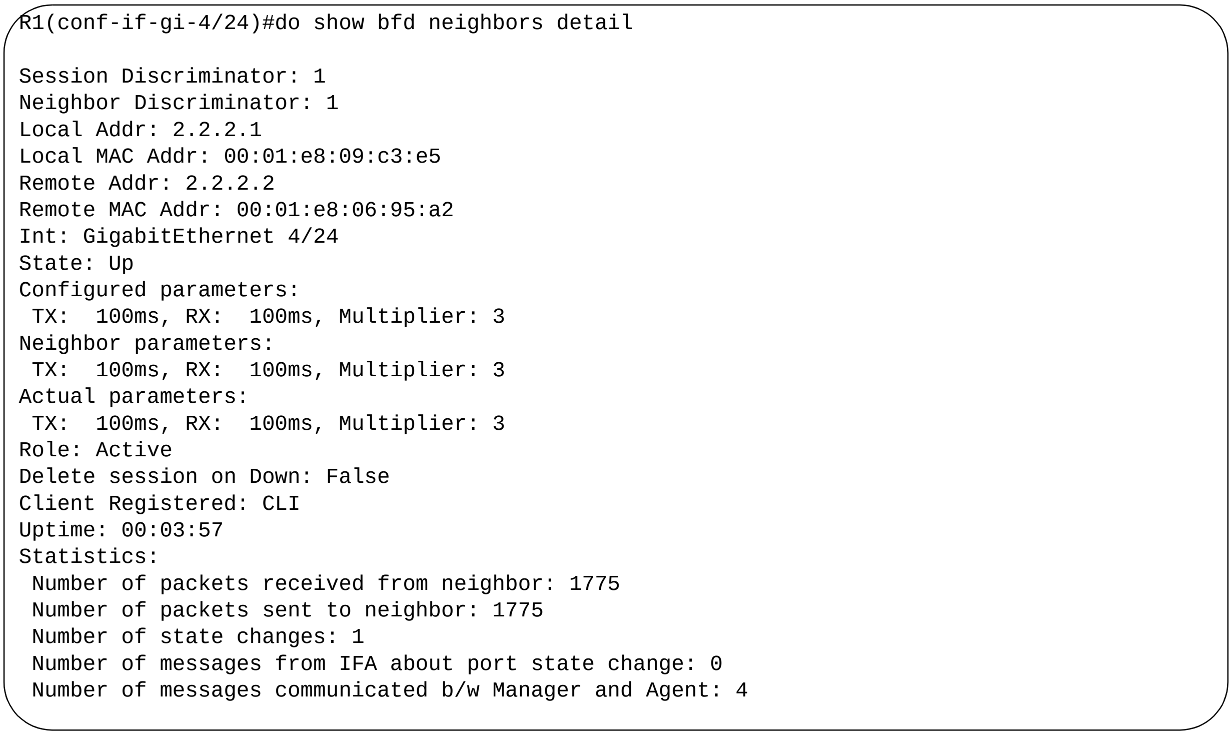

```
R1(conf-if-gi-4/24)#do show bfd neighbors detail

Session Discriminator: 1
Neighbor Discriminator: 1
Local Addr: 2.2.2.1
Local MAC Addr: 00:01:e8:09:c3:e5
Remote Addr: 2.2.2.2
Remote MAC Addr: 00:01:e8:06:95:a2
Int: GigabitEthernet 4/24
State: Up
Configured parameters:
 TX:  100ms, RX:  100ms, Multiplier: 3
Neighbor parameters:
 TX:  100ms, RX:  100ms, Multiplier: 3
Actual parameters:
 TX:  100ms, RX:  100ms, Multiplier: 3
Role: Active
Delete session on Down: False
Client Registered: CLI
Uptime: 00:03:57
Statistics:
 Number of packets received from neighbor: 1775
 Number of packets sent to neighbor: 1775
 Number of state changes: 1
 Number of messages from IFA about port state change: 0
 Number of messages communicated b/w Manager and Agent: 4
```

When both interfaces are configured for BFD, log messages are displayed indicating state changes, as shown in Message 8.

Message 8 BFD Session State Changes

```
R1(conf-if-gi-4/24)#00:36:01: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE: Changed session state to Down for
neighbor 2.2.2.2 on interface Gi 4/24 (diag: 0)
00:36:02: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE: Changed session state to Up for neighbor 2.2.2.2 on
interface Gi 4/24 (diag: 0)
```

Changing Physical Port Session Parameters

BFD sessions are configured with default intervals and a default role (active). The parameters that can be configured are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. These parameters are configured per interface; if you change a parameter, the change affects all physical port sessions on that interface.

To change session parameters on an interface:

Step	Task	Command Syntax	Command Mode
1	Change session parameters for all sessions on an interface.	**bfd interval** *milliseconds* **min_rx** *milliseconds* **multiplier** *value* **role** [**active** | **passive**]	INTERFACE

View session parameters using the **show bfd neighbors detail** command.

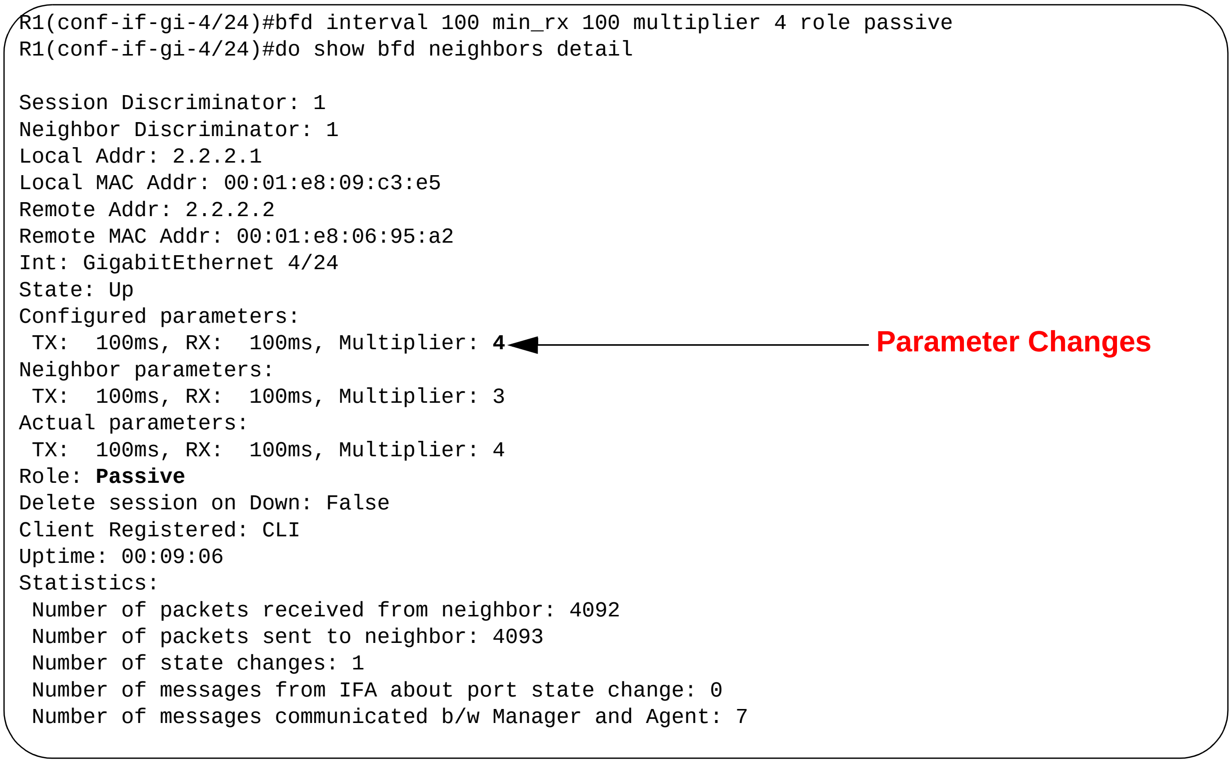

```
R1(conf-if-gi-4/24)#bfd interval 100 min_rx 100 multiplier 4 role passive
R1(conf-if-gi-4/24)#do show bfd neighbors detail

Session Discriminator: 1
Neighbor Discriminator: 1
Local Addr: 2.2.2.1
Local MAC Addr: 00:01:e8:09:c3:e5
Remote Addr: 2.2.2.2
Remote MAC Addr: 00:01:e8:06:95:a2
Int: GigabitEthernet 4/24
State: Up
Configured parameters:
 TX:  100ms, RX:  100ms, Multiplier: 4
Neighbor parameters:
 TX:  100ms, RX:  100ms, Multiplier: 3
Actual parameters:
 TX:  100ms, RX:  100ms, Multiplier: 4
Role: Passive
Delete session on Down: False
Client Registered: CLI
Uptime: 00:09:06
Statistics:
 Number of packets received from neighbor: 4092
 Number of packets sent to neighbor: 4093
 Number of state changes: 1
 Number of messages from IFA about port state change: 0
 Number of messages communicated b/w Manager and Agent: 7
```

Figure 390 Changing Session Parameters for Physical Ports

Disabling and Re-enabling BFD

BFD is enabled on all interfaces by default, though sessions are not created unless explicitly configured. If BFD is disabled, all of the sessions on that interface are placed in an Administratively Down state (Message 9), and the remote systems are notified of the session state change (Message 10).

To disable BFD on an interface:

Step	Task	Command Syntax	Command Mode
1	Disable BFD on an interface.	**no bfd enable**	INTERFACE

Message 9 Disabling BFD on a Local Interface

```
R1(conf-if-gi-4/24)#01:00:52: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE: Changed session state to Ad Dn for
neighbor 2.2.2.2 on interface Gi 4/24 (diag: 0)
```

Message 10 Remote System State Change due to Local State Admin Down

```
R2>01:32:53: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE: Changed session state to Down for neighbor 2.2.2.1
on interface Gi 2/1 (diag: 7)
```

To re-enable BFD on an interface:

Step	Task	Command Syntax	Command Mode
1	Enable BFD on an interface.	**bfd enable**	INTERFACE

Configuring BFD for Static Routes

BFD gives systems a link state detection mechanism for static routes. With BFD, systems are notified to remove static routes from the routing table as soon as the link state change occurs, rather than having to wait until packets fail to reach their next hop.

Configuring BFD for static routes is a three-step process:

1. Enable BFD globally. See Enabling BFD Globally on page 622.
2. On the local system, establish a session with the next hop of a static route. See page 626.
3. On the remote system, establish a session with the physical port that is the origin of the static route. See Establishing a Session on Physical Ports on page 622.

Related Configuration Tasks

- Change session parameters. See page 627.
- Disable BFD for all static routes. See page 627.

Establishing Sessions for Static Routes

Sessions are established for all neighbors that are the next hop of a static route.

Figure 391 Enabling BFD for Static Routes

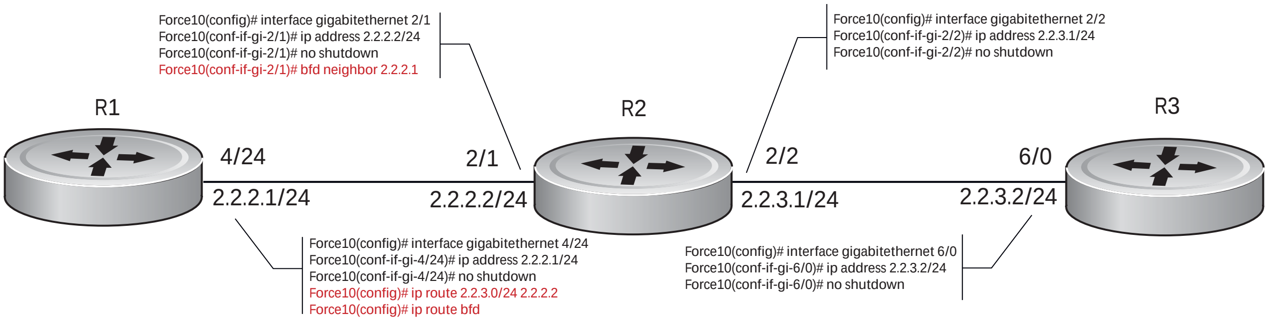

```
Force10(config)# interface gigabitethernet 2/1
Force10(conf-if-gi-2/1)# ip address 2.2.2.2/24
Force10(conf-if-gi-2/1)# no shutdown
Force10(conf-if-gi-2/1)# bfd neighbor 2.2.2.1
```

```
Force10(config)# interface gigabitethernet 2/2
Force10(conf-if-gi-2/2)# ip address 2.2.3.1/24
Force10(conf-if-gi-2/2)# no shutdown
```

```
Force10(config)# interface gigabitethernet 4/24
Force10(conf-if-gi-4/24)# ip address 2.2.2.1/24
Force10(conf-if-gi-4/24)# no shutdown
Force10(config)# ip route 2.2.3.0/24 2.2.2.2
Force10(config)# ip route bfd
```

```
Force10(config)# interface gigabitethernet 6/0
Force10(conf-if-gi-6/0)# ip address 2.2.3.2/24
Force10(conf-if-gi-6/0)# no shutdown
```

fnC0039mp

To establish a BFD session:

Step	Task	Command Syntax	Command Mode
1	Establish BFD sessions for all neighbors that are the next hop of a static route.	**ip route bfd**	CONFIGURATION

Verify that sessions have been created for static routes using the command **show bfd neighbors**, as shown in Figure 392. View detailed session information using the command **show bfd neighbors detail**, as shown in Figure 390.

Figure 392 Viewing Established Sessions for Static Routes

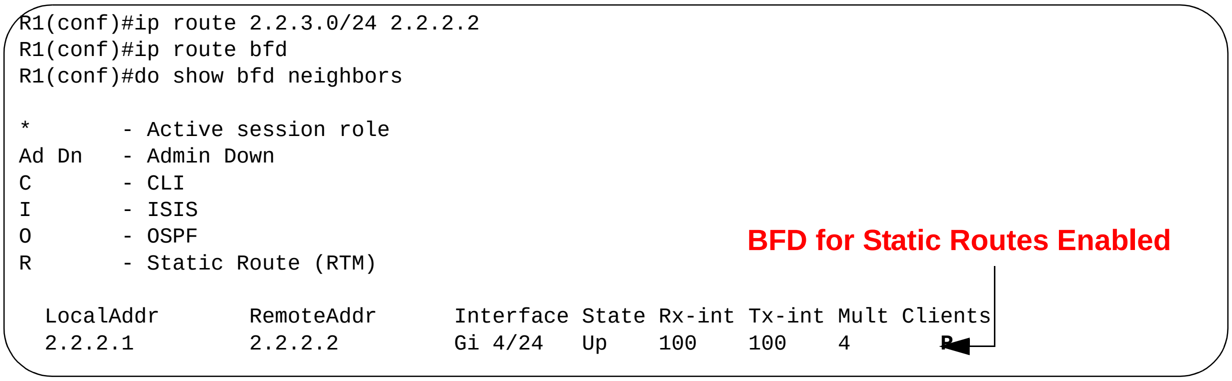

```
R1(conf)#ip route 2.2.3.0/24 2.2.2.2
R1(conf)#ip route bfd
R1(conf)#do show bfd neighbors

*       - Active session role
Ad Dn   - Admin Down
C       - CLI
I       - ISIS
O       - OSPF
R       - Static Route (RTM)

  LocalAddr       RemoteAddr      Interface State Rx-int Tx-int Mult Clients
  2.2.2.1         2.2.2.2         Gi 4/24   Up    100    100    4       R
```

BFD for Static Routes Enabled

Changing Static Route Session Parameters

BFD sessions are configured with default intervals and a default role. The parameters that can be configured are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. These parameters are configured for all static routes; if you change a parameter, the change affects all sessions for static routes.

To change parameters for static route sessions:

Step	Task	Command Syntax	Command Mode
1	Change parameters for all static route sessions.	**ip route bfd interval** *milliseconds* **min_rx** *milliseconds* **multiplier** *value* **role** [**active** | **passive**]	CONFIGURATION

View session parameters using the command **show bfd neighbors detail**, as shown in Figure 390 on page 625.

Disabling BFD for Static Routes

If BFD is disabled, all static route BFD sessions are torn down. A final Admin Down packet is sent to all neighbors on the remote systems, and those neighbors change to the Down state (Message 10 on page 625).

To disable BFD for static routes:

Step	Task	Command Syntax	Command Mode
1	Disable BFD for static routes.	**no ip route bfd**	CONFIGURATION

Configuring BFD for OSPF

When using BFD with OSPF, the OSPF protocol registers with the BFD manager on the RPM. BFD sessions are established with all neighboring interfaces participating in OSPF. If a neighboring interface fails, the BFD agent on the line card notifies the BFD manager, which in turn notifies the OSPF protocol that a link state change occurred.

Configuring BFD for OSPF is a two-step process:

1. Enable BFD globally. See Enabling BFD Globally on page 622.
2. Establish sessions for all or particular OSPF neighbors. See page 629.

Related Configuration Tasks

- Change session parameters. See page 630.
- Disable BFD sessions for OSPF. See page 630.

Establishing Sessions with OSPF Neighbors

BFD sessions can be established with all OSPF neighbors at once, or sessions can be established with all neighbors out of a specific interface. Sessions are only established when the OSPF adjacency is in the full state.

Figure 393 Establishing Sessions with OSPF Neighbors

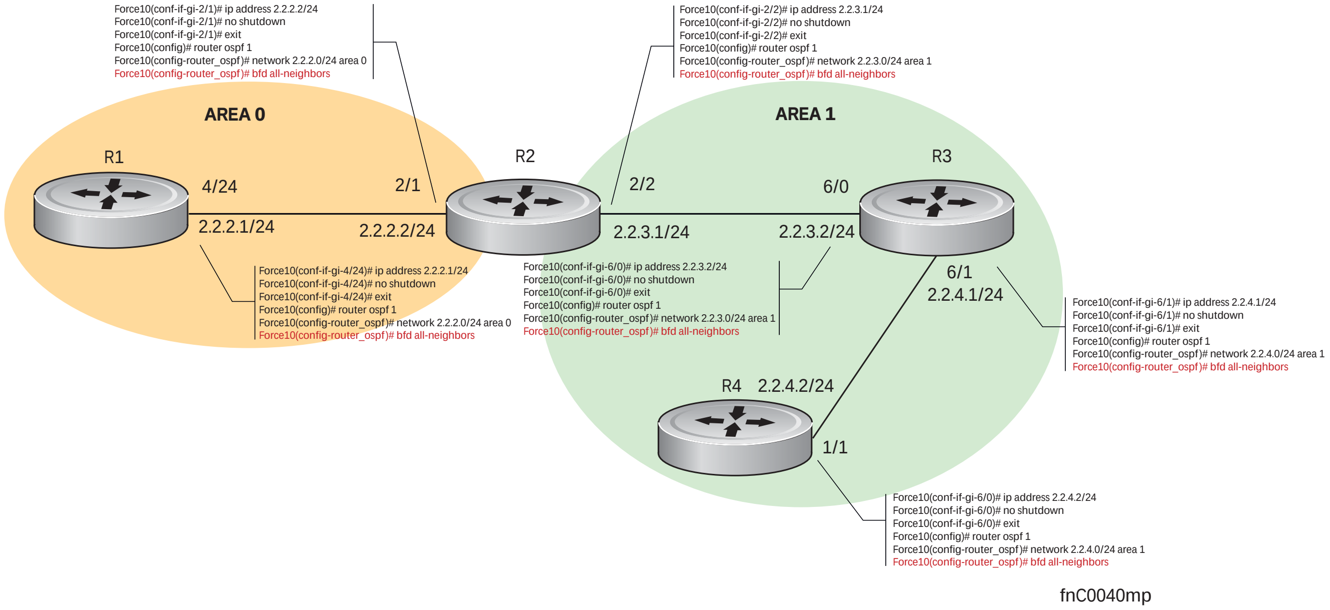

```
Force10(conf-if-gi-2/1)# ip address 2.2.2.2/24
Force10(conf-if-gi-2/1)# no shutdown
Force10(conf-if-gi-2/1)# exit
Force10(config)# router ospf 1
Force10(config-router_ospf)# network 2.2.2.0/24 area 0
Force10(config-router_ospf)# bfd all-neighbors
```

```
Force10(conf-if-gi-2/2)# ip address 2.2.3.1/24
Force10(conf-if-gi-2/2)# no shutdown
Force10(conf-if-gi-2/2)# exit
Force10(config)# router ospf 1
Force10(config-router_ospf)# network 2.2.3.0/24 area 1
Force10(config-router_ospf)# bfd all-neighbors
```

```
Force10(conf-if-gi-4/24)# ip address 2.2.2.1/24
Force10(conf-if-gi-4/24)# no shutdown
Force10(conf-if-gi-4/24)# exit
Force10(config)# router ospf 1
Force10(config-router_ospf)# network 2.2.2.0/24 area 0
Force10(config-router_ospf)# bfd all-neighbors
```

```
Force10(conf-if-gi-6/0)# ip address 2.2.3.2/24
Force10(conf-if-gi-6/0)# no shutdown
Force10(conf-if-gi-6/0)# exit
Force10(config)# router ospf 1
Force10(config-router_ospf)# network 2.2.3.0/24 area 1
Force10(config-router_ospf)# bfd all-neighbors
```

```
Force10(conf-if-gi-6/1)# ip address 2.2.4.1/24
Force10(conf-if-gi-6/1)# no shutdown
Force10(conf-if-gi-6/1)# exit
Force10(config)# router ospf 1
Force10(config-router_ospf)# network 2.2.4.0/24 area 1
Force10(config-router_ospf)# bfd all-neighbors
```

```
Force10(conf-if-gi-6/0)# ip address 2.2.4.2/24
Force10(conf-if-gi-6/0)# no shutdown
Force10(conf-if-gi-6/0)# exit
Force10(config)# router ospf 1
Force10(config-router_ospf)# network 2.2.4.0/24 area 1
Force10(config-router_ospf)# bfd all-neighbors
```

To establish BFD with all OSPF neighbors:

Step	Task	Command Syntax	Command Mode
1	Establish sessions with all OSPF neighbors.	**bfd all-neighbors**	ROUTER-OSPF

To establish BFD for all OSPF neighbors on a single interface:

Step	Task	Command Syntax	Command Mode
1	Establish sessions with all OSPF neighbors on a single interface.	**ip ospf bfd all-neighbors**	INTERFACE

View the established sessions using the command **show bfd neighbors**, as shown in Figure 394.

Figure 394 Viewing Established Sessions for OSPF Neighbors

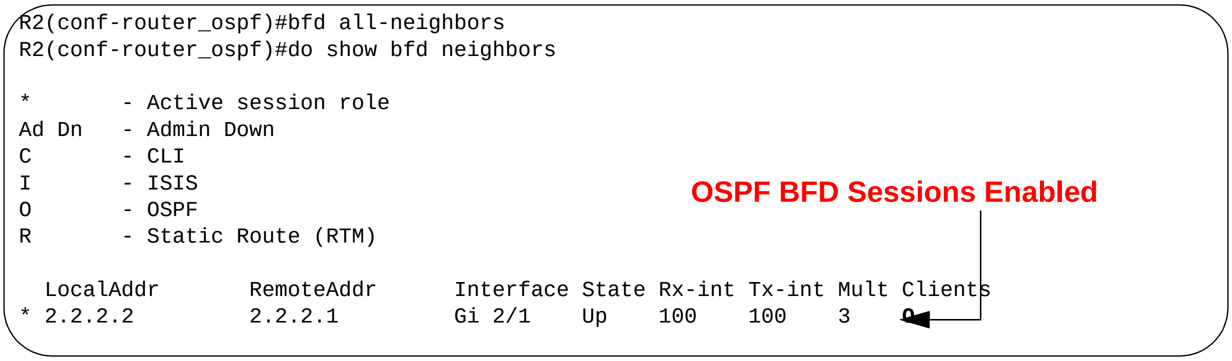

```
R2(conf-router_ospf)#bfd all-neighbors
R2(conf-router_ospf)#do show bfd neighbors

*       - Active session role
Ad Dn   - Admin Down
C       - CLI
I       - ISIS
O       - OSPF
R       - Static Route (RTM)

  LocalAddr       RemoteAddr      Interface State Rx-int Tx-int Mult Clients
* 2.2.2.2         2.2.2.1         Gi 2/1    Up    100    100    3    O
```

OSPF BFD Sessions Enabled

Changing OSPF Session Parameters

BFD sessions are configured with default intervals and a default role. The parameters that can be configured are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. These parameters are configured for all OSPF sessions or all OSPF sessions on a particular interface; if you change a parameter globally, the change affects all OSPF neighbors sessions. If you change a parameter at interface level, the change affects all OSPF sessions on that interface.

To change parameters for all OSPF sessions:

Step	Task	Command Syntax	Command Mode
1	Change parameters for OSPF sessions.	**bfd all-neighbors interval** *milliseconds* **min_rx** *milliseconds* **multiplier** *value* **role** [**active** | **passive**]	ROUTER-OSPF

To change parameters for OSPF sessions on an interface:

Step	Task	Command Syntax	Command Mode
1	Change parameters for all OSPF sessions on an interface.	**ip ospf bfd all-neighbors interval** *milliseconds* **min_rx** *milliseconds* **multiplier** *value* **role** [**active** | **passive**]	INTERFACE

View session parameters using the command **show bfd neighbors detail**, as shown in Figure 390 on page 625.

Disabling BFD for OSPF

If BFD is disabled globally, all sessions are torn down, and sessions on the remote system are placed in a Down state. If BFD is disabled on an interface, sessions on the interface are torn down, and sessions on the remote system are placed in a Down state (Message 10 on page 625). Disabling BFD does not trigger a change in BFD clients; a final Admin Down packet is sent before the session is terminated.

To disable BFD sessions with all OSPF neighbors:

Step	Task	Command Syntax	Command Mode
1	Disable BFD sessions with all OSPF neighbors.	**no bfd all-neighbors**	ROUTER-OSPF

To disable BFD sessions with all OSPF neighbors out of an interface:

Step	Task	Command Syntax	Command Mode
1	Disable BFD sessions with all OSPF neighbors out of an interface	**ip ospf bfd all-neighbors disable**	INTERFACE

Configuring BFD for IS-IS

When using BFD with IS-IS, the IS-IS protocol registers with the BFD manager on the RPM. BFD sessions are then established with all neighboring interfaces participating in IS-IS. If a neighboring interface fails, the BFD agent on the line card notifies the BFD manager, which in turn notifies the IS-IS protocol that a link state change occurred.

Configuring BFD for IS-IS is a two-step process:

1. Enable BFD globally. See Enabling BFD Globally on page 622.
2. Establish sessions for all or particular IS-IS neighbors. See page 632.

Related Configuration Tasks

- Change session parameters. See page 633.
- Disable BFD sessions for IS-IS. See page 633.

Establishing Sessions with IS-IS Neighbors

BFD sessions can be established for all IS-IS neighbors at once or sessions can be established for all neighbors out of a specific interface.

Figure 395 Establishing Sessions with IS-IS Neighbors

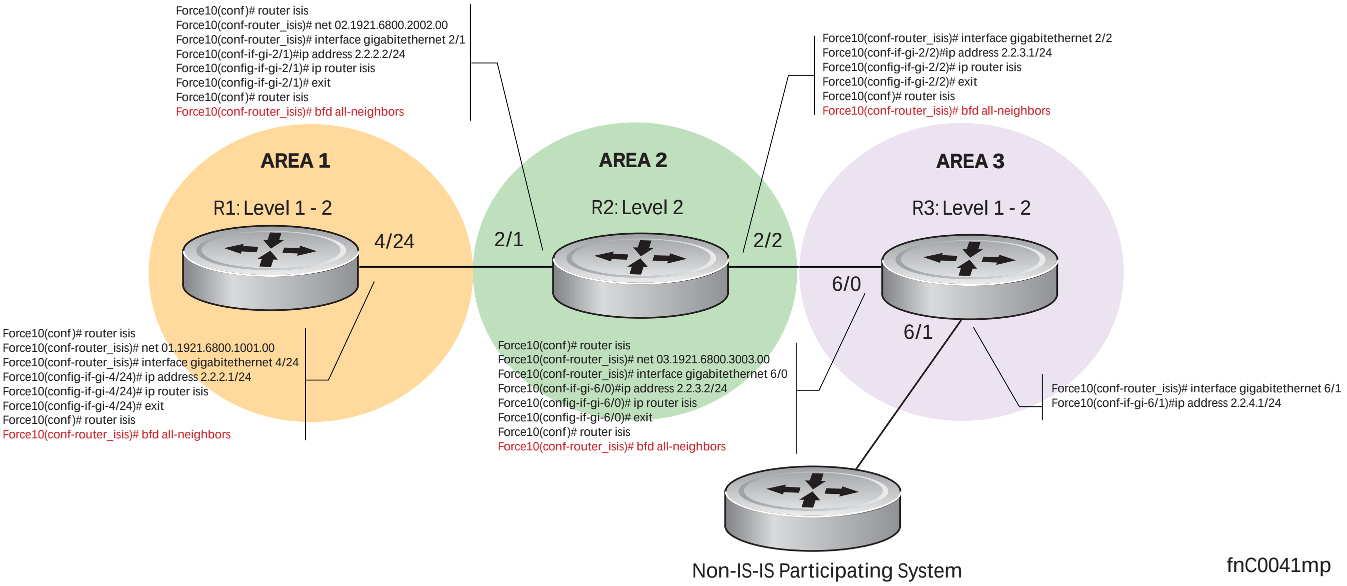

```
Force10(conf)# router isis
Force10(conf-router_isis)# net 02.1921.6800.2002.00
Force10(conf-router_isis)# interface gigabitethernet 2/1
Force10(conf-if-gi-2/1)#ip address 2.2.2.2/24
Force10(config-if-gi-2/1)# ip router isis
Force10(config-if-gi-2/1)# exit
Force10(conf)# router isis
Force10(conf-router_isis)# bfd all-neighbors
```

```
Force10(conf-router_isis)# interface gigabitethernet 2/2
Force10(conf-if-gi-2/2)#ip address 2.2.3.1/24
Force10(config-if-gi-2/2)# ip router isis
Force10(config-if-gi-2/2)# exit
Force10(conf)# router isis
Force10(conf-router_isis)# bfd all-neighbors
```

```
Force10(conf)# router isis
Force10(conf-router_isis)# net 01.1921.6800.1001.00
Force10(conf-router_isis)# interface gigabitethernet 4/24
Force10(config-if-gi-4/24)# ip address 2.2.2.1/24
Force10(config-if-gi-4/24)# ip router isis
Force10(config-if-gi-4/24)# exit
Force10(conf)# router isis
Force10(conf-router_isis)# bfd all-neighbors
```

```
Force10(conf)# router isis
Force10(conf-router_isis)# net 03.1921.6800.3003.00
Force10(conf-router_isis)# interface gigabitethernet 6/0
Force10(conf-if-gi-6/0)#ip address 2.2.3.2/24
Force10(config-if-gi-6/0)# ip router isis
Force10(config-if-gi-6/0)# exit
Force10(conf)# router isis
Force10(conf-router_isis)# bfd all-neighbors
```

```
Force10(conf-router_isis)# interface gigabitethernet 6/1
Force10(conf-if-gi-6/1)#ip address 2.2.4.1/24
```

To establish BFD with all IS-IS neighbors:

Step	Task	Command Syntax	Command Mode
1	Establish sessions with all IS-IS neighbors.	**bfd all-neighbors**	ROUTER-ISIS

To establish BFD with all IS-IS neighbors out of a single interface:

Step	Task	Command Syntax	Command Mode
1	Establish sessions with all IS-IS neighbors out of an interface.	**isis bfd all-neighbors**	INTERFACE

View the established sessions using the command **show bfd neighbors**, as shown in Figure 396.

Figure 396 Viewing Established Sessions for IS-IS Neighbors

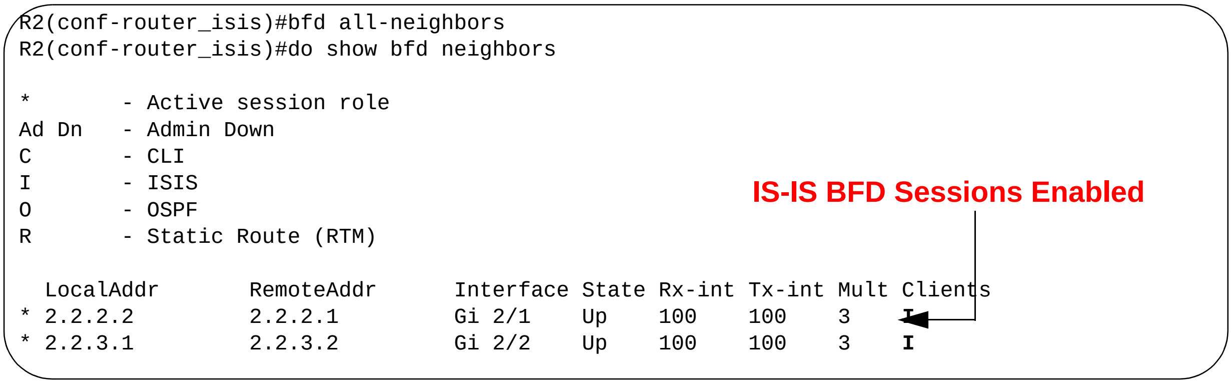

```
R2(conf-router_isis)#bfd all-neighbors
R2(conf-router_isis)#do show bfd neighbors

*       - Active session role
Ad Dn   - Admin Down
C       - CLI
I       - ISIS
O       - OSPF
R       - Static Route (RTM)

  LocalAddr       RemoteAddr      Interface State Rx-int Tx-int Mult Clients
* 2.2.2.2         2.2.2.1         Gi 2/1    Up    100    100    3    I
* 2.2.3.1         2.2.3.2         Gi 2/2    Up    100    100    3    I
```

Changing IS-IS Session Parameters

BFD sessions are configured with default intervals and a default role. The parameters that can be configured are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. These parameters are configured for all IS-IS sessions or all IS-IS sessions out of an interface; if you change a parameter globally, the change affects all IS-IS neighbors sessions. If you change a parameter at interface level, the change affects all IS-IS sessions on that interface.

To change parameters for all IS-IS sessions:

Step	Task	Command Syntax	Command Mode
1	Change parameters for all IS-IS sessions.	**bfd all-neighbors interval** *milliseconds* **min_rx** *milliseconds* **multiplier** *value* **role** [**active** | **passive**]	ROUTER-ISIS

To change parameters for IS-IS sessions on an interface:

Step	Task	Command Syntax	Command Mode
1	Change parameters for all IS-IS sessions out of an interface.	**isis bfd all-neighbors interval** *milliseconds* **min_rx** *milliseconds* **multiplier** *value* **role** [**active** | **passive**]	INTERFACE

View session parameters using the command **show bfd neighbors detail**, as shown in Figure 390 on page 625.

Disabling BFD for IS-IS

If BFD is disabled globally, all sessions are torn down, and sessions on the remote system are placed in a Down state. If BFD is disabled on an interface, sessions on the interface are torn down, and sessions on the remote system are placed in a Down state (Message 10 on page 625). Disabling BFD does not trigger a change in BFD clients; a final Admin Down packet is sent before the session is terminated.

To disable BFD sessions with all IS-IS neighbors:

Step	Task	Command Syntax	Command Mode
1	Disable BFD sessions with all IS-IS neighbors.	**no bfd all-neighbors**	ROUTER-ISIS

To disable BFD sessions with all IS-IS neighbors out of an interface:

Step	Task	Command Syntax	Command Mode
1	Disable BFD sessions with all IS-IS neighbors out of an interface	**isis bfd all-neighbors disable**	INTERFACE

Configuring BFD for VRRP

When using BFD with VRRP, the VRRP protocol registers with the BFD manager on the RPM. BFD sessions are established with all neighboring interfaces participating in VRRP. If a neighboring interface fails, the BFD agent on the line card notifies the BFD manager, which in turn notifies the VRRP protocol that a link state change ocurred.

Configuring BFD for VRRP is a three-step process:

1. Enable BFD globally. See Enabling BFD Globally on page 622.
2. Establish VRRP BFD sessions with all VRRP-participating neighbors.
3. On the master router, establish a VRRP BFD sessions with the backup routers. See page 634.

Related Configuration Tasks

- Change session parameters. See page 636.
- Disable or re-enable BFD on an interface. See page 629.

Establishing Sessions with all VRRP Neighbors

BFD sessions can be established for all VRRP neighbors at once, or a session can be established with a particular neighbor.

Figure 397 Establishing Sessions with VRRP Neighbors

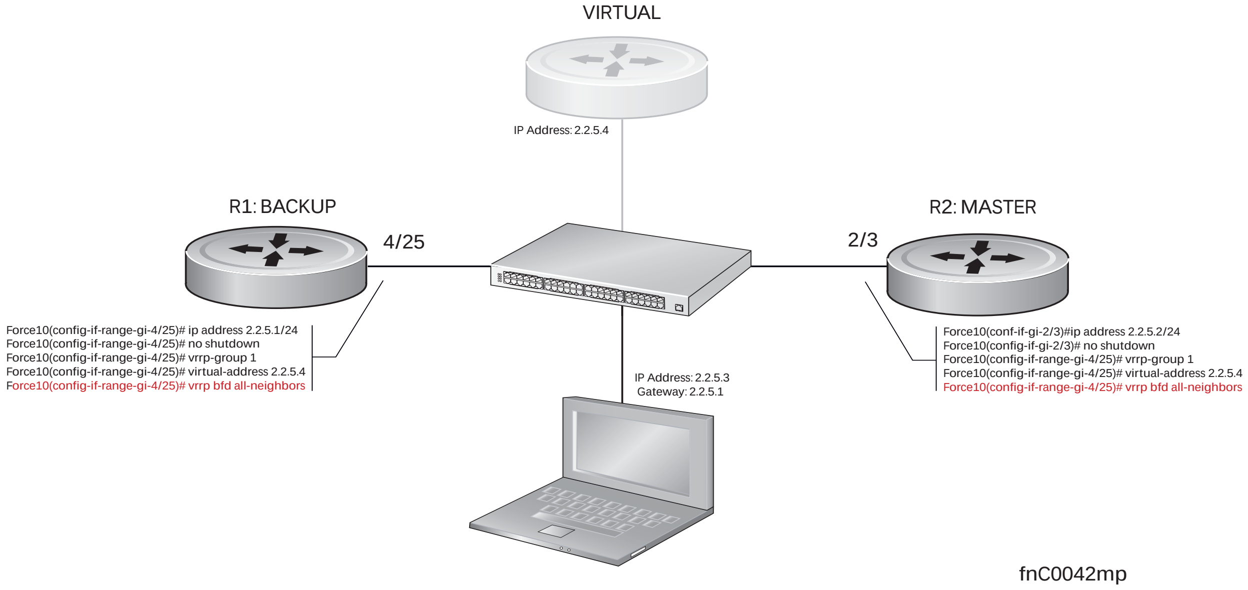

```
Force10(config-if-range-gi-4/25)# ip address 2.2.5.1/24
Force10(config-if-range-gi-4/25)# no shutdown
Force10(config-if-range-gi-4/25)# vrrp-group 1
Force10(config-if-range-gi-4/25)# virtual-address 2.2.5.4
Force10(config-if-range-gi-4/25)# vrrp bfd all-neighbors
```

```
Force10(conf-if-gi-2/3)#ip address 2.2.5.2/24
Force10(config-if-gi-2/3)# no shutdown
Force10(config-if-range-gi-4/25)# vrrp-group 1
Force10(config-if-range-gi-4/25)# virtual-address 2.2.5.4
Force10(config-if-range-gi-4/25)# vrrp bfd all-neighbors
```

To establish sessions with all VRRP neighbors:

Step	Task	Command Syntax	Command Mode
1	Establish sessions with all VRRP neighbors.	**vrrp bfd all-neighbors**	INTERFACE

Establishing VRRP Sessions on VRRP Neighbors

The master router does not care about the state of the backup router, so it does not participate in any VRRP BFD sessions. Therefore, VRRP BFD sessions on the backup router cannot change to the UP state. The master router must be configured to establish an individual VRRP session the backup router.

To establish a session with a particular VRRP neighbor:

Step	Task	Command Syntax	Command Mode
1	Establish a session with a particular VRRP neighbor.	**vrrp bfd neighbor** *ip-address*	INTERFACE

View the established sessions using the command **show bfd neighbors**, as shown in Figure 398.

Figure 398 Viewing Established Sessions for VRRP Neighbors

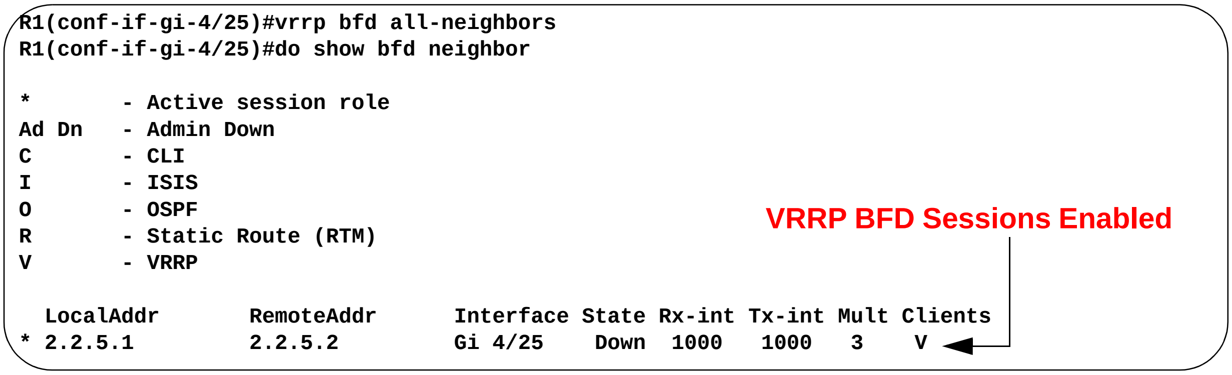

```
R1(conf-if-gi-4/25)#vrrp bfd all-neighbors
R1(conf-if-gi-4/25)#do show bfd neighbor

*       - Active session role
Ad Dn   - Admin Down
C       - CLI
I       - ISIS
O       - OSPF
R       - Static Route (RTM)
V       - VRRP

  LocalAddr       RemoteAddr      Interface State Rx-int Tx-int Mult Clients
* 2.2.5.1         2.2.5.2         Gi 4/25     Down  1000   1000   3    V
```

VRRP BFD Sessions Enabled

Session state information is also shown in the **show vrrp** command output, as shown in Figure 399.

Figure 399 Viewing Established Sessions for VRRP Neighbors

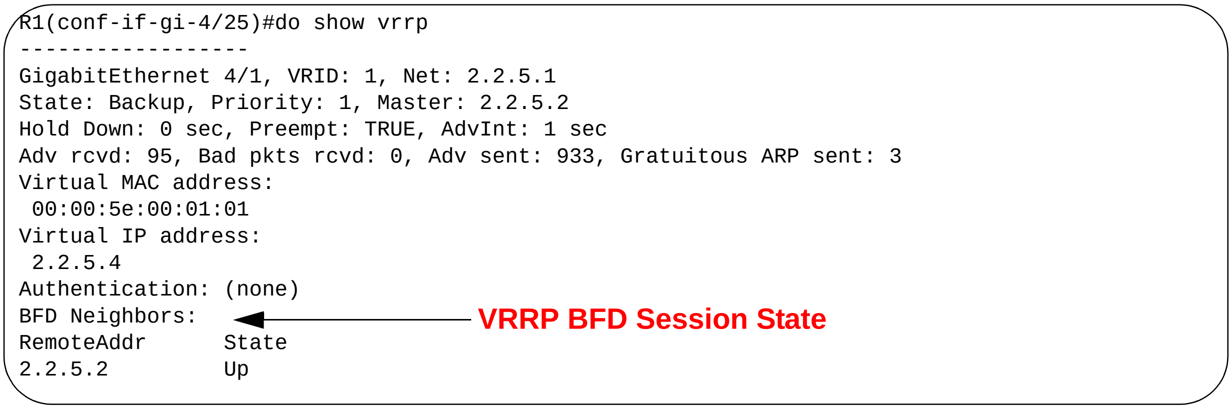

```
R1(conf-if-gi-4/25)#do show vrrp
------------------
GigabitEthernet 4/1, VRID: 1, Net: 2.2.5.1
State: Backup, Priority: 1, Master: 2.2.5.2
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 95, Bad pkts rcvd: 0, Adv sent: 933, Gratuitous ARP sent: 3
Virtual MAC address:
 00:00:5e:00:01:01
Virtual IP address:
 2.2.5.4
Authentication: (none)
BFD Neighbors:
RemoteAddr      State
2.2.5.2         Up
```

VRRP BFD Session State

Changing VRRP Session Parameters

BFD sessions are configured with default intervals and a default role. The parameters that can be configured are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. You can change parameters for all VRRP sessions for a particular neighbor.

To change parameters for all VRRP sessions:

Step	Task	Command Syntax	Command Mode
1	Change parameters for all VRRP sessions.	**vrrp bfd all-neighbors interval** *milliseconds* **min_rx** *milliseconds* **multiplier** *value* **role** [**active** | **passive**]	INTERFACE

To change parameters for a particular VRRP session:

Step	Task	Command Syntax	Command Mode
1	Change parameters for a particular VRRP session.	**vrrp bfd neighbor** *ip-address* **interval** *milliseconds* **min_rx** *milliseconds* **multiplier** *value* **role** [**active** | **passive**]	INTERFACE

View session parameters using the command **show bfd neighbors detail**, as shown in Figure 390 on page 625.

Disabling BFD for VRRP

If any or all VRRP sessions are disabled, the sessions are torn down. A final Admin Down control packet is sent to all neighbors and sessions on the remote system change to the Down state (Message 10 on page 625).

To disable all VRRP sessions on an interface:

Step	Task	Command Syntax	Command Mode
1	Disable all VRRP sessions on an interface.	**no vrrp bfd all-neighbors**	INTERFACE

To disable all VRRP sessions in a particular VRRP group:

Step	Task	Command Syntax	Command Mode
1	Disable all VRRP sessions in a VRRP group.	**bfd disable**	VRRP

To disable a particular VRRP session:

Step	Task	Command Syntax	Command Mode
1	Disable a particular VRRP session on an interface.	**no vrrp bfd neighbor** *ip-address*	INTERFACE

Configuring BFD for VLANs

BFD on Force10 systems is a Layer 3 protocol. Therefore, BFD is used with routed VLANs. BFD on VLANs is analagous to BFD on physical ports. If no routing protocol is enabled, and a remote system fails, the local system does not remove the connected route until the first failed attempt to send a packet. If BFD is enabled, the local system removes the route when it stops receiving periodic control packets from the remote system.

There is one BFD Agent for VLANs and Port-channels, which resides on RP2 as opposed to the other agents which are on the line card. Therefore, the 100 total possible sessions that this agent can maintain is shared for VLANs and Port-channels.

Configuring BFD for VLANs is a two-step process:

1. Enable BFD globally on all participating routers. See Enabling BFD Globally on page 622.
2. Establish sessions with VLAN neighbors. See page 638.

Related Configuration Tasks

- Change session parameters. See page 639.
- Disable BFD for VLANs. See page 629.

Establishing Sessions with VLAN Neighbors

To establish a session, BFD must be enabled at interface level on both ends of the link, as shown in Figure 400. The session parameters do not need to match.

Figure 400 Establishing Sessions with VLAN Neighbors

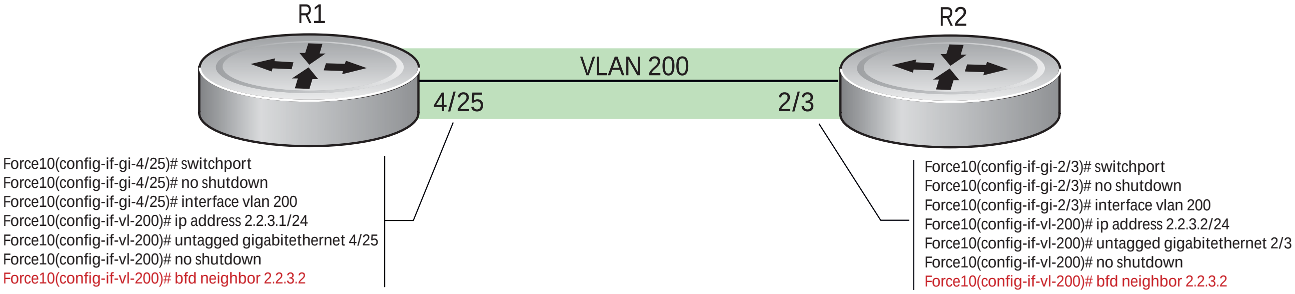

```
Force10(config-if-gi-4/25)# switchport
Force10(config-if-gi-4/25)# no shutdown
Force10(config-if-gi-4/25)# interface vlan 200
Force10(config-if-vl-200)# ip address 2.2.3.1/24
Force10(config-if-vl-200)# untagged gigabitethernet 4/25
Force10(config-if-vl-200)# no shutdown
Force10(config-if-vl-200)# bfd neighbor 2.2.3.2
```

```
Force10(config-if-gi-2/3)# switchport
Force10(config-if-gi-2/3)# no shutdown
Force10(config-if-gi-2/3)# interface vlan 200
Force10(config-if-vl-200)# ip address 2.2.3.2/24
Force10(config-if-vl-200)# untagged gigabitethernet 2/3
Force10(config-if-vl-200)# no shutdown
Force10(config-if-vl-200)# bfd neighbor 2.2.3.2
```

fnC0043mp

To establish a BFD session with a VLAN neighbor:

Step	Task	Command Syntax	Command Mode
1	Establish sessions with a VLAN neighbor.	**bfd neighbor** *ip-address*	INTERFACE VLAN

View the established sessions using the command **show bfd neighbors**, as shown in Figure 401.

Figure 401 Viewing Established Sessions for VLAN Neighbors

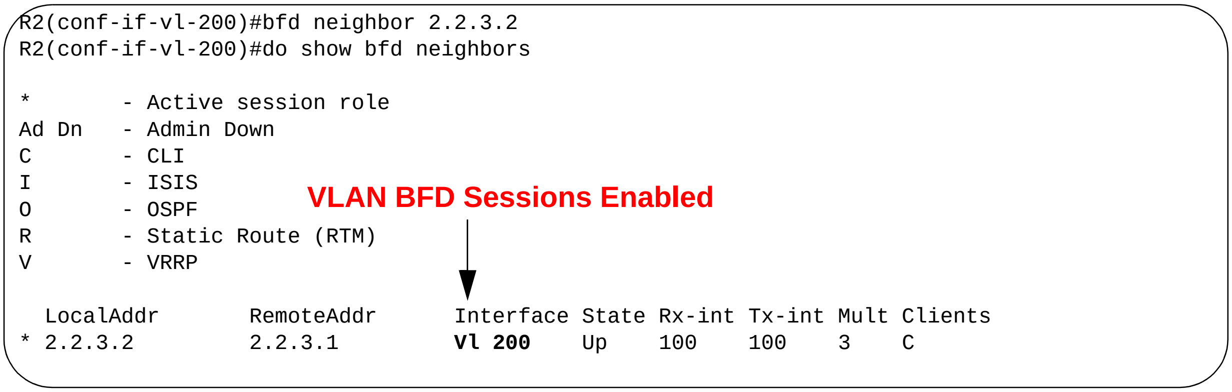

```
R2(conf-if-vl-200)#bfd neighbor 2.2.3.2
R2(conf-if-vl-200)#do show bfd neighbors

*       - Active session role
Ad Dn   - Admin Down
C       - CLI
I       - ISIS
O       - OSPF
R       - Static Route (RTM)
V       - VRRP

  LocalAddr       RemoteAddr      Interface State Rx-int Tx-int Mult Clients
* 2.2.3.2         2.2.3.1         Vl 200    Up    100    100    3    C
```

VLAN BFD Sessions Enabled

Changing Session Parameters

BFD sessions are configured with default intervals and a default role. The parameters that can be configured are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. These parameters are configured per interface; if a configuration change is made, the change affects all sessions on that interface.

To change session parameters on an interface:

Step	Task	Command Syntax	Command Mode
1	Change session parameters for all sessions on an interface.	**bfd interval** *milliseconds* **min_rx** *milliseconds* **multiplier** *value* **role** [**active** | **passive**]	INTERFACE VLAN

View session parameters using the command **show bfd neighbors detail**, as shown in Figure 390 on page 625.

Disabling BFD for VLANs

If BFD is disabled on an interface, sessions on the interface are torn down. A final Admin Down control packet is sent to all neighbors, and sessions on the remote system change to the Down state (Message 10 on page 625).

To disable BFD on a VLAN interface:

Step	Task	Command Syntax	Command Mode
1	Disable all sessions on a VLAN interface.	**no bfd enable**	INTERFACE VLAN

Configuring BFD for Port-Channels

BFD on Port-channels is analogous to BFD on physical ports. If no routing protocol is enabled, and a remote system fails, the local system does not remove the connected route until the first failed attempt to send a packet. If BFD is enabled, the local system removes the route when it stops receiving periodic control packets from the remote system.

There is one BFD Agent for VLANs and Port-channels, which resides on RP2 as opposed to the other agents which are on the line card. Therefore, the 100 total possible sessions that this agent can maintain is shared for VLANs and Port-channels.

Configuring BFD for Port-channels is a two-step process:

1. Enable BFD globally on all participating routers. See Enabling BFD Globally on page 622.
2. Enable BFD at interface level at both ends of the Port-channel. See page 640.

Related Configuration Tasks

- Change session parameters. See page 639.
- Disable BFD a Port-channel. See page 629.

Establishing Sessions on Port-channels

To establish a session, BFD must be enabled at interface level on both ends of the link, as shown in Figure 387. The session parameters do not need to match.

Figure 402 Establishing Sessions on Port-channels

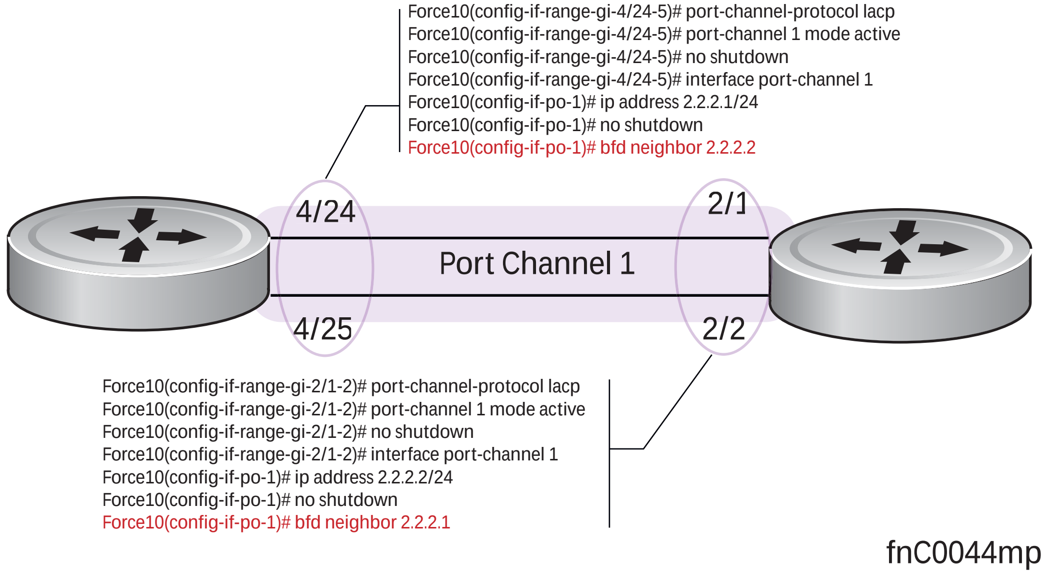
```
Force10(config-if-range-gi-4/24-5)# port-channel-protocol lacp
Force10(config-if-range-gi-4/24-5)# port-channel 1 mode active
Force10(config-if-range-gi-4/24-5)# no shutdown
Force10(config-if-range-gi-4/24-5)# interface port-channel 1
Force10(config-if-po-1)# ip address 2.2.2.1/24
Force10(config-if-po-1)# no shutdown
Force10(config-if-po-1)# bfd neighbor 2.2.2.2
```

```
Force10(config-if-range-gi-2/1-2)# port-channel-protocol lacp
Force10(config-if-range-gi-2/1-2)# port-channel 1 mode active
Force10(config-if-range-gi-2/1-2)# no shutdown
Force10(config-if-range-gi-2/1-2)# interface port-channel 1
Force10(config-if-po-1)# ip address 2.2.2.2/24
Force10(config-if-po-1)# no shutdown
Force10(config-if-po-1)# bfd neighbor 2.2.2.1
```

To establish a session on a Port-channel:

Step	Task	Command Syntax	Command Mode
1	Establish a session on a Port-channel.	**bfd neighbor** *ip-address*	INTERFACE PORT-CHANNEL

View the established sessions using the command **show bfd neighbors**, as shown in Figure 396.

Figure 403 Viewing Established Sessions for VLAN Neighbors

```
R2(conf-if-po-1)#bfd neighbors 2.2.2.1
R2(conf-if-po-1)#do show bfd neighors

*       - Active session role
Ad Dn   - Admin Down
C       - CLI
I       - ISIS
O       - OSPF
R       - Static Route (RTM)
V       - VRRP

  LocalAddr       RemoteAddr       Interface State Rx-int Tx-int Mult Clients
* 2.2.2.2         2.2.2.1          Po 1      Up    100    100    3    C
```

Port-channel BFD Sessions Enabled

Changing Port-Channel Session Parameters

BFD sessions are configured with default intervals and a default role. The parameters that can be configured are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. These parameters are configured per interface; if you change a parameter, the change affects all sessions on that interface.

To change session parameters on an interface:

Step	Task	Command Syntax	Command Mode
1	Change session parameters for all sessions on a Port-channel interface.	**bfd interval** *milliseconds* **min_rx** *milliseconds* **multiplier** *value* **role** [**active** | **passive**]	INTERFACE PORT-CHANNEL

View session parameters using the command **show bfd neighbors detail**, as shown in Figure 390 on page 625.

Disabling BFD for Port-Channels

If BFD is disabled on an interface, sessions on the interface are torn down. A final Admin Down control packet is sent to all neighbors, and sessions on the remote system are placed in a Down state (Message 10 on page 625).

To disable BFD for a Port-channel:

Step	Task	Command Syntax	Command Mode
1	Disable BFD for a port-channel.	**no bfd enable**	INTERFACE PORT-CHANNEL

Configuring Protocol Liveness

Protocol Liveness is a feature that notifies the BFD Manager when a client protocol is disabled. When a client is disabled, all BFD sessions for that protocol are torn down. Neighbors on the remote system receive an Admin Down control packet and are placed in the Down state (Message 10 on page 625).

To enable Protocol Liveness:

Step	Task	Command Syntax	Command Mode
1	Enable Protocol Liveness	**bfd protocol-liveness**	CONFIGURATION

TroubleShooting BFD

Examine control packet field values using the command **debug bfd detail**. Figure 404 shows a three-way handshake using this command.

Figure 404 debug bfd detail Command Output

```
R1(conf-if-gi-4/24)#00:54:38: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE: Changed session state
to Down for neighbor 2.2.2.2 on interface Gi 4/24 (diag: 0)
00:54:38 : Sent packet for session with neighbor 2.2.2.2 on Gi 4/24
TX packet dump:
    Version:1, Diag code:0, State:Down, Poll bit:0, Final bit:0, Demand bit:0
    myDiscrim:4, yourDiscrim:0, minTx:1000000, minRx:1000000, multiplier:3, minEchoRx:0
00:54:38 : Received packet for session with neighbor 2.2.2.2 on Gi 4/24
RX packet dump:
    Version:1, Diag code:0, State:Init, Poll bit:0, Final bit:0, Demand bit:0
    myDiscrim:6, yourDiscrim:4, minTx:1000000, minRx:1000000, multiplier:3, minEchoRx:0
00:54:38: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE: Changed session state to Up for neighbor
2.2.2.2 on interface Gi 4/24 (diag: 0)
```

Examine control packets in hexadecimal format using the command **debug bfd packet**.

Figure 405 debug bfd packet Command Output

```
RX packet dump:
         20 c0 03 18 00 00 00 05 00 00 00 04 00 01 86 a0
         00 01 86 a0 00 00 00 00
00:34:13 : Sent packet for session with neighbor 2.2.2.2 on Gi 4/24
TX packet dump:
         20 c0 03 18 00 00 00 04 00 00 00 05 00 01 86 a0
         00 01 86 a0 00 00 00 00
00:34:14 : Received packet for session with neighbor 2.2.2.2 on Gi 4/24
RX packet dump:
         20 c0 03 18 00 00 00 05 00 00 00 04 00 01 86 a0
         00 01 86 a0 00 00 00 00
00:34:14 : Sent packet for session with neighbor 2.2.2.2 on Gi 4/24
TX packet dump:
```

The output for the command **debug bfd event** is the same as the log messages that appear on the console by default.

Chapter 38 802.1x

Protocol Overview

802.1x is a method of port security. A device connected to a port that is enabled with 802.1x is disallowed from sending or receiving packets on the network until its identity can be verified (through a username and password, for example). This feature is named for its IEEE specification.

802.1x employs Extensible Authentication Protocol (EAP) to transfer a device’s credentials to an authentication server (typically RADIUS) via a mandatory intemediary network access device, in this case, a Force10 switch. The network access device mediates all communication between the end-user device and the authentication server so that the network remains secure. The network access device uses EAP over Ethernet (EAPOL) to communicate with the end-user device and EAP over RADIUS to communicate with the server.

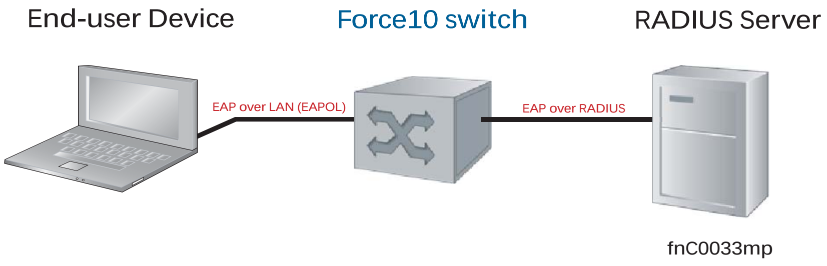

Figure 406 and Figure 407 show how EAP frames are encapsulated in Ethernet and Radius frames.

Figure 406 EAPOL Frame Format

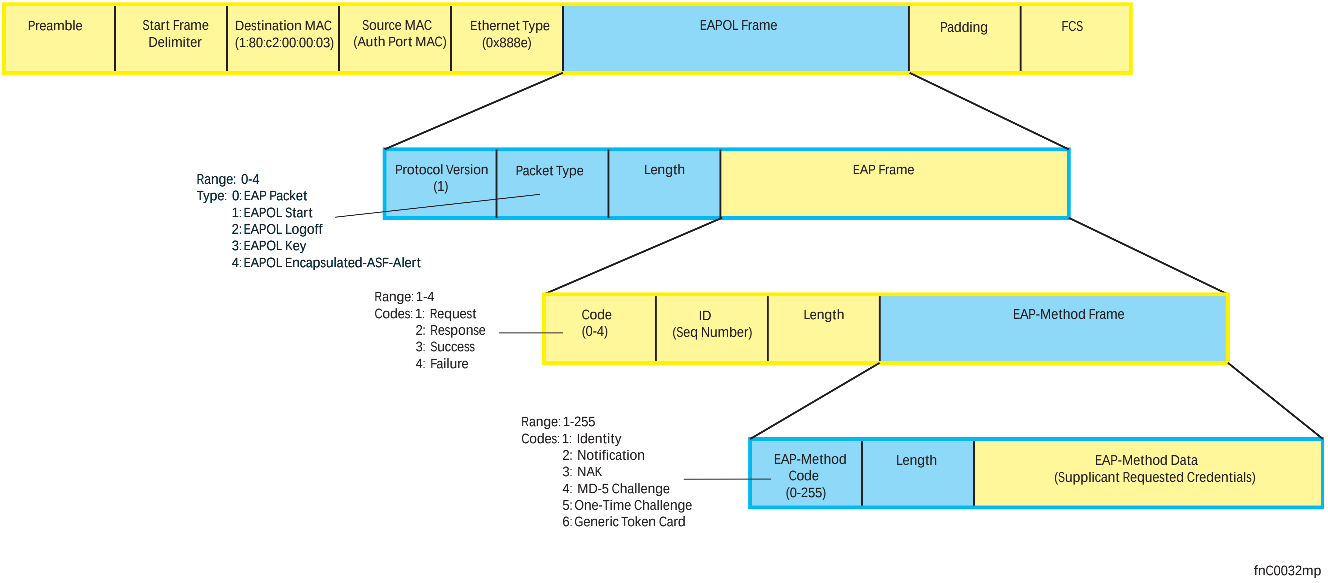

Figure 407 RADIUS Frame Format

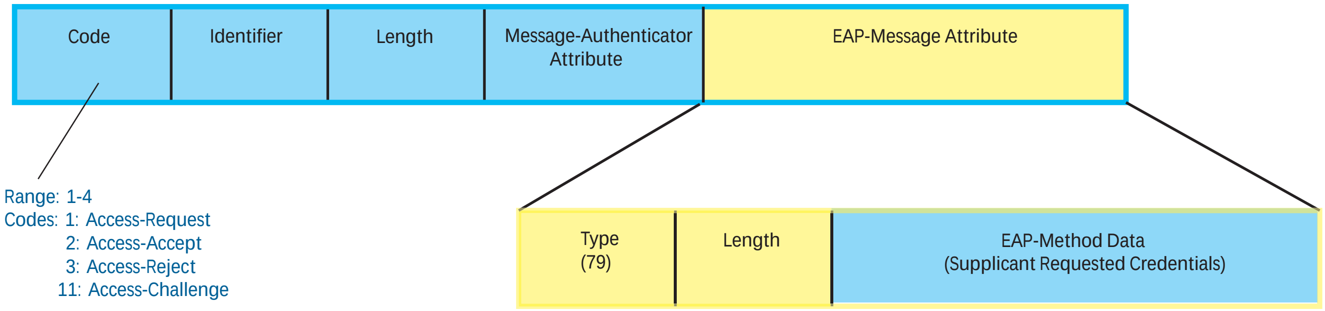

fnC0034mp

The authentication process involves three devices:

- The device attempting to access the network is the **supplicant**. The supplicant is not allowed to communicate on the network until the port is authorized by the authenticator. It can only communicate with the authenticator in response to 802.1x requests.
- The device with which the supplicant communicates is the **authenticator**. The authenicator is the gate keeper of the network. It translates and forwards requests and responses between the authentication server and the supplicant. The authenticator also changes the status of the port based on the results of the authentication process. The Force10 switch is the authenticator.
- The **authentication-server** selects the authentication method, verifies the information provided by the supplicant, and grants it network access privileges.

Ports can be in one of two states:

- Ports are in an **unauthorized** state by default. In this state, non-802.1x traffic cannot be forwarded in or out of the port.
- The authenticator changes the port state to **authorized** if the server can authenticate the supplicant. In this state, network traffic can be forwarded normally.

Note: The Force10 switches place 802.1x-enabled ports in the unathorized state by default.

The Port-authentication Process

The authentication process begins when the authenticator senses that a link status has changed from down to up:

1. When the authenticator senses a link state change, it requests that the supplicant identify itself using an EAP Identity Request Frame.
2. The supplicant responds with its identity in an EAP Response Identity frame.
3. The authenticator decapsulates the EAP Response from the EAPOL frame, encapulates it in a RADIUS Access-Request frame, and forwards the frame to the authentication server.

4. The authentication server replies with an Access-Challenge. The Access-Challenge is request that the supplicant prove that it is who it claims to be, using a specified method (an EAP-Method). The challenge is translated and forwarded to the supplicant by the authenticator.
5. The supplicant can negotiate the authentication method, but if it is acceptable, the supplicant provides the requested challenge information in an EAP Response, which is translated and forwarded to the authentication server as another Access-Request.
6. If the identity information provided by the supplicant is valid, the authentication server sends an Access-Accept frame in which network privileges are specified. The authenticator changes the port state to authorized, and forwards an EAP Success frame. If the identity information is invalid, the server sends and Access-Reject frame. The port state remains unauthorized, and the authenticator forwards EAP Failure frame.

Figure 408 802.1x Authentication Process

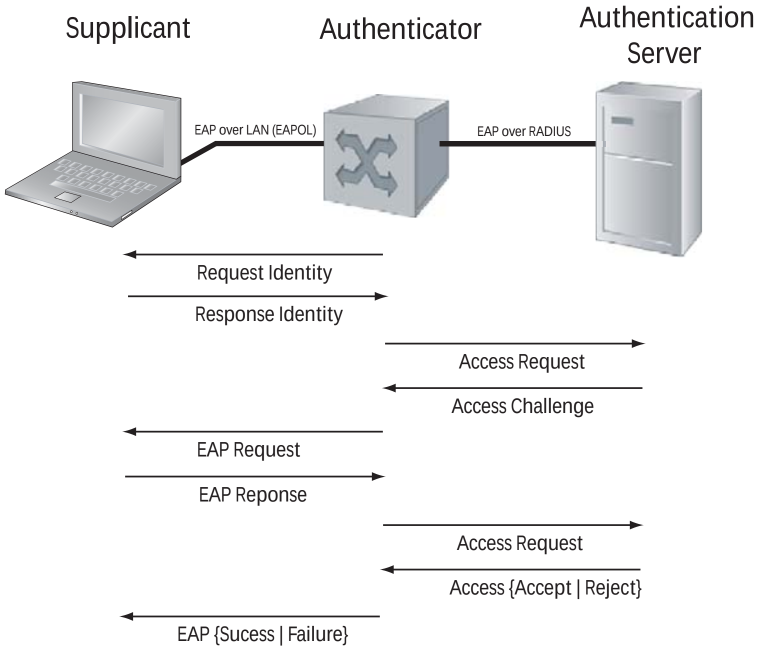

Configuring 802.1x

Configuring 802.1x on a port is a two-step process:

1. Enable 802.1x globally. See page 648.
2. Enable 802.1x on an interface. See page 648.

Related Configuration Tasks

- Configuring Request Identity Re-transmissions on page 650
- Configuring Port-control on page 652
- Re-authenticating a Port on page 653
- Configuring Timeouts on page 654

Important Points to Remember

- E-Series and C-Series support only RADIUS as the authentication server.
- 802.1x is not supported on port-channels or port-channel members.

Enabling 802.1x

802.1x must be enabled globally and at interface level.

Figure 409 Enabling 802.1x

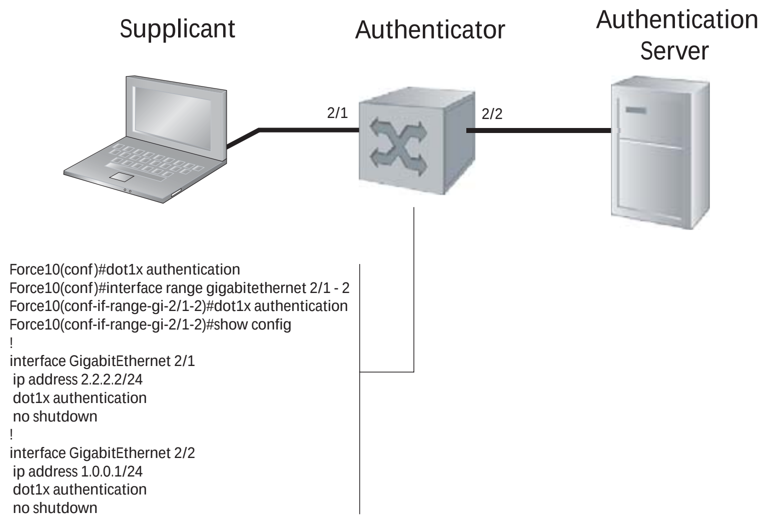

```
Force10(conf)#dot1x authentication
Force10(conf)#interface range gigabitethernet 2/1 - 2
Force10(conf-if-range-gi-2/1-2)#dot1x authentication
Force10(conf-if-range-gi-2/1-2)#show config
!
interface GigabitEthernet 2/1
 ip address 2.2.2.2/24
 dot1x authentication
 no shutdown
!
interface GigabitEthernet 2/2
 ip address 1.0.0.1/24
 dot1x authentication
 no shutdown
```

To enable 802.1x:

Step	Task	Command Syntax	Command Mode
1	Enable 802.1x globally.	**dot1x authentication**	CONFIGURATION
2	Enter INTERFACE mode on an interface or a range of interfaces.	**interface** [**range**]	INTERFACE
3	Enable 802.1x on an interface or a range of interfaces.	**dot1x authentication**	INTERFACE

Verify that 802.1x is enabled globally and at interface level using the command **show running-config | find dot1x** from EXEC Privilege mode, as shown in Figure 410.

Figure 410 Verifying 802.1x Global Configuration

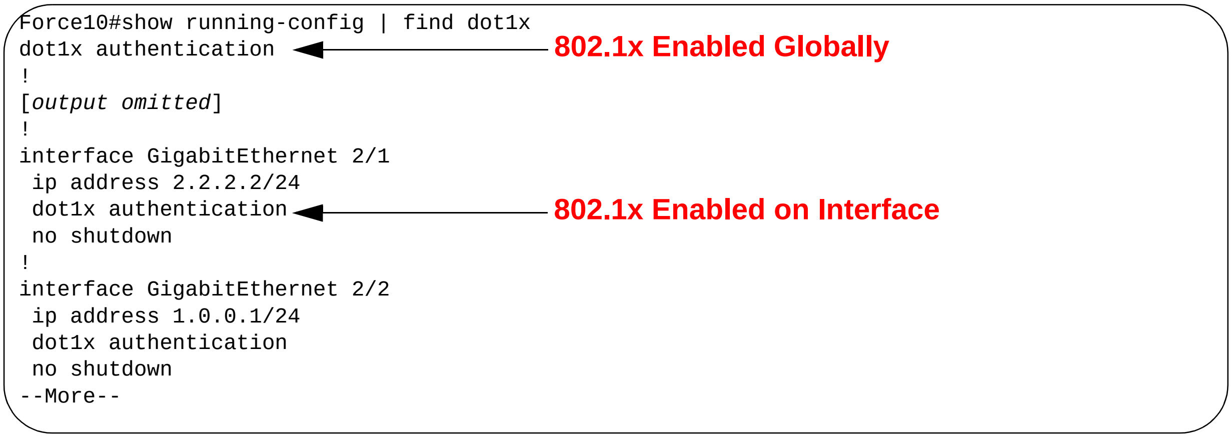

```
Force10#show running-config | find dot1x
dot1x authentication
!
[output omitted]
!
interface GigabitEthernet 2/1
 ip address 2.2.2.2/24
 dot1x authentication
 no shutdown
!
interface GigabitEthernet 2/2
 ip address 1.0.0.1/24
 dot1x authentication
 no shutdown
--More--
```

802.1x Enabled Globally

802.1x Enabled on Interface

View 802.1x configuration information for an interface using the command **show dot1x interface**, as shown in Figure 411.

Figure 411 Verifying 802.1x Interface Configuration

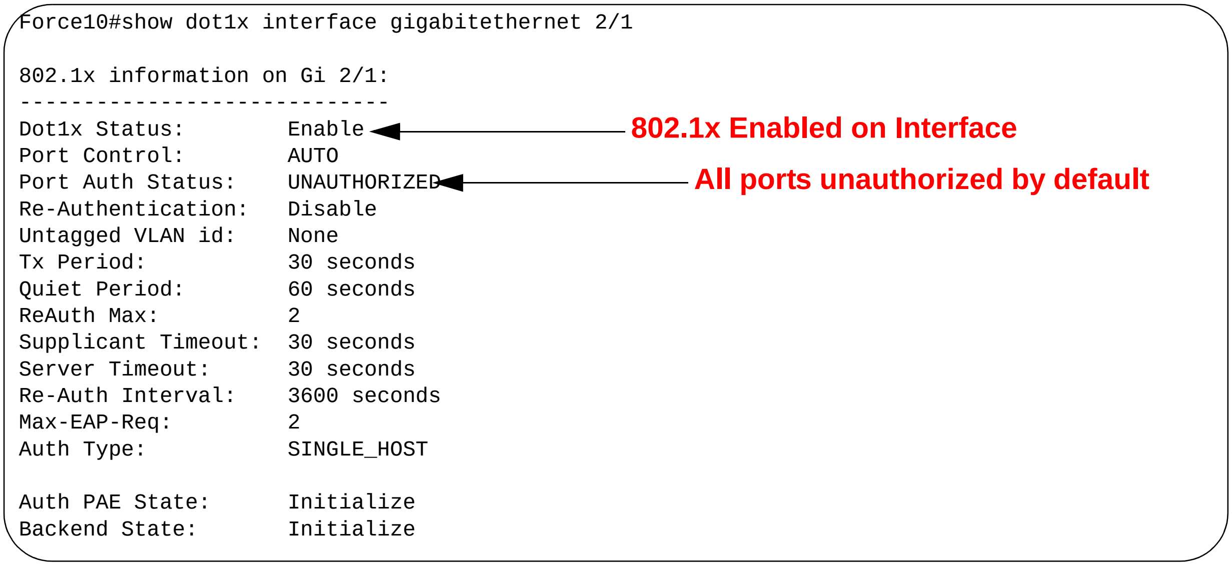

```
Force10#show dot1x interface gigabitethernet 2/1

802.1x information on Gi 2/1:
-----------------------------
Dot1x Status:          Enable
Port Control:          AUTO
Port Auth Status:      UNAUTHORIZED
Re-Authentication:     Disable
Untagged VLAN id:      None
Tx Period:             30 seconds
Quiet Period:          60 seconds
ReAuth Max:            2
Supplicant Timeout:    30 seconds
Server Timeout:        30 seconds
Re-Auth Interval:      3600 seconds
Max-EAP-Req:           2
Auth Type:             SINGLE_HOST

Auth PAE State:        Initialize
Backend State:         Initialize
```

802.1x Enabled on Interface

All ports unauthorized by default

Configuring Request Identity Re-transmissions

If the authenticator sends a Request Identity frame, but the supplicant does not respond, the authenticator waits 30 seconds and then re-transmits the frame. The amount of time that the authenticator waits before re-transmitting and the maximum number of times that the authenticator re-transmits are configurable.

Note: There are several reasons why the supplicant might fail to respond; the supplicant might have been booting when the request arrived, or there might be a physical layer problem.

To configure the amount of time that the authenticator waits before re-transmitting an EAP Request Identity frame:

Step	Task	Command Syntax	Command Mode
1	Configure the amount of time that the authenticator waits before re-transmitting an EAP Request Identity frame.	**dot1x tx-period** *number* Range: 1-31536000 (1 year) Default: 30	INTERFACE

To configure a maximum number of Request Identity re-transmisions:

Step	Task	Command Syntax	Command Mode
1	Configure a maximum number of times that a Request Identity frame can be re-transmitted by the authenticator.	**dot1x max-eap-req** *number* Range: 1-10 Default: 2	INTERFACE

Figure 412 shows configuration information for a port for which the authenticator re-transmits an EAP Request Identity frame after 90 seconds and re-transmits a maximum of 10 times.

Configuring a Quiet Period after a Failed Authentication

If the supplicant fails the authentication process, the authenticator sends another Request Identity frame after 30 seconds by default, but this period can be configured.

Note: The quiet period (**dot1x quiet-period**) is an transmit interval for after a failed authentication where as the Request Identity Re-transmit interval (**dot1x tx-period**) is for an unresponsive supplicant.

To configure the quiet period after a failed authentication:

Step	Task	Command Syntax	Command Mode
1	Configure the amount of time that the authenticator waits to re-transmit a Request Identity frame after a failed authentication.	**dot1x quiet-period** *seconds* Range: 1-65535 Default: 60	INTERFACE

Figure 412 shows configuration information for a port for which the authenticator re-transmits an EAP Request Identity frame:

- after 90 seconds and a maximum of 10 times for an unresponsive supplicant
- Re-transmits an EAP Request Identity frame

Figure 412 Configuring a Request Identity Re-transmissions

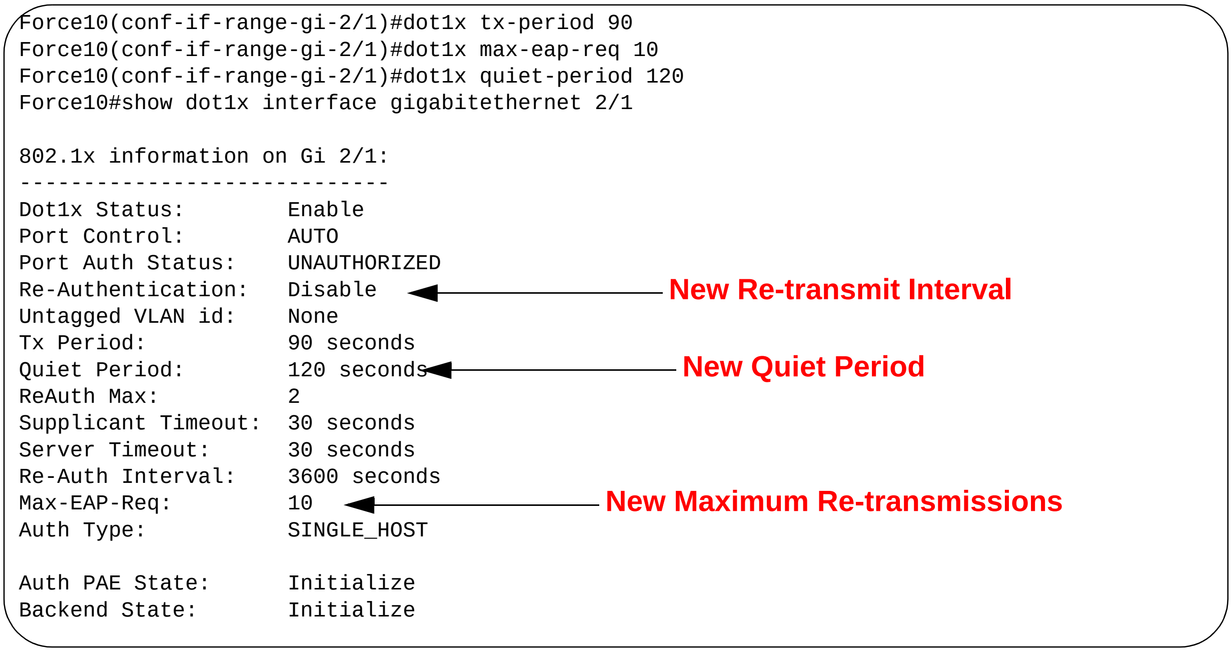

```
Force10(conf-if-range-gi-2/1)#dot1x tx-period 90
Force10(conf-if-range-gi-2/1)#dot1x max-eap-req 10
Force10(conf-if-range-gi-2/1)#dot1x quiet-period 120
Force10#show dot1x interface gigabitethernet 2/1

802.1x information on Gi 2/1:
-----------------------------
Dot1x Status:          Enable
Port Control:          AUTO
Port Auth Status:      UNAUTHORIZED
Re-Authentication:     Disable
Untagged VLAN id:      None
Tx Period:             90 seconds
Quiet Period:          120 seconds
ReAuth Max:            2
Supplicant Timeout:    30 seconds
Server Timeout:        30 seconds
Re-Auth Interval:      3600 seconds
Max-EAP-Req:           10
Auth Type:             SINGLE_HOST

Auth PAE State:        Initialize
Backend State:         Initialize
```

Forcibly Authorizing or Unauthorizing a Port

IEEE 802.1x requires that a port can be manually placed into any of three states:

- **ForceAuthorized** is an authorized state. A device connected to this port in this state is never subjected to the authentication process, but is allowed to communicate on the network. Placing the port in this state is same as disabling 802.1x on the port.

- **ForceUnauthorized** an unauthorized state. A device connected to a port in this state is never subjected to the authentication process and is not allowed to communicate on the network. Placing the port in this state is the same as shutting down the port. Any attempt by the supplicant to initiate authentication is ignored.
- **Auto** is an unauthorized state by default. A device connected to this port is this state is subjected to the authentication process. If the process is successful, the port is authorized and the connected device can communicate on the network. All ports are placed in the **auto** state by default.

To place a port in one of these three states:

Step	Task	Command Syntax	Command Mode
1	Place a port in the ForceAuthorized, ForceUnauthorized, or Auto state.	**dot1x port-control** {**force-authorized** | **force-unauthorized** | **auto**} Default: auto	INTERFACE

Figure 413 shows configuration information for a port that has been force-authorized.

Figure 413 Configuring Port-control

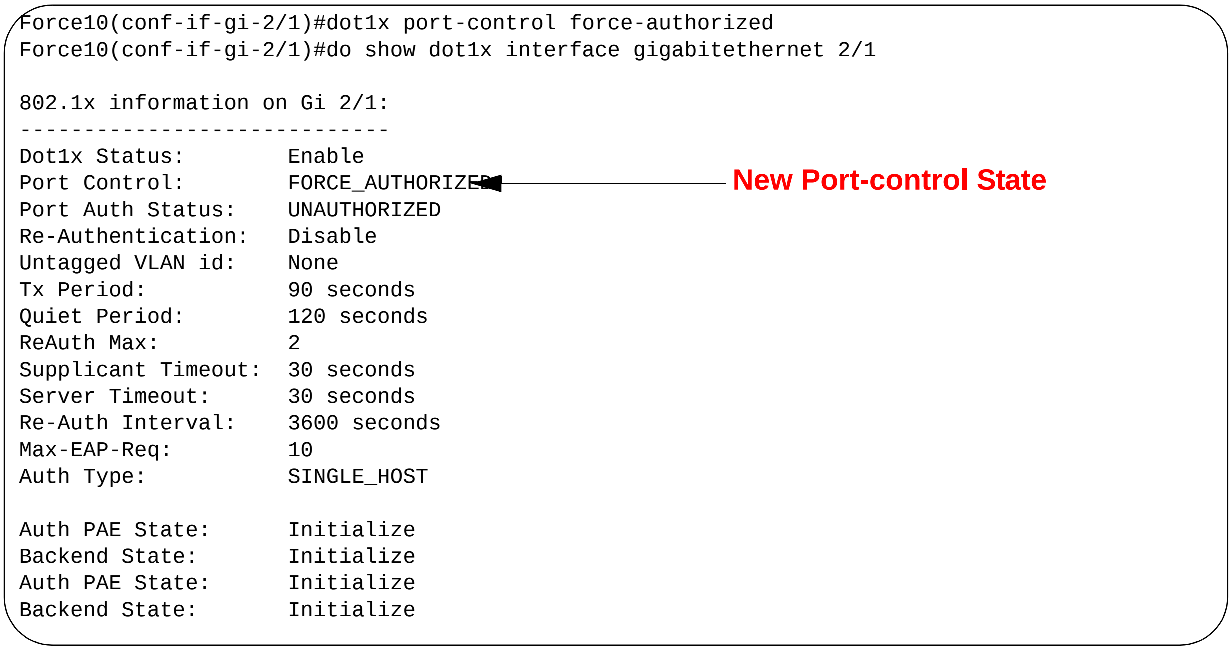

```
Force10(conf-if-gi-2/1)#dot1x port-control force-authorized
Force10(conf-if-gi-2/1)#do show dot1x interface gigabitethernet 2/1

802.1x information on Gi 2/1:
-----------------------------
Dot1x Status:          Enable
Port Control:          FORCE_AUTHORIZED
Port Auth Status:      UNAUTHORIZED
Re-Authentication:     Disable
Untagged VLAN id:      None
Tx Period:             90 seconds
Quiet Period:          120 seconds
ReAuth Max:            2
Supplicant Timeout:    30 seconds
Server Timeout:        30 seconds
Re-Auth Interval:      3600 seconds
Max-EAP-Req:           10
Auth Type:             SINGLE_HOST

Auth PAE State:        Initialize
Backend State:         Initialize
Auth PAE State:        Initialize
Backend State:         Initialize
```

Re-authenticating a Port

Periodic Re-authentication

After the supplicant has been authenticated, and the port has been authorized, the authenticator can be configured to re-authenticates the supplicant periodically. If re-authentication is enabled, the supplicant is required to re-authenticate every 3600 seconds, but this interval can be configured. A maximum number of re-authentications can be configured as well.

To configure a re-authentication or a reauthentication period:

Step	Task	Command Syntax	Command Mode
1	Configure the authenticator to periodically re-authenticate the supplicant.	**dot1x reauthentication** [**interval**] *seconds* Range: 1-65535 Default: 60	INTERFACE

To configure a maximum number of re-authentications:

Step	Task	Command Syntax	Command Mode
1	Configure the maximum number of times that the supplicant can be reauthenticated.	**dot1x reauth-max** *number* Range: 1-10 Default: 2	INTERFACE

Figure 414 Configuring a Reauthentiction Period

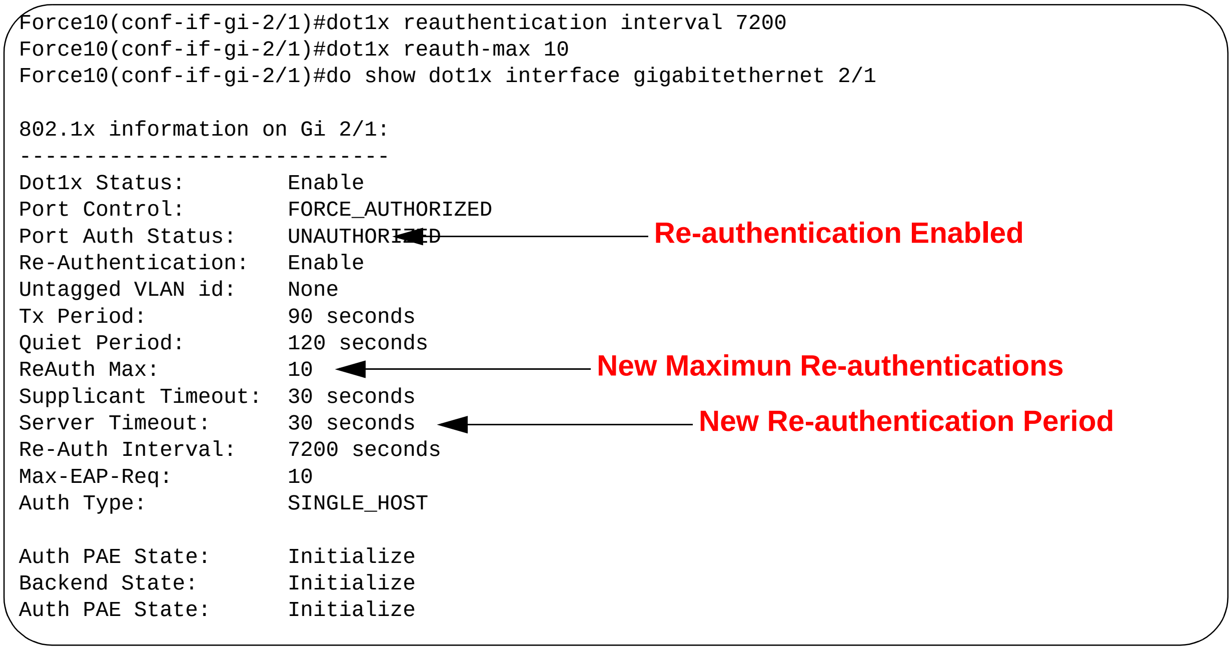

```
Force10(conf-if-gi-2/1)#dot1x reauthentication interval 7200
Force10(conf-if-gi-2/1)#dot1x reauth-max 10
Force10(conf-if-gi-2/1)#do show dot1x interface gigabitethernet 2/1

802.1x information on Gi 2/1:
-----------------------------
Dot1x Status:          Enable
Port Control:          FORCE_AUTHORIZED
Port Auth Status:      UNAUTHORIZED
Re-Authentication:     Enable
Untagged VLAN id:      None
Tx Period:             90 seconds
Quiet Period:          120 seconds
ReAuth Max:            10
Supplicant Timeout:    30 seconds
Server Timeout:        30 seconds
Re-Auth Interval:      7200 seconds
Max-EAP-Req:           10
Auth Type:             SINGLE_HOST

Auth PAE State:        Initialize
Backend State:         Initialize
Auth PAE State:        Initialize
```

Re-authentication Enabled

New Maximun Re-authentications

New Re-authentication Period

Configuring Timeouts

If the supplicant or the authentication server is unresponsive, the authenticator terminates the authentication process after 30 seconds. This amount of time that the authenticator waits for a response can be configured.

To terminate the authentication process due to an unresponsive supplicant:

Step	Task	Command Syntax	Command Mode
1	Terminate the authentication process due to an unresponsive supplicant.	**dot1x supplicant-timeout** *seconds* Range: 1-300 Default: 30	INTERFACE

To terminate the authentication process due to an unresponsive authentication server:

Step	Task	Command Syntax	Command Mode
1	Terminate the authentication process due to an unresponsive authentication server.	**dot1x server-timeout** *seconds* Range: 1-300 Default: 30	INTERFACE

Figure 415 shows configuration information for a port for which the authenticator terminates the authentication process for an unresponsive supplicant or server after 15 seconds.

Figure 415 Configuring a Timeout

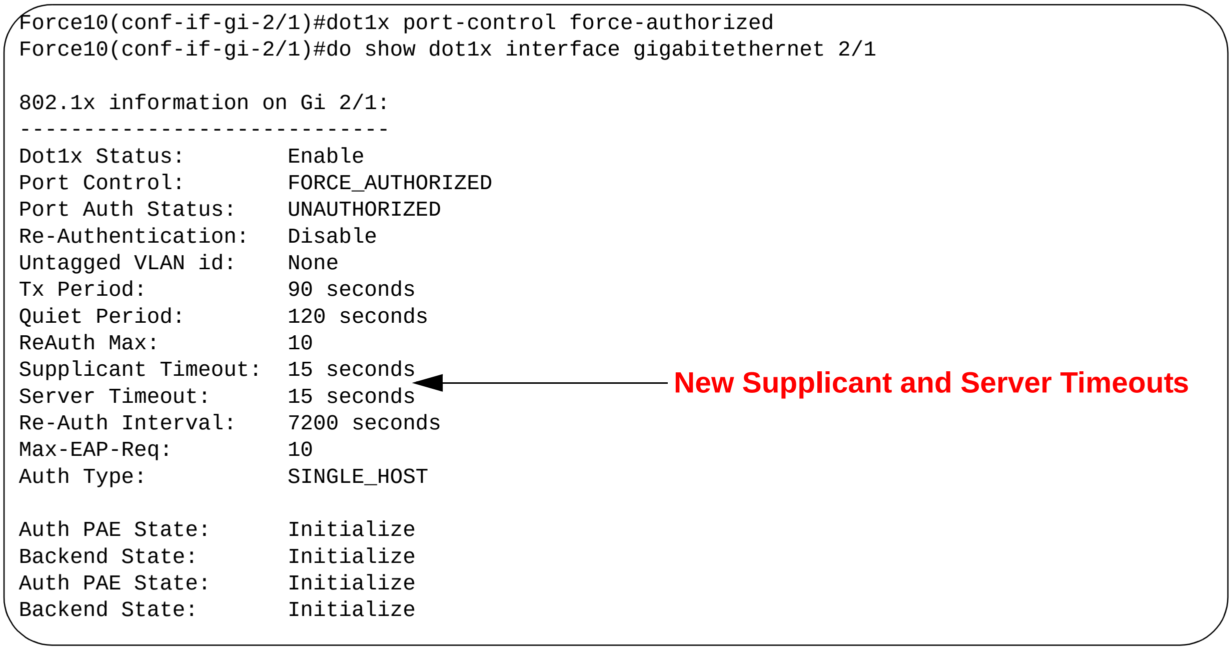

```
Force10(conf-if-gi-2/1)#dot1x port-control force-authorized
Force10(conf-if-gi-2/1)#do show dot1x interface gigabitethernet 2/1

802.1x information on Gi 2/1:
-----------------------------
Dot1x Status:         Enable
Port Control:         FORCE_AUTHORIZED
Port Auth Status:     UNAUTHORIZED
Re-Authentication:    Disable
Untagged VLAN id:     None
Tx Period:            90 seconds
Quiet Period:         120 seconds
ReAuth Max:           10
Supplicant Timeout:   15 seconds
Server Timeout:       15 seconds
Re-Auth Interval:     7200 seconds
Max-EAP-Req:          10
Auth Type:            SINGLE_HOST

Auth PAE State:       Initialize
Backend State:        Initialize
Auth PAE State:       Initialize
Backend State:        Initialize
```

New Supplicant and Server Timeouts

Chapter 39

C-Series Debugging and Diagnostics

C-Series	✓
E-Series	NO

Platform Specific Feature: C-Series Debugging and Diagnostics is supported on C-Series only.

In addition to standard manageability features such as LEDs, SNMP alarms and traps, and Syslogging, the C-Series supports several diagnostic and debugging features that are crucial to isolating and resolving support issues during the operations and maintenance phase.

- Switch Fabric Link Monitoring on page 658
- Runtime Hardware Status Monitoring on page 660
- Bootup Diagnostics on page 662
- Environmental Monitoring on page 662
- Trace Logs on page 664
- Advanced Debugging Commands on page 667
- Offline Diagnostics on page 672

Switch Fabric Overview

The switch fabric is formed through the installed RPMs and line cards via C-Series Switch Fabric (CSF) ASICs.

Each RPM includes four CSFs, each of which provides eight Backplane Data (BDP) links, one link for each line card slot. In total, an RPM provides 32 BDP links of forwarding capacity.

Each line card includes two CSFs. Six of the eight links on the CSFs are used as follows:

- Up to four ports—ports 1 to 4—connect to the Forwarding Processors (FP). These ports are referred to as the Internal Dataplane (IDP) link.
- Ports 5-8 connect to the RPMs. These ports are referred to as the BDP links.

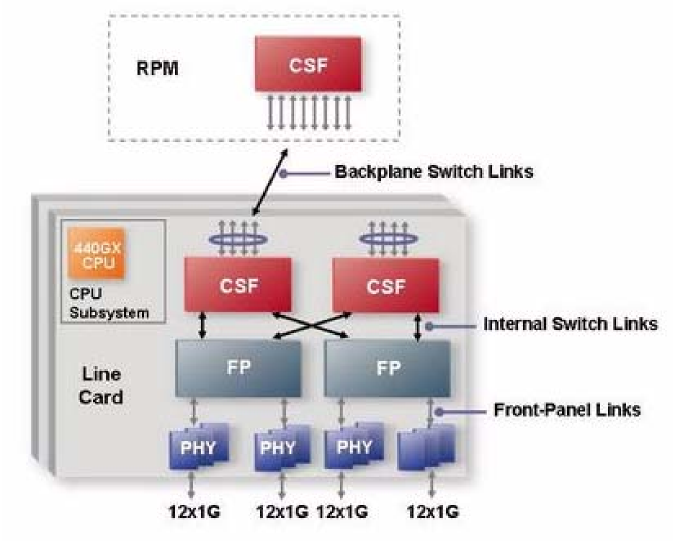

Figure 416 Architecture Diagram of the 1x48GE Line Card

The number of FPs varies with the line card type, as shown in Table 51.

Table 51 FPs, CSFs, and IDP Links by Line Card Type

Line Card Type	# of FPs	# of CSFs	IDP Links Used on CSF
48x1GE	2	2	• Ports 1 and 2 connect to the FPs. • Ports 3 and 4 are unused.
96x1GE	4	2	Ports 1-4 connect to the FPs.
4x10GE	2	2	Ports 1-4 connect to the FPs.

Switch Fabric Link Monitoring

FTOS Switch Manager (SWMGR) task monitors the BDP links on the RPM. This task also monitors the overall state of the switch fabric and reports any changes via Syslog messages.

FTOS Switch Agent (SWAGT) monitors the IDP and BDP links on the line cards.

FTOS Link Monitoring task continually polls the status of the IDP and BDP links. If it finds an open link, the system brings down the link and reports the condition via a message similar to the one shown in Table 52.

Table 52 FTOS Link Monitoring Syslog Message Example

```
Mar 12 21:01:18: %RPM1-P:CP %SWMGR-1-BDP_LINK_DETECT: Backplane datapath link status for RSM0
Switch fabric unit# 0 port# 0 => DOWN
!- Describes only the state of the port.
Mar 10 16:58:28: %RPM1-P:CP %SWMGR-1-IDP_LINK_DETECT: Internal datapath link status for Linecard#5
Switch unit# 1 port# 24 and Switch fabric unit# 3 port# 1=> DOWN
!- Describes the state of the link.
```

Note: These messages are not reported when a line card is reset by a user command.

If a backplane link on a line card goes down, the RPM side of the link stays up to avoid duplicate reporting.

Bringing down an IDP or BDP link causes the card to be powered-off and placed into a "card problem - port pipe problem" state. Use the **show linecard** command to view the status of a line card.

If a single BDP link to the active RPM is down, the line card will be placed in an error state. Use the **show switch links** command to view the status of the dataplane links, as shown in Figure 417.

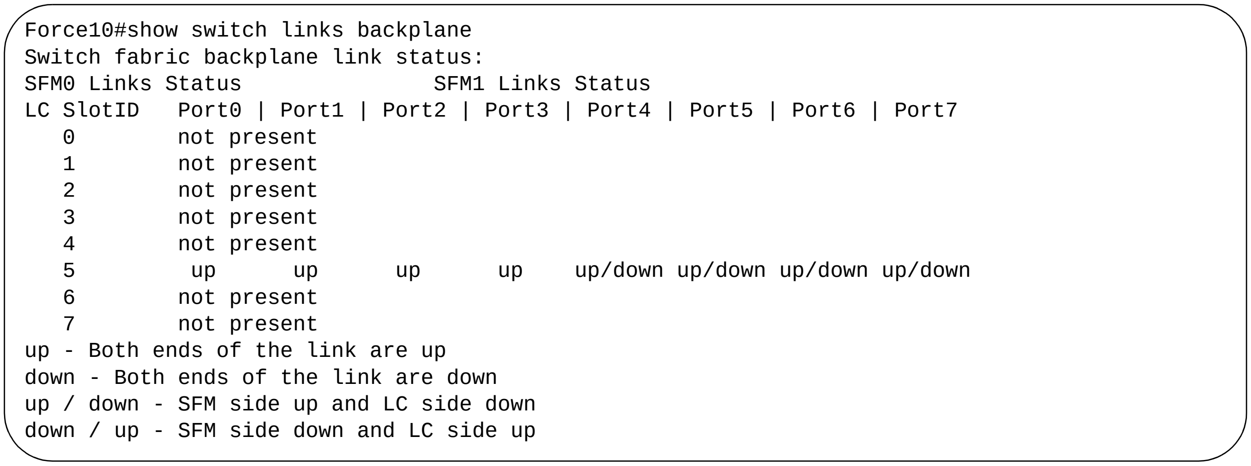

```
Force10#show switch links backplane
Switch fabric backplane link status:
SFM0 Links Status               SFM1 Links Status
LC SlotID   Port0 | Port1 | Port2 | Port3 | Port4 | Port5 | Port6 | Port7
   0        not present
   1        not present
   2        not present
   3        not present
   4        not present
   5         up      up      up      up    up/down up/down up/down up/down
   6        not present
   7        not present
up - Both ends of the link are up
down - Both ends of the link are down
up / down - SFM side up and LC side down
down / up - SFM side down and LC side up
```

Figure 417 show switch links backplane Command Example

To monitor the status of a virtual SFM, use the **show sfm** command shown in Figure 418. The system reports an "active" status if all CSF ASICs on the RPM initialize successfully, whether or not any line cards are installed and the BDP links are up.

```
Force10#show sfm
Switch Fabric State:  up
-- SFM 0 --
Status        : active
Module Type   : SFM - Switch Fabric Module
Up Time       : 1 day, 6 hr, 0 min
-- SFM 1 --
Status        : not present
```

Figure 418 show sfm Command Example

Use the FTOS Syslogging feature to monitor the overall status of the switch fabric. Changes in switch fabric status are reported via messages similar those in Table 53.

Table 53 Switch Fabric Status Change Syslog Message Example

00:00:13: %RPM1-P:CP %TSM-6-SFM_FULL_PARTIAL_STATE: SW_FAB_UP_1 SFM in the system 00:00:13: %RPM1-P:CP %TSM-6-SFM_SWITCHFAB_STATE: Switch Fabric: UP

Runtime Hardware Status Monitoring

The FTOS Poll Manager (POLLMGR) process reads the key status registers on hardware sub-components to pro-actively identify and report a hardware fault. An example Syslog message is shown in Table 54.

Table 54 Poll Manager Syslog Message Example

%RPM1-P:CP %POLLMGR-2-BPL_IRC_ERR: Back Plane Link Error

The Poll Manager runs automatically in the background and cannot be disabled. The possible Poll Manager Syslog messages are given in Table 55.

Table 55 Poll Manager Syslog Message Description

Message	Description
POLLMGR-2-ALT_RPM_STATE	Reports that either the standby RPM is not present or has been detected.
POLLMGR-2-USER_FLASH_DETECT	Reports the status of the flash disks. If reported during boot up, this message indicates either: a External flash disk missing in slot0: b Internal flash disk missing in flash: If reported during runtime, this message indicates either: a External flash disk removed from slot0: b Internal flash disk removed from flash:

Table 55 Poll Manager Syslog Message Description

Message	Description
POLLMGR-2-POLLMGR_RPM_ECC_ERR_DETECT	Indicates that the system detected a single-bit ECC memory error in the RPM CPU memory (SDRAM). The system tracks the number of multi-bit errors and resets the system after a certain number of such errors are recorded. Upon reset, the system writes a failure trace file to the TRACE_LOG directory for analysis by Force10 Networks.
POLLMGR-2-POLLMGR_BPL_IRC_ERR	Indicates that the system detected an error on the internal IPC switch subsystem connection between the two RPMs. This connection is referred to as Inter-RPM Communication (IRC). When a number of consecutive IRC heartbeat messages are lost, the system will declare an IRC timeout via a Syslog message and reset the system. This message suggests that a hardware fault on the RPM may have caused the IRC timeout. To troubleshoot this issue: • Verify that the RPMs are fully inserted. • Try a swap test of the RPMs. • Capture the output of the following show hardware commands: •**show hardware rpm** *number* **cpu party-bus statistics** •**show hardware rpm** *number* **mac counters** •**show hardware rpm** *number* **mac port-statistics rpm** *number* (of alternate RPM)
POLLMGR-2-POLLMGR_PTYBUS_LINK_SCAN	Indicates the internal IPC party bus connection to a line card has changed to down. IPC, or inter-process communication, is the protocol used among the RPM and line card CPUs to exchange information. The underlying IPC subsystem uses internal Ethernet links. To troubleshoot this condition: • Capture the output of the **show hardware linecard cpu party-bus statistics** command, and forward it to Force10 Networks.

Figure 419 illustrates the IPC subsystem, including the IRC links between the RPMs, and the relevant troubleshooting commands.

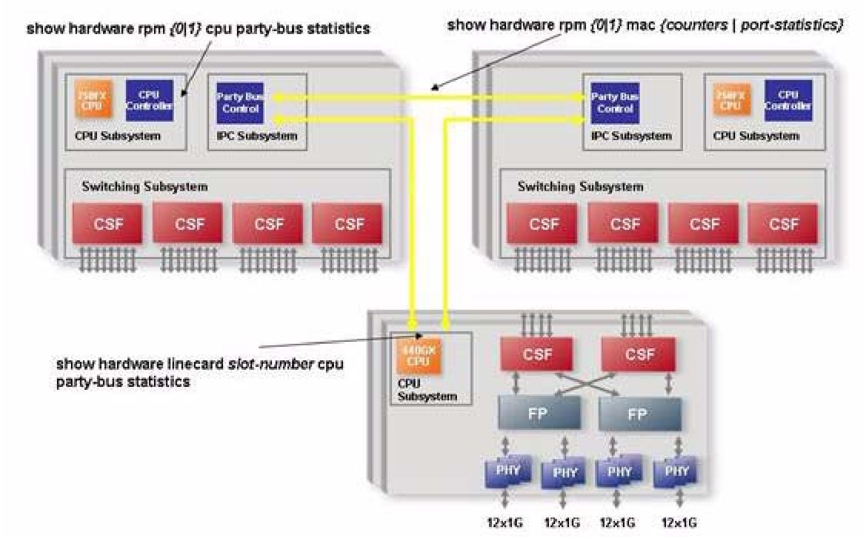

Figure 419 IPC Sub-system

Bootup Diagnostics

During bootup and reset of a card, diagnostics check the status of key hardware sub-components and verify that all ASICs initialize successfully.

Recognizing Bootup Failures

Any detected failures or errors during bootup are reported via Syslog. The messages in Table 56 and Table 57 might be reported for line card failures and RPMs, respectively.

Table 56 Line Card Boot Up Failure Syslog Messages

%CHAGT-5-LINK_STATUS_DOWN: Link status bad for port pipe [number] on line card [number]
%CHAGT-5-PORT PIPE DOWN: Port pipe [number] down or errored on line card [number]

Table 57 RPM Boot Up Failure Syslog Messages

%CHMGR-2-RPM_ISOLATED: Active RPM is unable to talk to line cards
%CHMGR-3-RAM_STANDBY_RPM_FAULT: Secondary RPM fault detected
%CHMGR-3-RPM_POST_PORTPIPE_FAIL: RPM port pipe fails on boot up test
%CHMGR-3-RPM_DRIVER_OPEN_FAIL: RPM driver open fails on boot up
%CHMGR-3-RPM_SWITCH_OPEN_FAIL: RPM switch driver fails on boot up
%CHMGR-3-RPM_POST_RTC_FAIL: RPM RTC fails on boot up test

Troubleshooting Bootup Failures

If these messages are seen, collect the output of the **show console lp** and **show tech** commands and contact the Force10 Networks Technical Assistance Center.

Environmental Monitoring

All C-Series components use environmental monitoring hardware to detect overtemperature, undervoltage, and overvoltage conditions. Use the **show environment** command to monitor the components for any major or minor alarm conditions. The output in Figure 420 displays the environment status of the RPM.

```
Force10#show environment rpm

--  RPM Environment Status  --
Slot  Status        Temp   Voltage
--------------------------------------
  0   active        33C    ok
  1   not present
```

Figure 420 show environment rpm Command Example

Recognizing an Overtemperature Condition

An overtemperature condition occurs, for one of two reasons:

- The card genuinely is too hot.
- A sensor has malfunctioned.

Inspect cards adjacent to the one reporting the condition to discover the cause.

- If directly adjacent cards are not normal temperature, suspect a genuine overheating condition.
- If directly adjacent cards are normal temperature, suspect a faulty sensor.

When the system detects a genuine over-temperature condition, it powers off the card. To recognize this condition, look for the system messages in Table 58.

Table 58 Over Temperature Condition System Messages

CHMGR-2-MAJOR_TEMP: Major alarm: chassis temperature high (temperature reaches or exceeds threshold of [value]C) CHMGR-2-TEMP_SHUTDOWN_WARN: WARNING! temperature is [value]C; approaching shutdown threshold of [value]C

To view the programmed alarm thresholds levels, including the shutdown value, execute the **show alarms threshold** command shown in Figure 421.

```
Force10#show alarms threshold

--  Temperature Limits (deg C)  --
---------------------------------------------------------------
          Minor     Minor Off   Major    Major Off   Shutdown
Linecard   75        70         80         77         85
RPM        65        60         75         70         80
Force10#
```

Figure 421 show alarms threshold Command Example

Troubleshooting an Overtemperature Condition

To troubleshoot an over-temperature condition:

1. Use the **show environment** commands to monitor the temperature levels.
2. Check air flow through the system. On the C-Series, air flows sideways from right to left. Ensure the air ducts are clean and that all fans are working correctly.
3. Once the software has determined that the temperature levels are within normal limits, the card can be re-powered safely. Use the **power-on** command in EXEC mode to bring the line card back online.

In addition, Force10 requires that you install blanks in all slots without a line card to control airflow for adequate system cooling.

Note: Exercise care when removing a card; if it has exceeded the major or shutdown thresholds, the card could be hot to the touch.

Recognizing an Under-voltage Condition

If the system detects an under-voltage condition and declares an alarm. To recognize this condition, look for the system messages in Table 59.

Table 59 Under-voltage Condition System Messages

%CHMGR-1-CARD_SHUTDOWN: Major alarm: Line card 2 down - auto-shutdown due to under voltage

This message in Table 59 indicates that the specified card is not receiving enough power. In response, the system first shuts down Power over Ethernet (PoE). If the under-voltage condition persists, line cards are shut down, then RPMs.

Troubleshooting an Under-voltage Condition

To troubleshoot an under-voltage condition, check that the correct number of power supplies are installed and their Status LEDs are lit.

Trace Logs

In addition to the syslog buffer, FTOS buffers trace messages which are continuously written by various FTOS software tasks to report hardware and software events and status information. Each trace message provides the date, time, and name of the FTOS process. All messages are stored in a ring buffer and can be saved to a file either manually or automatically upon failover.

- There are three trace buffers for CP: software, hardware, and command-history.
- There are two trace buffers for LP: software and hardware.

Buffer Full Condition

When the trace ring buffer is full, the trace lines are saved as a file into the flash (e.g. hw_trace_RPM0CP.0). When the buffer fills for the second time, it is saved as hw_trace_RPM0CP.1, and so on until hw_trace_RPM0CP.4. From the sixth time onwards, when the trace buffer fills, the second trace file (hw_trace_RPM0CP.1) is overwritten.

Trace file hw_trace_RPM0CP.0 is not overwritten so that chassis bootup message are preserved.

FTOS uses a similar approach to saving the various trace messages for CP and all LPs.

The CP and LP trace file names are:

Note: Line card 1 is taken as the example for the following filenames.

- **CP [SW trace]** : sw_trace_RPM0CP.0, sw_trace_RPM0CP.1, sw_trace_RPM0CP.2, sw_trace_RPM0CP.3 and sw_trace_RPM0CP.4
- **CP [HW trace]** : hw_trace_RPM0CP.0, hw_trace_RPM0CP.1, hw_trace_RPM0CP.2, hw_trace_RPM0CP.3 and hw_trace_RPM0CP.4
- **LP [SW trace]** : sw_trace_LPX.0, sw_trace_LP1.1, sw_trace_LP1.2, sw_trace_LP1.3 and sw_trace_LP1.4
- **LP [HW trace]** : hw_trace_LPX.0, hw_trace_LP1.1, hw_trace_LP1.2, hw_trace_LP1.3 and hw_trace_LP1.4

Trace files are saved in the directory flash:/TRACE_LOG_DIR/TRACE_CURR_BOOT. Upon a system reload this directory is renamed flash:/TRACE_LOG_DIR/TRACE_LAST_BOOT, and a fresh empty flash:/TRACE_LOG_DIR/TRACE_CURR_BOOT directory is created.

Manual Reload Condition

When the chassis is reloaded manually (through the CLI), trace messages in all of the buffers (software and hardware) in CP and linecards are saved to the flash as reload_traceRPM0_CP and reload_traceLP1 in flash:/TRACE_LOG_DIR/TRACE_CURR_BOOT. After reload, you can see these files in flash:/TRACE_LOG_DIR/TRACE_LAST_BOOT..

When the trace messages are being saved on reload, Message 11 is displayed.

Message 11 Saving Trace Messages

```
Starting to save trace messages… Done.
```

The CP and LP trace file names at chassis reload are:

- **CP:** reload_traceRPM0_CP
- **LP:** reload_traceLP1

CP/RP1/RP2 Software Exceptions

When a RPM resets due to an RP1 or RP2 software exception, the linecard trace files are saved to flash:/TRACE_LOG_DIR directory.

The CP and LP trace file names in the case of a software exception are:

- **CP:** failure_trace_RPM1_CP
- **LP:** failure_trace_RPM1_LP1

With a dual RPM system, linecard trace logs are saved in the case of a CP, RP1 or RP2 crash. Linecard trace logs are saved in the new primary RPM. Here the name of the line card trace file helps in identifying a failed RPM. That is, if RPM0 fails, then the trace files are saved in RPM1 with filename failure_trace_RPM0_LP1. This is because the failover happens before the linecard traces are saved.

Viewing Trace Buffer Content

The command-history trace feature captures all commands entered by all users of the system with a time stamp and writes these messages to a dedicated trace log buffer. The system generates a trace message for each executed command. No password information is saved to the file.

To view the command-history trace, use the **show command-history** command, as shown in Figure 422.

```
C300-TAC-B6#show command
[12/5 10:57:8]: CMD-(CLI):service password-encryption
[12/5 10:57:12]: CMD-(CLI):hostname Force10
[12/5 10:57:12]: CMD-(CLI):ip telnet server enable
[12/5 10:57:12]: CMD-(CLI):line console 0
[12/5 10:57:12]: CMD-(CLI):line vty 0 9
[12/5 10:57:13]: CMD-(CLI):boot system rpm0 primary flash://FTOS-CB-1.1.1.2E2.bin
```

Figure 422 show command-history Command Example

Writing the Contents of the Trace Buffer

The trace logs are saved to automatically but you can save the contents of a buffer manually via the CLI.

To manually write the contents of a trace buffer on CP to a file on the flash:

Step	Task	Command Syntax	Command Mode
1	Write the buffered trace log to flash.	**upload trace-log cp** [**cmd-history** | **hw-trace** | **sw-trace**]	EXEC privilege

To manually write the contents of a trace buffer on LP to a file on the flash:

Step	Task	Command Syntax	Command Mode
1	Write the buffered trace log to flash.	**upload trace-log** [**rp1** | **rp2** | **linecard**] *number* [**hw-trace** | **sw-trace**]	EXEC privilege

Advanced Debugging Commands

The C-Series supports advanced debugging commands for isolating suspected software and hardware support issues.

debug Commands

The debug command tree provides packet- and event-level debugging for major protocols, as shown in Figure 423.

```
Force10#debug ?
aaa                      AAA debug information
arp                      IP ARP debug information
cpu-traffic-stats        Start collecting CPU traffic statistics
dot1x                    Dot1x debug information
ftpserver                FTP Server debug information
ifm                      IFM Debug information
ip                       IP debug information
ntp                      NTP debug information
ppp                      PPP debug commands
radius                   RADIUS debug information
spanning-tree            Spanning tree debug information
tacacs+                  TACACS+ debug information
```

Figure 423 debug Command Tree

show hardware Commands

The **show hardware** command tree consists of EXEC privilege commands that have been created or changed specially for use with the C-Series. These commands display information from a hardware sub-component, such as an FP or CSF ASIC, and from hardware-based feature tables.

Table 60 lists the **show hardware** commands available as of the latest FTOS version.

Note: show hardware commands should only be used under the guidance of Force10 Networks Technical Assistance Center.

Table 60 show hardware Commands

Command	Description
show hardware interface phy	View link status information, including the transmitted and received auto-negotiation control words. Use the **registers** keyword to capture a dump of key PHY registers for your technical support representative.
show hardware drops	View internal packet-drop counters on a line card or RPM

Table 60 show hardware Commands

Command	Description
show hardware rpm mac counters **clear hardware rpm mac counters**	Enter the keyword **counters** keyword to view or clear the receive and transmit frame counters for the party bus switch in the IPC subsystem on the RPM.
show hardware rpm mac port-statistics **clear hardware rpm mac port-statistics**	Enter the keyword **port-statistics** to view or clear detailed Ethernet statistics for the specified port on the party bus switch.
show hardware rpm cpu management	View internal interface status information for the RPM CPU port which connects to the external management interface.
show hardware cpu party-bus **clear hardware cpu party-bus**	View or clear statistics for the party-bus port on the CPU of the specified line card or RPM.
show hardware cpu data-plane	View driver-level statistics for the data-plane port on the CPU for the specified line card or RPM.
show hardware unit	Views advanced counters, statistics, and register information for the FP and CSF ASICs.

Monitoring CPU and Memory Usage

Table 61 lists the commands to that display CPU and memory utilization.

Table 61 CPU and Memory Utilization show Commands

Command	Description
show process cpu cp	Lists all of the running processes on the RPM and CPU usage for the last 5 seconds, 1 minute, and 5 minutes. Use the number parameter to specify the number of process with the highest CPU usage to display.
show process cpu lp	Lists all of the running processes on the specified line card and CPU usage for the last 5 seconds, 1 minute and 5 minutes. Use the number parameter to specify the number of process with the highest CPU usage to display.
show process memory cp	Lists detailed memory statistics for all processes running on the RPM.
show process memory lp	Lists detailed memory statistics for all processes running on the line card in the specified slot.

Figure 424 displays memory utilization for Line Card 1 using the **show processes memory lp** command.

```
Force10#show processes memory lp 1
Total:   214684976, MaxUsed:    26029268, CurrentUsed:    25930900, CurrentFree:   188754076
    TaskName  TotalAllocated     TotalFreed       MaxHeld  CurrentHolding
   tRootTask         9014200         364056       8650144         8650144
    tTelnetd             144              0           144             144
  tTftpdTask            3584              0          3584            3584
   tPortmapd           18368              0         18368           18368
      tShell            3248              0          3248            3248
     isrTask           38256              0         38256           38256
    tExcTask               0          18472             0               0
         tme         3081210            144       3081066         3081066
         ipc           33812            144         33812           33668
       evagt           24192              0         24192           24192
       lcMgr           34048            464         34016           33584
         dla            3112           3088          1568              24
--More--
```

Figure 424 show processes memory lp Command Example

Monitoring CPU and Memory Usage via SNMP

To gather CPU and memory utilization via SNMP for the primary RPM, use the FORCE10-CHASSIS-MIB (available from iSupport) and the objects in the chRpmUtilTable, shown in Table 62.

Table 62 CPU and Memory Utilization Objects

MIB Object	MIB Object Name	Description
1.3.6.1.4.1.6027.3.1.1.3.7.1.3	chRpmCpuUtil5Sec	CPU utilization percentage for last 5 seconds
1.3.6.1.4.1.6027.3.1.1.3.7.1.4	chRpmCpuUtil1Min	CPU utilization percentage for last 1 minute
1.3.6.1.4.1.6027.3.1.1.3.7.1.5	chRpmCpuUtil5Min	CPU utilization percentage for last 5 minutes
1.3.6.1.4.1.6027.3.1.1.3.7.1.6	chRpmMemUsageUtil	Total memory utilization in percentage

Recognizing a High CPU Condition

Recognize a high CPU condition by any of the messages in Message 12.

Message 12 High CPU Condition

```
Feb 13 13:56:16: %RPM1-S:CP %CHMGR-5-TASK_CPU_THRESHOLD: Cpu usage above threshold for task
"sysAdmTsk"(100.00%) in CP.

Feb 13 13:56:20: %RPM1-S:CP %CHMGR-5-CPU_THRESHOLD: Overall cp cpu usage above threshold.
Cpu5SecUsage (100)

Feb 13 13:56:20: %RPM1-S:CP %CHMGR-5-TASK_CPU_THRESHOLD_CLR: Cpu usage drops below
threshold for task "sysAdmTsk"(0.00%) in CP.
```

Troubleshooting a High CPU Condition

If FTOS indicates a high CPU condition or you suspect one:

Step	Task	Command Syntax	Command Mode
1	Enable **debug cpu-traffic-stats**, and monitor the output with the **show cpu-traffic-stats** command. These commands indicate the physical interface through which the questionable traffic is arriving.	**debug cpu-traffic-stats** **show cpu-traffic-stats**	CONFIGURATION EXEC Privilege
2	Review the **show ip traffic** command output. This command displays the types of IP traffic destined to the CPU.	**show ip traffic**	EXEC Privilege

Monitoring Hardware Components Via SNMP

The SNMP traps and OIDs in Table 63 provide information on C-Series hardware components.

Table 63 SNMP Traps and OIDs

OID String	OID Name	Description
RPM		
.1.3.6.1.4.1.6027.3.1.1.3.8	chRPMMajorAlarmStatus	Fault status of the major alarm LED on the RPM
.1.3.6.1.4.1.6027.3.1.1.3.9	chRPMMinorAlarmStatus	Fault status of the minor alarm LED on the RPM
.1.3.6.1.4.1.6027.3.1.1.4.0.11	chAlarmRpmUp	Trap generated when the status of primary or secondary RPM changes to up and running
.1.3.6.1.4.1.6027.3.1.1.4.0.12	chAlarmRpmDown	Trap generated when the status of primary or secondary RPM changes to down, either by software reset or by being physically removed from the chassis
Line Card		

Table 63 SNMP Traps and OIDs

OID String	OID Name	Description
.1.3.6.1.4.1.6027.3.1.1.2.3.1.15	chSysCardOperStatus	Operational status of the card. • If the chSysCardAdminStatus is up, the valid state is ready—the card is present and ready and the chSysCardOperStatus status is up. • If the chSysCardAdminStatus is down the service states can be: • **offline**: the card is not used. • **cardNotmatch**: the card does not match what is configured • **cardProblem**: a hardware problem has been detected on the card. • **diagMode**: the card is in the diagnostic mode. **Note:** chSysCardFaultStatus is supported only the C-Series.
.1.3.6.1.4.1.6027.3.1.1.4.0.1	chAlarmCardDown	Trap reported when a card operational status changes to down
.1.3.6.1.4.1.6027.3.1.1.4.0.2	chAlarmCardUp	Trap reported when a card operational status changes to up
.1.3.6.1.4.1.6027.3.1.1.4.0.3	chAlarmCardReset	Trap reported when a card is reset
.1.3.6.1.4.1.6027.3.1.1.4.0.7	chAlarmCardProblem	Trap reported when a card operational status changes to card problem
.1.3.6.1.4.1.6027.3.1.1.1.4.0.5	chAlarmCardMismatch	Trap generated when the configured card does not match the installed card
.1.3.6.1.4.1.6027.3.1.1.1.4.0.6	chAlarmCardRemove	Trap generated when a card is removed
Power Supply Unit		
.1.3.6.1.4.1.6027.3.1.1.2.1.1.2	chSysPowerSupplyOperStatus	Each entry in the chSysPowerSupplyTable includes a set of objects which describe the status of a particular power supply.
.1.3.6.1.4.1.6027.3.1.1.4.0.13	chAlarmPowerSupplyDown	Trap generated when the power supply status changes to non-operational
.1.3.6.1.4.1.6027.3.1.1.4.0.17	chAlarmPowerSupplyClear	Trap generated when the power supply status changes to operational.
.1.3.6.1.4.1.6027.3.1.1.4.0.32	chAlarmMajorPS	Trap generated when a power supply major alarm is issued
.1.3.6.1.4.1.6027.3.1.1.4.0.33	chAlarmMajorPSClr	Trap generated when a power supply major alarm is cleared
.1.3.6.1.4.1.6027.3.1.1.4.0.34	chAlarmMinorPS	Trap generated when a power supply minor alarm is issued
.1.3.6.1.4.1.6027.3.1.1.4.0.35	chAlarmMinorPSClr	Trap generated when a power supply minor alarm is cleared

Table 63 SNMP Traps and OIDs

OID String	OID Name	Description
Fan Tray		
.1.3.6.1.4.1.6027.3.1.1.2.2.1.2	chSysFanTrayOperStatus	Each entry in the chSysFanTrayTable includes a set of objects that describe the status of a particular fan tray, as identified by the chSysFanTrayIndex
.1.3.6.1.4.1.6027.3.1.1.4.0.36	chAlarmMinorFanBad	Trap generated when the status of one or more fans changes to down, and generates a minor alarm
.1.3.6.1.4.1.6027.3.1.1.4.0.21	chAlarmMinorFanBadClear	Trap generated when the minor alarm on the one or more fans is cleared
.1.3.6.1.4.1.6027.3.1.1.4.0.16	chAlarmFanTrayDown	Trap generated when all fans are down and/or when the fan tray status changes to missing or down
.1.3.6.1.4.1.6027.3.1.1.4.0.20	chAlarmFanTrayClear	Trap generated when all fans and/or the fan tray status changes to operational

Offline Diagnostics

Note: As the SFM on the C-Series is a logical concept only, the FORCE10-CHASSIS-MIB SFM-related SNMP alarms and traps are not used.

The offline diagnostics test suite is useful for isolating faults and debugging hardware.

Diagnostics are invoked from the FTOS CLI. While diagnostics are running, the status can be monitored via the CLI. The tests results are written to a file in flash memory and can be displayed on screen. Detailed statistics for all tests are collected. These statistics include:

- last execution time
- first and last test pass time
- first and last test failure time
- total run count
- total failure count
- consecutive failure count
- error code.

The diagnostics tests are grouped into three levels:

- **Level 0** checks the device inventory and verifies the existence of the devices (e.g., device ID test).
- **Level 1** verifies that the devices are accessible via designated paths (e.g., line integrity tests) and tests the internal parts (e.g., registers) of the devices.

- **Level 2** performs on-board loopback tests on various data paths (e.g., data port pipe and Ethernet).

Configuration Task List

Note: This procedure assumes you have already loaded an FTOS image. These instructions illustrates the process of running offline diagnostics using line cards, but the process is the same for RPMs. Only the command keyword **linecard** must change to **rpm**. See the Command Line Reference Guide for details.

1. Take the line card offline. page 673.
2. Run offline diagnostics. page 674.
3. View offline diagnostic test results. page 674.
4. Bring the line card back online. page 677.

Important Points to Remember

- Offline diagnostics can only be run on offline line cards and on the standby route processor module (RPM). The primary RPM cannot be not tested.
- Diagnostics test only connectivity, not the entire data path.
- The complete diagnostics test suite normally runs for 4 to 6 minutes; the 48-port 1-Gigabit line card takes slightly longer than the 4-port 10-Gigabit line card.

Taking the Line Card Offline

Place the line card in an offline state using the **offline linecard** command, as shown in Figure 425.

Figure 425 offline linecard Command Example

```
Force10#offline linecard 5
00:50:05: %RPM0-P:CP %CHMGR-2-CARD_DOWN: Line card 5 down - card offline
00:50:05: %RPM0-P:CP %IFMGR-1-DEL_PORT: Removed port: Te 5/0-3
```

Use the **show linecard all** command to confirm offline status, as shown in Figure 426.

Figure 426 show linecard all Command Example

```
Force10#show linecard all
--  Line cards  --
Slot  Status        NxtBoot    ReqTyp   CurTyp   Version     Ports
---------------------------------------------------------------------------
  0   not present
  1   online        online     E48TB    E48TB    2.2.1.1     48
  2   not present
  3   not present
  4   not present
  5   not present
  6   offline       online     E48TB    E48TB    2.2.1.1     48
Force10#
```

Running Offline Diagnostics

Start diagnostics on the line card using the **diag linecard** command, as shown in Figure 427.

Figure 427 diag linecard Command Example

```
Force10#diag linecard 5
Force10#00:50:44: %EX4PB:5 %DIAGAGT-6-DA_DIAG_STARTED: Starting diags on slot 5
00:50:44 : Approximate time to complete these Diags ... 5 Min
Force10#
```

Viewing Offline Diagnostic Test Results

Use the **show diag** command to view a brief report of the test results, as shown in Figure 428.

Figure 428 show diag linecard Command Example

```
Force10#show diag linecard 5
Diag status of Linecard slot 5:
-------------------------------------------------------------
  Card is currently offline.
  Card alllevels diag issued at THU FEB 08, 2018 04:10:06 PM.
  Current diag status:            Card diags are in progress.
--------------------------------------------------------------
00:54:19 : Diagnostic test results are stored on file: flash:/TestReport
-LC-5.txt
00:54:19: %EX4PB:5 %DIAGAGT-6-DA_DIAG_DONE: Diags finished on slot 5
Force10#
```

Use the **show file flash:/*filename*** view the detailed test results in the test report saved to flash memory on the RPM. Use the command. Figure 429 shows the filename of the test results, and Figure 429 shows the contents of the file.

Note: Report any test failures to your Force10 Networks technical support engineer.

Figure 429 Viewing Offline Diagnostics Test Results

```
Force10#show diag linecard 5
Diag status of Linecard slot 5:
-------------------------------------------------------------------
    Card is currently offline.
    Card alllevels diag issued at THU FEB 08, 2018 04:10:05 PM.
    Current diag status:            Card diags are done.
    Duration of execution:          3 min 35 sec.
    Diagonostic test results located:       flash:/TestReport-LC-5.txt
-------------------------------------------------------------------
*********************************C-Series Diagnostics********************
LCM Board serial Number : 0060384
CPU Version : Line Processor: AMCC 440GX (rev D)
FPGA firmware Version : 1.20
Diag image based on build : CS-1-1-509
LCM Board Voltage levels - 3.260000 V, 2.480000 V, 1.770000 V, 1.500000 V, 1.200
000 V
LCM Board temperature : 29 Degree C
LCM present on Slot : 5.
*********************LCM EEPROM INFO**************************
********MFG INFO*******************
Data in Lp Eeprom is listed...........
Vendor Id: 00
Country Code: 01
Date Code: 01012007
Serial Number: 0060384
Part Number: 1234
Product Revision: 1
Product Order Number: LC-CB-10GE-4P
Card Id: 402
*********SW INFO*******************
Data in Lp Eeprom is listed...........
Chassis Type: 6
Chassis Mode: 4
Backplane version: 1
******************** Starting iteration 1. ********************
******************* LEVEL 0 DIAGNOSTICS *********************
Test 1 - NVRAM Access test .......................................... PASS
Test 3 - FPGA Access Test ........................................... PASS
Test 4 - Probing Test for volt/Temp sensor .......................... PASS
Test 6 - Probing for POE device 1 ..................................     NOT APPL
Test 7 - Probing for POE device 2 ..................................     NOT APPL
Test 8 - Probing for POE device 3 ..................................     NOT APPL
Test 9 - Probing for POE device 4 ..................................     NOT APPL
Test 10 - EEPROM access test ........................................ PASS
Test 11 - PCI CPU BRG 0 Level0 Test ................................. PASS
Test 12 - PCI F10 DEV 0 Level0 Test ................................. PASS
Test 13 - PCI F10 DEV 1 Level0 Test ................................. PASS
Test 14 - PCI F10 DEV 2 Level0 Test ................................     NOT APPL
Test 15 - PCI F10 DEV 3 Level0 Test ................................     NOT APPL
Test 16 - PCI F10 DEV 4 Level0 Test ................................. PASS
Test 17 - PCI F10 DEV 5 Level0 Test ................................. PASS
Test 21 - PHY MGT DEV 0 Level0 Test ................................. PASS
Test 22 - PHY IPC DEV 0 Level0 Test ................................. PASS
Test 23 - PHY IPC DEV 1 Level0 Test ................................. PASS
Test 26 - FPGA Flash Primary Test ..................................      PASS
Test 27 - FPGA Firmware Compare Test ...............................      PASS
*************** LEVEL 1 DIAGNOSTICS*****************************
Test 100 - Performing SDRAM ECC test ................................ PASS
Test 101 - SDRAM Pseudo Random test ................................. PASS
```

Figure 430 Viewing Offline Diagnostics Test Results (continued)

```
Test 107 - NVRAM Address Line test .................................. PASS
Test 108 - NVRAM Data Line Test ..................................... PASS
Test 110 - NVRAM Read Write test .................................... PASS
.Test 111 - FLASH Write Read test .................................... PASS
Test 112 - FPGA Registers Verification Test ......................... PASS
Test 113 - FPGA Level1 Test ......................................... PASS
Test 114 - FPGA Data bus walking 0 and 1's test ..................... PASS
Test 115 - Temp/volt monitor write read test ........................ PASS
Test 116 - Reg verification test POE manager 1 ......................    NOT APPL
Test 117 - Reg verification test POE manager 2 ......................    NOT APPL
Test 118 - Reg verification test POE manager 3 ......................    NOT APPL
Test 119 - Reg verification test POE manager 4 ......................    NOT APPL
Test 120 - EEPROM write read test ................................... PASS
Test 122 - Blinking Status LEDs Test ................................ PASS
Test 123 - PHY MGT DEV 0 Level1 Test ................................ PASS
Test 124 - PHY IPC DEV 2 level1 Test ................................ PASS
Test 125 - PHY IPC DEV 3 level1 Test ................................ PASS
Test 126 - Local Eeprom MFG block checksum test ..................... PASS
Test 127 - Local Eeprom SW block checksum test ...................... PASS
Test 128 - 3.3 V Brick Load test .................................... PASS
Test 129 - 2.5 V Brick Load test .................................... PASS
Test 130 - 1.5 V Brick Load test .................................... PASS
Test 131 - 1.25 V Brick Load test ................................... PASS
Test 132 - 1.8 V Brick Load test .................................... PASS
Test 133 - XFP Verification Test 0 .................................. PASS
Test 134 - XFP Verification Test 1 .................................. PASS
Test 135 - XFP Verification Test 2 .................................. PASS
Test 136 - XFP Verification Test 3 .................................. PASS
Test 137 - XFP Verification Test 4 ..................................    NOT APPL
Test 138 - XFP Verification Test 5 ..................................    NOT APPL
Test 139 - XFP Verification Test 6 ..................................    NOT APPL
Test 140 - XFP Verification Test 7 ..................................    NOT APPL
************* LEVEL 2 DIAGNOSTICS*********************************
Test 200 - MAC MGT DEV 0 Level2 Test ................................ PASS
Test 201 - MAC IPC DEV 0 Level2 Test ................................ PASS
Test 202 - MAC IPC DEV 1 Level2 Test ................................ PASS
Test 203 - PHY MGT DEV 0 Level2 Test ................................ PASS
Test 204 - PHY IPC DEV 0 Level2 Test ................................ PASS
Test 205 - PHY IPC DEV 1 Level2 Test ................................ PASS
Test 216 - F10-SFM Port wise Traffic Test with PHY Loopback ......... PASS
Test 218 - LCM SNAKE Test using CPU Traffic with MAC Loopback ....... PASS
Test 219 - LCM SNAKE Test using CPU Traffic with PHY Loopback ....... PASS
Test 220 - LCM Port wise Traffic Test using CPU traffic - MAC Loopbac PASS
Test 221 - LCM Port wise Traffic Test using CPU traffic - PHY Loopbac PASS
Test 222 - POE I2C Interface stress test on Unit - 0 ................    NOT APPL
Test 223 - POE I2C Interface stress test on Unit - 1 ................    NOT APPL
Test 224 - POE I2C Interface stress test on Unit - 2 ................    NOT APPL
Test 225 - POE I2C Interface stress test on Unit - 3 ................    NOT APPL
*********** FORCE10 C series Diagnostics END******************************
Number of Diagnostics performed 69
Number of Diagnostics failed 0
 End of Diags
 Duration of execution:        3 min 18 sec
-------------------------------------------------------------------
Force10#
```

Bringing the Line Card Online

Bring the card back online using the **online linecard** command. The card will be reset. Use the **show linecard all** command to verify the online status of the line card.

Chapter 40

E-Series Debugging and Diagnostics

C-Series	NO
E-Series	✓

Platform Specific Feature: E-Series Debugging and Diagnostics is supported on E-Series only.

In addition to the FTOS high availability features, E-Series and FTOS support several diagnostics and debug features that are integral components to delivering maximum uptime. These features consist of the following:

- System Health Checks on page 680
- SFM Channel Monitoring on page 686
- IPC Timeout Information Collection on page 688
- Debug Commands on page 690
- Show Hardware Commands on page 690
- Offline Diagnostics on page 691
- Parity Error Detection Reporting on page 692
- Trace Logs on page 694
- Recognizing a High CPU Condition on page 697

Note: These diagnostics and debugability features are available on TeraScale systems only, unless specifically noted.

Overview

The FTOS diagnostics and debugging features are a proactive approach to maximizing system uptime and reducing meantime to resolution (MTTR) when a problem occurs. This feature set includes a combination of proactive and reactive components designed to alert the user to network events, automatically collect information on the event, and allow the user to collect diagnostic information from the system.

- Proactive component
  - — The system health check detects, reports, and takes action on an error in real time.
  - — When an automatic corrective action is not appropriate, the system health check reports the detected anomaly, in real time, via a syslog message and/or SNMP.

- Reactive component
  - — When an error condition is asserted, appropriate show and debug commands are available to assist in identifying the condition as well as rapid fault isolation.

System Health Checks

An automatic runtime loopback test monitors the overall health status of the dataplane. This loopback test runs while the system's switch fabric is up; detecting potential blockages in the system's usual data transfer path.

Runtime dataplane loopback check

This is a dataplane loopback health check. Periodically, the primary RPM and each line card, in an online start, sends a packet through the dataplane channels, verifying the packet is returned, and then verifying the dataplane is functioning as expected (see Figure 431). Both portpipes on the line cards are tested.

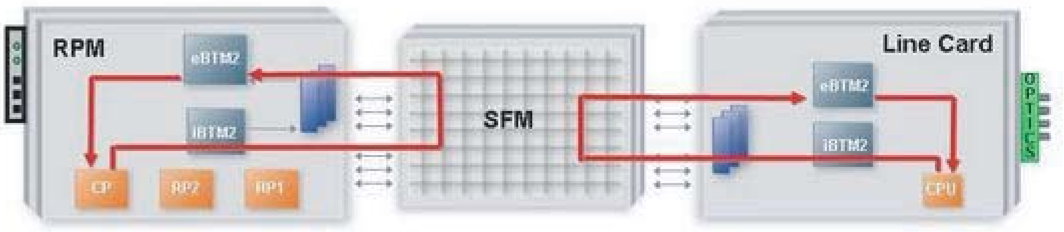

Figure 431 Dataplane Loopback

If three consecutive packets are lost, an error message is logged and then one of the following happens:

- The RPM-SFM runtime loopback test failure initiates an SFM *walk* whenever it is enabled, feasible and necessary. The system automatically places each SFM (in sequential order) in an offline state, runs the loopback test, and then places the SFM back in an active state. This continues until the system determines a working SFM combination. If no working combination is found, the system restores to the pre-walking SFM state and the switch fabric state remains up. No more SFM walks are conducted as long as the SFM settings remain unchanged (setting changes include SFM reset, power off/on, and hotswap). However, the runtime loopback tests will continue with failure messages being logged every five minutes.

Note: SFM walking assumes a chassis with the maximum number of SFMs in an active state.

The loopback runtime test results reflect the overall health status of the dataplane. SFM walking can help to identify a single faulty SFM which is persistently dropping all traffic. For any partial packet loss, the loopback test results can only indicate that there is partial packet loss on the dataplane.

When an automatic SFM walk is conducted, events are logged to indicate the start and completion of the SFM walk and the results. A complete system message set is shown below.

```
%TSM-2-RPM_LOOPBACK_FAIL: RPM-SFM dataplane loopback test failed
%TSM-2-SFM_WALK_START: Automatic SFM walk-through started
%TSM-6-RPM_LOOPBACK_PASS: RPM-SFM dataplane loopback test succeeded
%TSM-2-BAD_SFM_DISABLED: Bad SFM in slot 0 detected and disabled
%TSM-2-SFM_WALK_SUCCEED: Automatic SFM walk-through succeeded
```

- An SFM walk will not be able to identify multiple faulty SFMs, faulty linecards, or faulty RPM. In this case, the following event is logged.

```
%TSM-2-RPM_LOOPBACK_FAIL: RPM-SFM dataplane loopback test failed
%TSM-2-SFM_WALK_START: Automatic SFM walk-through started
%TSM-2-SFM_WALK_FAIL: Automatic SFM walk-through failed to identify single faulty SFM
```

- If a line card runtime loopback test fails, the system does *not* launch an SFM walk. A message is logged indicating the failure.

```
%TSM-2-RPM_LOOPBACK_FAIL: Linecard-SFM dataplane loopback test failed on linecard 6
```

The runtime dataplane loopback test is enabled by default. To disable this feature, use the following command.

Task	Command	Mode
Disable the runtime loopback test on the primary RPM and line cards. To re-enable, use the **no dataplane-diag disable loopback** command	**dataplane-diag disable loopback**	CONFIGURATION

Note: Disabling the runtime loopback test prevents the **sfm-walk** command and **sfm-bringdown** commands from taking effect.

Disable RPM-SFM Walk

If a full set of SFMs are online during the runtime loopback test and an RPM-SFM runtime loopback test failure occurs, an automatic SFM walk is launched in an attempt to determine if the failure is due to a faulty SFM. If confirmed, the single faulty SFM is identified and disabled by default.

To disable the automatic SFM walk that is launched after an RPM-SFM runtime loopback test failure, use the following command in CONFIGURATION mode.

Task	Command	Mode
Disable the automatic SFM walk that is launched after an RPM-SFM runtime loopback test failure. To re-enable the automatic SFM walk, use the **no dataplane-diag disable sfm-walk** command.	**dataplane-diag disable sfm-walk**	CONFIGURATION

Note: Disabling the **sfm-walk** command prevents the **sfm-bringdown** command from taking effect.

RPM-SFM Bring Down

If a full set of SFMs are online during the runtime loopback test and a RPM-SFM runtime loopback test failure occurs, an automatic SFM walk is launched in an attempt to determine if the failure is due to a faulty SFM. If confirmed, the single SFM is identified and disabled (bringdown) by default.

To disable the automatic bring-down of an SFM that is identified by the SFM walk during the RPM-SFM runtime loopback test, use the following CONFIGURATION mode command.

Task	Command	Mode
Disable the automatic bring down of the single faulty SFM identifed by the SFM walk during the RPM-SFM runtime loopback test. To re-enable the automatic bring down of an SFM, use the **no dataplane-diag disable sfm-bringdown** command.	**dataplane-diag disable sfm-bringdown**	CONFIGURATION

Manual Loopback Test

This manual dataplane loopback test is a supplemental test to the automatic runtime loopback test and can be initiated regardless if the runtime loopback test is enabled or disabled. Use this test to verify that the dataplane is actually functional *even* when a switch fabric status is down but there are at least (max-1) SFMs in active or diag failure state.

Task	Command	Mode
Execute a manual dataplane loopback test: • all-loopback – Both the RPM and the line card dataplane loopback test is done. • rpm-loopback – Only the RPM dataplane loopback test is done. This test can be run when the switch fabric is in either an operational or a non-operational state.	**diag sfm** [**all-loopback** | **rpm-loopback**]	EXEC

If the RPM-SFM or line card-SFM loopback test detects an SFM failure, an attempt is made to isolate a single faulty SFM by automatically *walking* the SFMs. For this failure case, error messages similar to the runtime loopback test error are generated.

Note: The dataplane runtime loopback configuration does not apply to this manual loopback test.

In the example in Figure 432, the manual loopback tests is successful, and no SFM failure is detected.

```
Force10#diag sfm all-loopback
Proceed with dataplane loopback test [confirm yes/no]:yes
SFM loopback test completed successfully.

Force10#
```

Figure 432 diag sfm all-loopback command Example

If the test passes when the switch fabric is down and there are at least (max-1) SFMs in the chassis, then the system will bring the switch fabric back up automatically. Like the runtime loopback test, the manual loopback test failure will not bring the switch fabric down.

Note: Line card-SFM loopback test failure, during the manual test, will trigger an SFM walk.

Power On/Off the SFM

If you suspect that an SFM is faulty and would like to manually disable it to determine whether any packet loss or forwarding issues are resolved, execute the following command.

Task	Command	Mode
Power on or off a specific SFM.	**power-**{**off** | **on**} **sfm** *slot-number*	EXEC

Note: Execute this command only during an offline diagnostics; this command may bring down the switch fabric.

When there are a full set of SFMs online, powering down one SFM will reduce the total bandwidth supported by the chassis, and may affect data flow. A warning message is issued at the command line that requires user confirmation to proceed with the command (Figure 433).

```
Force10#power-off sfm 0
SFM0 is active. Powering it off it might impact the data traffic.
Proceed with power-off [confirm yes/no]:yes
Feb 15 23:52:53: %RPM1-P:CP %CHMGR-2-MINOR_SFM: Minor alarm: only eight working
SFM
Force10#
```

Figure 433 power-off sfm command with data traffic warning message

Since this command is for diagnostic purposes, you can power off more than one SFM which may cause a switch fabric module to go down. A warning message is issued at the command line and requires user confirmation to proceed with the command (Figure 434).

```
Force10#power-off sfm 1
WARNING!! SFM1 is active. Powering it off it will cause Switch Fabric to go down!!
Proceed with power-off [confirm yes/no]:yes
Feb 16 00:03:19: %RPM1-P:CP %TSM-6-SFM_SWITCHFAB_STATE: Switch Fabric: DOWN
Feb 16 00:03:20: %RPM1-P:CP %CHMGR-0-MAJOR_SFM: Major alarm: Switch fabric down
Force10#
```

Figure 434 power-off sfm command with switch fabric down warning message

Once the SFM is powered off, the SFM status indicates that the SFM has been powered off by the user. Use the **show sfm all** command to display the status (Figure 435).

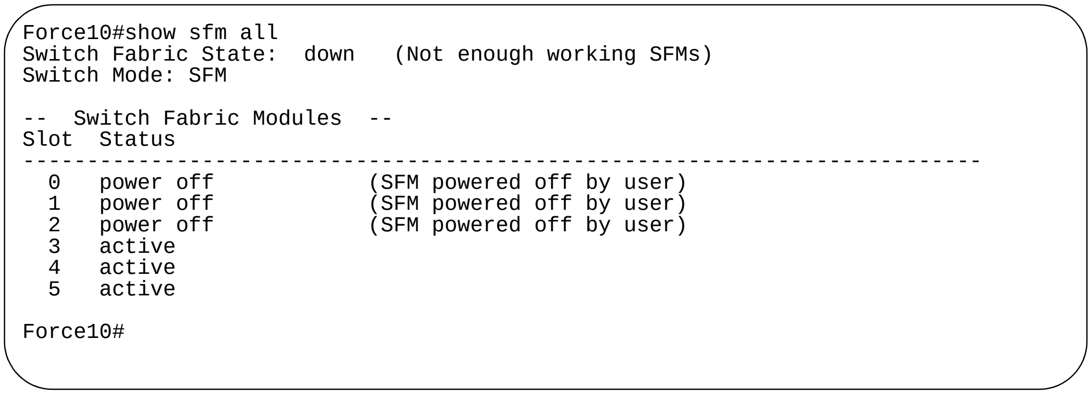

```
Force10#show sfm all
Switch Fabric State:  down   (Not enough working SFMs)
Switch Mode: SFM

--  Switch Fabric Modules  --
Slot  Status
---------------------------------------------------------------------------
  0   power off            (SFM powered off by user)
  1   power off            (SFM powered off by user)
  2   power off            (SFM powered off by user)
  3   active
  4   active
  5   active

Force10#
```

Figure 435 show sfm all command Example

Reset the SFM

When the SFM is taken offline due to an error condition, you can execute the **reset sfm** command and initiate a manual recovery.

Task	Command	Mode
Reset a specific SFM module (power-off and then power-on).	**reset sfm** *slot-number*	EXEC

When an error is detected on an SFM module, this command is a manual recovery mechanism. Since this command can be used with *live* traffic running, the switch fabric will not go down if the switch fabric is in an UP state. When there is a full set of SFMs online in the chassis, resetting one SFM will reduce the total bandwidth supported by the chassis and may effect data flow. A warning message is issued at the command line and requires user confirmation to proceed. (Figure 436)

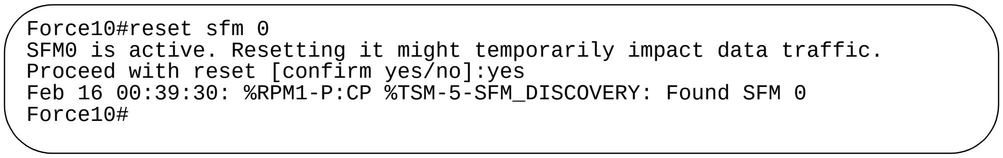

```
Force10#reset sfm 0
SFM0 is active. Resetting it might temporarily impact data traffic.
Proceed with reset [confirm yes/no]:yes
Feb 16 00:39:30: %RPM1-P:CP %TSM-5-SFM_DISCOVERY: Found SFM 0
Force10#
```

Figure 436 reset sfm command example

This command does not permit resetting any SFM when the system has (max-1) SFM and switch fabric is up (Figure 437)

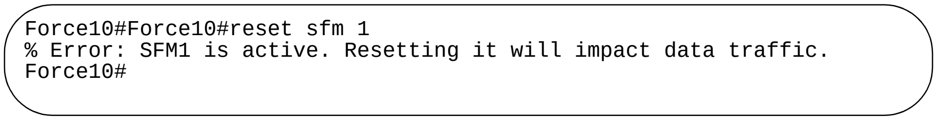

```
Force10#Force10#reset sfm 1
% Error: SFM1 is active. Resetting it will impact data traffic.
Force10#
```

Figure 437 reset sfm error message

➡ **Note:** Resetting an SFM in a power-off state is not permitted. Use the command **power-on sfm** to bring the SFM back to a power-on state.

SFM Channel Monitoring

In addition to monitoring the datapath, the SFM channels can be monitored using the Per-Channel Deskew FIFO Overflow (PCDFO) polling feature on all line cards and RPMs in both EtherScale and TeraScale E1200, E600, and E300 chassis. Like the datapath loopback feature, the PCDFO polling feature is enabled by default.

➡ **Note:** This feature is not supported on the E600i chassis.

The PCDFO polling feature monitors data received over the switch fabric. When a DFO error is detected, no automatic action is initiated by the system. The message issued is similar to:

```
%RPM1-P:CP %CHMGR-2-SFM_PCDFO: PCDFO error detected for SFM4
```

The following graphic illustrates the E600 and E1200 switch fabric architecture. Each ingress and egress Buffer and Traffic Management (BTM) ASIC maintains nine channel connections to the TeraScale Switch Fabric (TSF) ASIC.

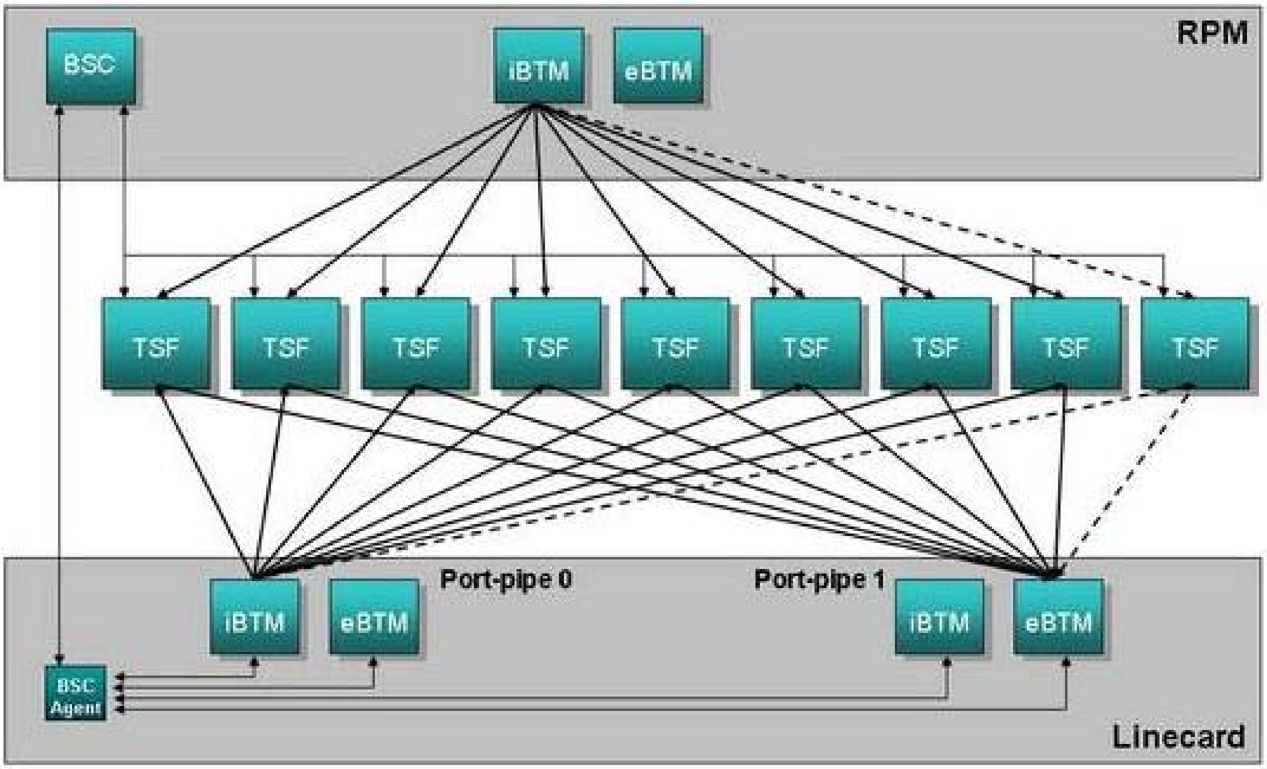

Responding to PCDFO Events

Troubleshooting PCDFO events requires applying some human intelligence to differentiate between transient and systematic failures.   PCDFO events can be caused by several factors, including:

- Backplane noise
- Data corruption
- Bad epoch timing
- Mis-configuration of backplane

There are two PCDFO error types: Transient and Systematic. Transient error are non-persistent events that occur as one-events during normal operation. Systematic errors are repeatable events. For example, some hardware device or component is malfunctioning in such a way that it persistently exhibits incorrect behavior.

For the transient case, PCDFO errors are not reported to the log. The hardware system automatically recovers from the error state, and the dataplane continues to function properly. In persistent case, PCDFO errors will appear in the log, and the error state is likely to remain if not handled.

With PCDFO error data alone, it is impossible to arrive at a conclusion which will pinpoint the cause for PCDFO error or reason for packets drop. For example, it is quite possible to have multiple line cards/RPM show different channels with PCDFO error. Nonetheless, PCDFO status is a very useful data point as an indication of the health of the dataplane, particularly when an error is persistent.

To disable the PCDFO polling feature, use the following command in CONFIGURATION mode.

Task	Command	Mode
Disable the PCDFO polling feature. To re-enable, use the **no dataplane-diag disable dfo-reporting** command.	**dataplane-diag disable dfo-reporting**	CONFIGURATION

Detection of a PCDFO event causes the system to generate a message similar to the following.

```
%RPM1-P:CP %CHMGR-2-SFM_PCDFO: PCDFO error detected for SFM #
```

Events are logged when PCDFO error first occurs on any SFM and when PCDFO error pattern changes.

No automatic action is taken by the system when a DFO error is detected. If such an error is reported, note the SFM slot number identified in the message and contact Force10 technical support. In addition, to confirm that the identified SFM needs to be replaced, use the **diag sfm all-loopback** to execute a manual dataplane loopback test.

IPC Timeout Information Collection

Each RPM consists of three CPUs:

- Control Processor (CP)
- Routing Processor 1 (RP1)
- Routing Processor 2 (RP2)

The three CPUs use Fast Ethernet connections to communicate to each other and to the line card CPUs using Inter-Processor Communication (IPC). The CP monitors the health status of the other processors using heartbeat messaging exchange.

```
%RPM1-P:CP %IPC-2-STATUS: target rp2 not responding
%RPM0-S:CP %RAM-6-FAILOVER_REQ: RPM failover request from active peer: Auto failover on failure
%RPM0-S:CP %RAM-6-ELECTION_ROLE: RPM0 is transitioning to Primary RPM.
%RPM0-P:CP %TSM-6-SFM_SWITCHFAB_STATE: Switch Fabric: UP
```

FTOS automatically saves critical information about the IPC failure to NVRAM. Such information includes:

- Status counters on the internal Ethernet interface
- Traffic profile of the inter-CPU bus

Upon the next boot, this information is uploaded to a file in the CRASH_LOG directory. Use the following command sequence beginning in EXEC mode to capture this file for analysis by the Force10 Networks TAC.

Step	Task	Command	Mode
1	Display the directories in flash memory. The output should include: **1 drwx 2048 Jan 01 1980 00:00:06 CRASH_LOG_DIR**	**dir flash:**	EXEC
2	Change to the CRASH_LOG directory.	**cd CRASH_LOG_DIR**	EXEC
3	View any saved files in the CRASH_LOG directory. The naming convention is: *sysinfo_RPMIDProcessorID _ timestamp* For example: **sysinfo_RPM1CP_20060616_013125** **sysinfo_RPM1RP1_20060616_013248** **sysinfo_RPM1RP2_20060616_013249**	**dir**	EXEC privilege
4	View the contents of the file.	**show file flash://CRASH_LOG_DIR/ [*file_name*]**	EXEC privilege

In a dual RPM system, the two RPMs send synchronization messages via inter-RPM communication (IRC). As described in the High Availability chapter, an RPM failover can be triggered by loss of the heartbeat (similar to a keepalive message) between the two RPMs. FTOS reports this condition via syslog messages, as follows:

```
20:29:07: %RPM1-S:CP %IRC-4-IRC_WARNLINKDN: Keepalive packet 7 to peer RPM is lost
20:29:07: %RPM1-S:CP %IRC-4-IRC_COMMDOWN: Link to peer RPM is down
%RPM1-S:CP %RAM-4-MISSING_HB: Heartbeat lost with peer RPM. Auto failover on heart beat lost.
%RPM1-S:CP %RAM-6-ELECTION_ROLE: RPM1 is transitioning to Primary RPM.
```

FTOS automatically saves critical information, about the IRC failure, to NVRAM. Use the same three-step procedure to capture this file for analysis by Force10 Networks.

FTOS actually saves up to three persistent files depending upon the type of failure. When reporting an RPM failover triggered by a loss of the IPC or IRC heartbeats, look for failure records in the following directories:

- Application or kernel core dump RP in the CORE_DUMP_DIR
- CP trace log file (look for a filename with the phrase “failure_trace”) in the TRACE_LOG_DIR
- RP and/or CP sysinfo file in the CRASH_LOG_DIR, as explained above

Debug Commands

FTOS supports an extensive suite of debug commands for troubleshooting specific problems while working with Force10 Networks technical support staff. All debug commands are entered in privileged EXEC mode. See the FTOS Command Reference for details.

Show Hardware Commands

The show hardware command tree consists of privileged EXEC commands created or changed specially for use with the E-Series. These commands display information from a hardware sub-component, such as the Buffer and Traffic Management (BTM) ASIC and the Forwarding and Packet Classification (FPC) ASIC. They should be used only under the guidance of Force10 Networks technical support staff.

The following table lists the show hardware commands. For detailed information on these and other commands, see the *FTOS Command Line Interface Reference* document.

Command	Description
show hardware rpm *slot-number* **mac counters** [**port** *port-number*] **clear hardware rpm** *slot-number* **mac counters**	View or clear the receive- and transmit-counters for the party-bus control switch on the IPC subsystem of the RPM.

show hardware rpm *slot-number* **cp** {**data-plane** | **management-port**} | **party-bus**} {**counters** | **statistics**}
show hardware rpm *slot-number* {**rp1** | **rp2**} {**data-plane** | **party-bus**} {**counters** | **statistics**}
	Display advanced debugging information for the RPM processors.

show hardware linecard *number* **port-set** *pipe-number* **fpc forward** {**counters** | **drops** | **spi** {**err-counters** | *spichannel#* **counters**} | **status**}	Display receive and transmit counters, error counters and status registers for the forwarding functional area of the FPC (flexible packet classification engine).
show hardware linecard *number* **port-set** *pipe-number* **fpc lookup detail**	Display diagnostic and debug information related to the lookup functional area of the Flexible Packet Classification (FPC).

Offline Diagnostics

These diagnostics can be useful for isolating faults and debugging TeraScale hardware installed in a chassis.

Diagnostics are invoked from the FTOS CLI. While diagnostics are running, the status can be monitored via the CLI. The tests results are written to a file in flash memory and can be displayed on screen. Detailed statistics for all tests are collected and include:

- last execution time
- first test pass time and last test pass time
- first test failure time and last test failure time
- total run count
- total failure count
- consecutive failure count
- error code

The diagnostics tests are grouped into three levels:

Level 0—Check the inventory of devices. Verify the existence of devices (e.g., device ID test).

Level 1—Verify the devices are accessible via designated paths (e.g., line integrity tests). Test the internal parts (e.g., registers) of devices.

Level 2—Perform on-board loopback tests on various data paths (e.g., data port pipe and Ethernet).

Important Points to Remember

- Offline diagnostics can be run only on an offline line card and on a standby route processor module (RPM). The primary RPM is not tested.
- Diagnostics test only connectivity and not the entire data path.
- A line card must be put into an offline state before diagnostics are run.
- Complete diagnostics test suite normally runs for 5 to 7 minutes on a single port-pipe line card and 12 to 15 minutes on a dual port-pipe line card. Running diagnostics on LC-EF-GE-90M cards may take slightly longer.

Offline Configuration Task List

Use the following steps to run offline diagnostics on the E-Series. This procedure assumes the FTOS image is installed.

1. Place the line card in an offline state with the **offline linecard** command. Use the **show linecard** command to confirm the new status.

```
Foce10#offline line 4
Mar 27 05:18:26: %RPM0-P:CP %CHMGR-2-CARD_DOWN: Line card 4 down - card offline
Mar 27 05:18:26: %RPM0-P:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Te 4/3
```

2. Start diagnostics on the line card with the **diag** command. The system will confirm that diagnostics tests are running by displaying the syslog message shown below.

   ```
   Force10#diag linecard 4 ?
   alllevels               Execute level 0-2 diags (default)
   level0                  Execute level 0 diags
   level1                  Execute level 1 diags
   level2                  Execute level 2 diags
   terminate               Stops the running test

   Force10#diag linecard 4
   Mar 27 01:54:00: %E12PD3:2 %DIAGAGT-6-DA_DIAG_STARTED: Starting diags on slot 4
   Mar 27 02:05:47: %E12PD3:2 %DIAGAGT-6-DA_DIAG_DONE: Diags finished on slot 4
   ```

3. Execute the **show diag** command to view a report of the test results.

   ```
   Force10#show diag linecard 4
   Diag status of Linecard slot 4:
   -------------------------------------------------------------------
       Card is currently offline.
       Card level0 diag issued at TUE Mar 27, 2007 05:19:35 AM.
       Current diag status:             Card diags are done (FAIL).
       Duration of execution:           0 min 0 sec.
       Number of diags performed:       39
       Number of diags passed:          36
       Number of diags failed:          3
       Number of notification received: 80
       Last notification received at:   TUE Mar 27, 2007 05:19:35 AM
   ```

4. Report any test failures to your Force10 Networks technical support engineer.
5. Bring the card back online with the **online linecard** {*slot#*} command. The card will be reset.

Parity Error Detection Reporting

There are two classes of Parity Errors that devices can encounter: Transient and Real.

Transient Parity Error

A transient parity error implies a read value was somehow ‘corrupted in transit’; however, the actual memory is not corrupted. The response to a transient error is to simply monitor the line card operation to see if the report returns. As no memory is actually corrupted, there is no memory to recover.

A transient parity error is reported anytime that a parity error is seen without being able to determining the offending address. This means that the parity error was seen upon entry into the scanning routine and when the SRAM scan completed, no address location read resulted in another parity error indication. The two following messages indicate that this type of error was seen. The first message is put into the sysDrvLog and the hardware log, the second message goes into the hardware log, and the third message is a syslog message displayed on the system console. The following message is an example sysDrvLog Message:

```
Task(ppdT2lSramParityScan): FPC PP 0 Transient Parity Error: 0xd50000 = 0x11.0 Address Unknown
```

The following hardware log message is different than the sysDrvLog message:

```
[1/1 4:0:20] ****** HW ERR: POLLER-(ppdT2lSramParityScan):FPC PP 0 Transient Parity Error:
0xd50000 = 0x11.0 Address Unknown
```

The third log message, which is the syslog message. is shown below:

```
Mar 25 11:09:30: %RPM1-P:CP %CHMGR-2-CARD_PARITY_ERR: Linecard 6 pp 0 FPC SRAM Transient Error
parity error:  FPC DDR Bank [191:128]
```

The line card status for the transient error is not indicated as a transient error until five transient errors have occurred. Note that the text “Last Event” indicates that this is the last type of event seen. If a real parity event occurs again, the Parity Status will change to the status information for a real parity error.

```
Force10#show linecard 6

--  Line card 6 --
Status        : online
Next Boot     : online
Required Type : E48TF - 48-port 10/100/1000Base-T line card with RJ-45 interfaces (EF)
Current Type  : E48TF - 48-port 10/100/1000Base-T line card with RJ-45 interfaces (EF)
Hardware Rev  : Base  - 1.1  PP0 - 1.0  PP1 - 1.0
Num Ports     : 48
Up Time       : 3 min, 44 sec
FTOS Version  : 6.5.4.1
Jumbo Capable : yes
Boot Flash    : A: 2.3.1.3     B: 2.3.1.3 [booted]
Memory Size   : 268435456 bytes
Temperature   : 42C
Power Status  : AC
Voltage       : ok
Serial Number : 0045149
Part Number   : 7520016602 Rev 06
Vendor Id     : 04
Date Code     : 01442005
Country Code  : 01
Parity Status : Last Event - FPC DDR Bank [191:128], Transient Failure, Running Count 5
```

Real Parity Error

A Real Parity Error implies a persistent corrupted memory location and the only means of recovery is to reboot the line card.

This type of parity error is reported anytime a direct correlation is seen between the read of an address location and the changing of the parity error status indication via a chip register read. Because this scan is done on a “live” system, a verification check is conducted on each parity error indication. The offending location is read a second time to ensure that no one else (possibly hardware) caused a parity error in another location. The parity error status register is then read again. If a parity error indication is shown, this address location is marked as having a parity error.

The messages shown below are generated for each “real” parity error. The first is put into the sysDrvLog and the hardware log, the second message goes into the hardware log, and the third message is a syslog message displayed on the system console. The following message is an example sysDrvLog Message:

```
Task(ppdT2lSramParityScan): FPC Parity Error Reported PP 0 Address 0x11ffe00 Index 0x7fff0
Status Register 0x11.0x0
```

The following hardware log message is different than the sysDrvLog message:

```
[1/1 4:1:14] ****** HW ERR: POLLER-(ppdT2lSramParityScan):FPC Parity Error Reported PP 0
Address 0x11ffe00 Index 0x7fff0 Status Register 0x11.0x0
```

The third log message which is a syslog message is shown below:

```
Mar 25 11:13:21: %RPM1-P:CP %CHMGR-2-CARD_PARITY_ERR: Linecard 6 pp 0 FPC SRAM Hard parity
error:  FPC DDR Bank [191:128] Address 0x11ffe00 Index 0x7fff0 Check the Hardware Log
```

It should also be noted that there may be more than one real parity error but only one syslog message is displayed for each scan of the SRAM in response to a noted parity error. To see if there have been more than just the single real parity error, the hardware log should be queried.

```
Force10#show linecard 6

--  Line card 6 --
Status        : online
Next Boot     : online
Required Type : E48TF - 48-port 10/100/1000Base-T line card with RJ-45 interfaces (EF)
Current Type  : E48TF - 48-port 10/100/1000Base-T line card with RJ-45 interfaces (EF)
Hardware Rev  : Base  - 1.1  PP0 - 1.0  PP1 - 1.0
Num Ports     : 48
Up Time       : 6 min, 30 sec
FTOS Version  : 6.5.4.1
Jumbo Capable : yes
Boot Flash    : A: 2.3.1.3     B: 2.3.1.3 [booted]
Memory Size   : 268435456 bytes
Temperature   : 43C
Power Status  : AC
Voltage       : ok
Serial Number : 0045149
Part Number   : 7520016602 Rev 06
Vendor Id     : 04
Date Code     : 01442005
Country Code  : 01
Parity Status : Last Event - FPC DDR Bank [191:128], Real Failure, Address 0x11ffe00
```

Trace Logs

In addition to the syslog buffer, FTOS buffers trace messages which are continuously written by various FTOS software tasks to report hardware and software events and status information. Each trace message provides the date, time, and name of the FTOS process. All messages are stored in a ring buffer and can be saved to a file either manually or automatically upon failover.

- There are three trace buffers for CP: software, hardware, and command-history.
- There are two trace buffers for LP: software and hardware.

Buffer Full Condition

When the trace ring buffer is full, the trace lines are saved as a file into the flash (e.g. hw_trace_RPM0CP.0). When the buffer fills for the second time, it is saved as hw_trace_RPM0CP.1, and so on until hw_trace_RPM0CP.4. From the sixth time onwards, when the trace buffer fills, the second trace file (hw_trace_RPM0CP.1) is overwritten.

Trace file hw_trace_RPM0CP.0 is not overwritten so that chassis bootup message are preserved.

FTOS uses a similar approach to saving the various trace messages for CP and all LPs.

The CP and LP trace file names are:

Note: Line card 1 is taken as the example for the following filenames.

- **CP [SW trace]** : sw_trace_RPM0CP.0, sw_trace_RPM0CP.1, sw_trace_RPM0CP.2, sw_trace_RPM0CP.3 and sw_trace_RPM0CP.4
- **CP [HW trace]** : hw_trace_RPM0CP.0, hw_trace_RPM0CP.1, hw_trace_RPM0CP.2, hw_trace_RPM0CP.3 and hw_trace_RPM0CP.4
- **LP [SW trace]** : sw_trace_LPX.0, sw_trace_LP1.1, sw_trace_LP1.2, sw_trace_LP1.3 and sw_trace_LP1.4
- **LP [HW trace]** : hw_trace_LPX.0, hw_trace_LP1.1, hw_trace_LP1.2, hw_trace_LP1.3 and hw_trace_LP1.4

Trace files are saved in the directory flash:/TRACE_LOG_DIR/TRACE_CURR_BOOT. Upon a system reload this directory is renamed flash:/TRACE_LOG_DIR/TRACE_LAST_BOOT, and a fresh empty flash:/TRACE_LOG_DIR/TRACE_CURR_BOOT directory is created.

Manual Reload Condition

When the chassis is reloaded manually (through the CLI), trace messages in all of the buffers (software and hardware) in CP and linecards are saved to the flash as reload_traceRPM0_CP and reload_traceLP1 in flash:/TRACE_LOG_DIR/TRACE_CURR_BOOT. After reload, you can see these files in flash:/ TRACE_LOG_DIR/TRACE_LAST_BOOT..

When the trace messages are being saved on reload, Message 13 is displayed.

Message 13 Saving Trace Messages

```
Starting to save trace messages… Done.
```

The CP and LP trace file names at chassis reload are:

- **CP:** reload_traceRPM0_CP
- **LP:** reload_traceLP1

CP/RP1/RP2 Software Exceptions

When a RPM resets due to an RP1 or RP2 software exception, the linecard trace files are saved to flash:/ TRACE_LOG_DIR directory.

The CP and LP trace file names in the case of a software exception are:

- **CP:** failure_trace_RPM1_CP
- **LP:** failure_trace_RPM1_LP1

With a dual RPM system, linecard trace logs are saved in the case of a CP, RP1 or RP2 crash. Linecard trace logs are saved in the new primary RPM. Here the name of the line card trace file helps in identifying a failed RPM. That is, if RPM0 fails, then the trace files are saved in RPM1 with filename failure_trace_RPM0_LP1. This is because the failover happens before the linecard traces are saved.

Viewing Trace Buffer Content

The command-history trace feature captures all commands entered by all users of the system with a time stamp and writes these messages to a dedicated trace log buffer. The system generates a trace message for each executed command. No password information is saved to the file.

To view the command-history trace, use the **show command-history** command, as shown in Figure 438.

```
C300-TAC-B6#show command
[12/5 10:57:8]: CMD-(CLI):service password-encryption
[12/5 10:57:12]: CMD-(CLI):hostname Force10
[12/5 10:57:12]: CMD-(CLI):ip telnet server enable
[12/5 10:57:12]: CMD-(CLI):line console 0
[12/5 10:57:12]: CMD-(CLI):line vty 0 9
[12/5 10:57:13]: CMD-(CLI):boot system rpm0 primary flash://FTOS-CB-1.1.1.2E2.bin
```

Figure 438 show command-history Command Example

Writing the Contents of the Trace Buffer

The trace logs are saved to automatically but you can save the contents of a buffer manually via the CLI.

To manually write the contents of a trace buffer on CP to a file on the flash:

Step	Task	Command Syntax	Command Mode
1	Write the buffered trace log to flash.	**upload trace-log cp** [**cmd-history** | **hw-trace** | **sw-trace**]	EXEC privilege

To manually write the contents of a trace buffer on LP to a file on the flash:

Step	Task	Command Syntax	Command Mode
1	Write the buffered trace log to flash.	**upload trace-log** [**rp1** | **rp2** | **linecard**] *number* [**hw-trace** | **sw-trace**]	EXEC privilege

Recognizing a High CPU Condition

Recognize a high CPU condition by any of the messages in Message 14.

Message 14 High CPU Condition

```
Feb 13 13:56:16: %RPM1-S:CP %CHMGR-5-TASK_CPU_THRESHOLD: Cpu usage above threshold for task
"sysAdmTsk"(100.00%) in CP.

Feb 13 13:56:20: %RPM1-S:CP %CHMGR-5-CPU_THRESHOLD: Overall cp cpu usage above threshold.
Cpu5SecUsage (100)

Feb 13 13:56:20: %RPM1-S:CP %CHMGR-5-TASK_CPU_THRESHOLD_CLR: Cpu usage drops below
threshold for task "sysAdmTsk"(0.00%) in CP.
```

Chapter 41 Streamline Upgrade

This section covers the following topics:

Upgrading the Boot Image

The upgrade process is a two-part process:

1.Upgrading the boot image(s)
2.Upgrading the boot profile by downloading runtime image files (if appropriate)

The boot image is stored onto the BOOT FLASH device. When the system boots properly, two identical images are created in BOOT FLASH.

When the RPM reloads or reboots, two possible events can happen:

- The new boot image loads properly, and the new boot image overwrites the old image in BOOT FLASH.
- The new image fails to load and is overwritten by the old image (after messages are generated, etc.) Although the upgrade did not occur, the system is still able to boot to the boot CLI.

Upgrading the Runtime Image

If the new runtime image is stored on a remote file system, the upgrade is simply a matter of changing the boot profile to point to the appropriate locations.

If the new runtime image is to be executed from a local file system, then upgrading means copying the file to the local file system and changing the boot profile as needed.

Streamlined Upgrade of the Runtime Image

The streamlined upgrade process copies the runtime image to a local file system and changes the boot profile as needed. This process avoids the possibility of operator error.

Once the copy command is run, you must reset the standby RPM and manually failover the RPM to the other RPM:

Step	Command Syntax	Command Mode	Usage
1	**copy** *ftp:file_url_rpm1flash:filename* **boot-image synchronize-rpm** [**external**]	EXEC privilege	Copy the runtime image to the local file system.

Upgrading FPGAs on the C-Series

C-Series	✓
E-Series	NO

Platform Specific Feature: The section Upgrading FPGAs on the C-Series applies to C-Series only.

There are two FPGAs on C-Series line cards and RPMs. One is the primary FPGA (A), and one is the backup (B). The primary FPGA image might occasionally require an upgrade. Perform an upgrade only when the system instructs you to to do so; contact the Technical Assistance Center if you have any questions.

Verifying that an FPGA Upgrade is Required

The system displays a syslog message during bootup (Message 15) if an FPGA image requires an upgrade.

Message 15 FPGA Upgrade Required on C-Series

```
% Error: Incompatible FPGA version detected, mandatory upgrade needed.
```

You can also use the command **show revision**, as shown in Figure 439. The running and required version numbers differ if the image requires an upgrade.

Figure 439 Identifing C-Series Cards that need an FPGA Upgrade

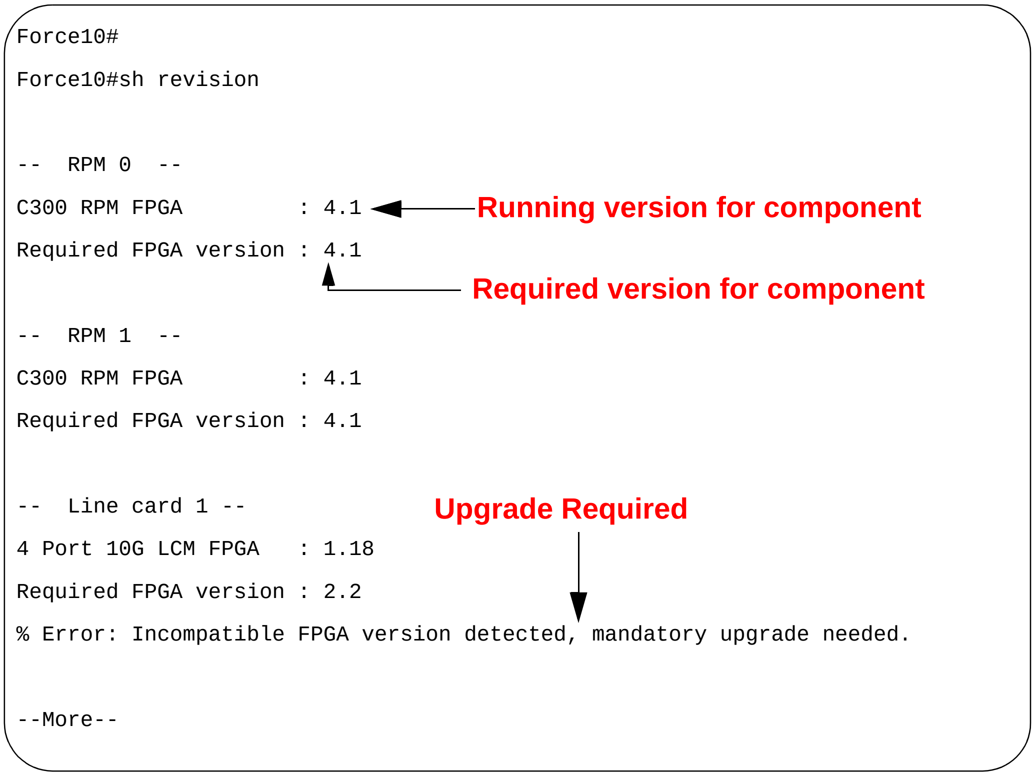

```
Force10#
Force10#sh revision

--  RPM 0  --
C300 RPM FPGA          : 4.1
Required FPGA version : 4.1

--  RPM 1  --
C300 RPM FPGA          : 4.1
Required FPGA version : 4.1

--  Line card 1 --
4 Port 10G LCM FPGA   : 1.18
Required FPGA version : 2.2
% Error: Incompatible FPGA version detected, mandatory upgrade needed.

--More--
```

Determine if the running FTOS version has the required image using the command **show os-version**, as shown in Figure 440.

Figure 440 Determining if the Running FTOS Image Contains the Required FPGA Version

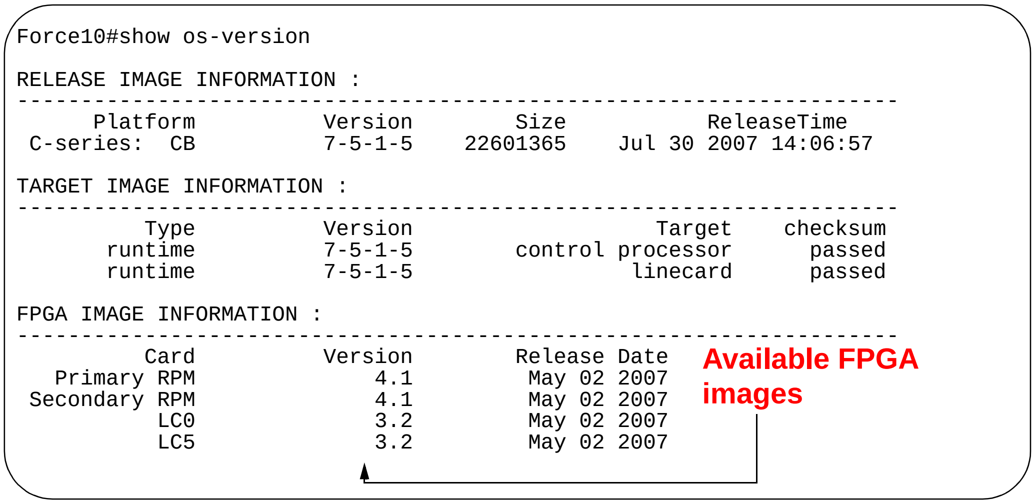

```
Force10#show os-version
RELEASE IMAGE INFORMATION :
---------------------------------------------------------------------
      Platform          Version       Size       ReleaseTime
 C-series:  CB          7-5-1-5      22601365    Jul 30 2007 14:06:57

TARGET IMAGE INFORMATION :
---------------------------------------------------------------------
          Type          Version                  Target     checksum
       runtime          7-5-1-5          control processor    passed
       runtime          7-5-1-5                   linecard    passed

FPGA IMAGE INFORMATION :
---------------------------------------------------------------------
          Card          Version          Release Date
   Primary RPM              4.1          May 02 2007
 Secondary RPM              4.1          May 02 2007
           LC0              3.2          May 02 2007
           LC5              3.2          May 02 2007
```

Upgrading the FPGA on a Line Card

Complete the upgrade procedure from the console. The FPGA cannot be upgraded via SSH or Telnet. Line cards must be upgraded one at a time. Only the primary FPGA image can be upgraded.

To upgrade the FPGA image on a line card:

Step	Task	Command	Command Mode
1	Shut down all of the interfaces on the card.	**shutdown**	INTERFACE
2	Upgrade the FPGA image **upgrade primary-fpga-flash**.	**upgrade primary-fpga-flash** *linecard number*	EXEC Privilege
3	Power cycle the line card.	**reset** *linecard number* **power-cycle**	EXEC Privilege

Note: You must power-cycle the card so that it can upload the new FPGA image.

Figure 441 and Figure 442 show the output of an upgrade using the command **upgrade primary-fpga-flash**.

- Figure 441 shows a successful upgrade.
- Figure 442 shows a failed upgrade because the flash selector is set to B.

Figure 441 Upgrading the C-Series Primary FPGA Image

```
Force10#upgrade primary-fpga-flash linecard 0


Proceed to upgrade primary fpga flash for linecard 0 from version 3.2 to 3.2
[yes/no]: yes

FPGA image upgrade is in progress. Please do NOT power off the card.

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!

linecard 0 primary fpga flash upgraded to version 3.2.


Force10#
```

Figure 442 Failed C-Series FPGA Upgrade due to Flash Selector set to B

```
Force10#upgrade primary-fpga-flash linecard 0

% Error: Fpga flash selection set to B. Please contact technical support.

Force10#
```

Upgrading the FPGA on an RPM

The FPGA images on the RPMs can be upgraded in parallel. Force10 recommends upgrading both RPMs to prevent having to repeat the procedure at another time.

To upgrade the RPM FPGA image:

Step	Task	Command Syntax	Command Mode
1	Shut down all of the interfaces on the card.	**shutdown**	INTERFACE
2	Upgrade the FPGA image on the standby RPM.	**upgrade primary-fpga-flash rpm**	EXEC Privilege
3	Power-cyle the RPM.	**reset rpm** *number* **power-cycle**	EXEC Privilege

Note: You must power-cycle the card so that the new FPGA image can be uploaded.

Restoring the FPGA

In case of an image failure on the primary FPGA, the image on the backup FPGA can be loaded onto the primary using the following commands:

Note: Only the primary FPGA can be restored.

Step	Task	Command Syntax	Command Mode
1	Shut down all of the interfaces on the card.	**shutdown**	INTERFACE
2	Restore the image on the primary FPGA for an RPM or line card.	**restore fpga-flash** [**rpm** | **linecard**] *number*	EXEC Privilege
3	Power-cyle the card.	**reset** [**rpm** | **linecad**] *number* **power-cycle**	EXEC Privilege

Note: If the FPGA is not restored after the power-cycle, contact the Force10 Technical Assistance Center.

The flash selector may be set to either A or B. Figure 441 shows a restoration of the primary flash using the command **restore fpga-flash**.

Figure 443 Restoring the C-Series Primary FPGA Flash

```
Force10#restore fpga-flash linecard 0

Proceed to restore primary fpga flash for linecard 0 [yes/no]: y
FPGA image restore is in progress. Please do NOT power off the card.
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!
linecard 0 primary fpga flash image restored.
```

Upgrading FPGAs on the E-Series

C-Series	NO
E-Series	✓

Platform Specific Feature: The section Upgrading FPGAs on the E-Series applies to E-Series only.

E-Series Switch Fabric Modules (CC-E-SFM3) have two two FPGAs that the SFMs use to boot. One is the primary FPGA (A) and one is the backup (B), which is used only if the primary fails. The primary FPGA might occasionally require an upgrade when performance enhancements become available.

Perform an upgrade only when the system instructs you to to do so; contact the Technical Assistance Center if you have any questions.

Note: Upgrading the SFM FPGA image brings the SFM down, and can bring the entire switch fabric down depending on the platform. Force10 recommends that you perform this upgrade only during a maintenance window.

Verifying that an FPGA Upgrade is Required

During SFM bootup, the SFM compares the loaded FPGA image version against the FPGA version in the FTOS image. If there is a mismatch, the FGPA image requires an upgrade, and the system displays a syslog message (Message 15). Since this check is performed upon SFM bootup, Message 15 is displayed upon a hot insertion as well as during a system bootup.

Message 16 FPGA Upgrade Required on E-Series

```
%RPM0-P:CP %TSM-5-SFM_REVISION_UPGRADE: Upgrade to FPGA revision 0x7.7.7
```

A syslog message is displayed if the SFM cannot boot from the primary FPGA and boots from the backup. Use this message as an alert to schedule an SFM upgrade during the start of your next maintenance window.

Upgrading the FPGA on the SFM

Only active SFMs can be upgraded, and an SFM reset is required on upgraded SFMs after the upgrade. Only one SFM upgrade operation can be initialized at any given time. If an upgrade request is made while one is already in progress, the system displays Message 17.

Message 17 FPGA Upgrade Already In Progress

```
%% Error: SFM upgrade already in progress.
```

To upgrade the FPGA image the SFM:

Step	Task	Command	Command Mode
1	Upgrade one or all SFM FPGA images.	**upgrade sfm** [*slot* | **all**] [**autoreset**] • Enter the keyword **autoreset** to invoke a dialog that prompts you to reset all of the relevant SFMs. This relieves you of having to execute separate **reset** command after the upgrade.	EXEC Privilege
2	Enter YES at the prompt to proceed with the upgrade		
3	When the upgrade is complete, • If you entered the keyword **autoreset** in Step 1, after the upgrade is complete, enter YES at the prompt to reset all SFMs. • If you did not use the keyword **autoreset** in Step 1, reset all upgraded SFMs using the command shown.	**reset sfm** [*slot* | **all**]	EXEC Privilege

➡ **Note:** You must reset the SFMs so that they can upload the new FPGA image. If the **all** option is selected in Step 1 or Step 3, the entire switch fabric is brought down, and traffic flow is disrupted.

Figure 444 shows the output of the command **upgrade sfm** command using the **autoreset** option.

Figure 444 Upgrading E-Series SFM FPGAs

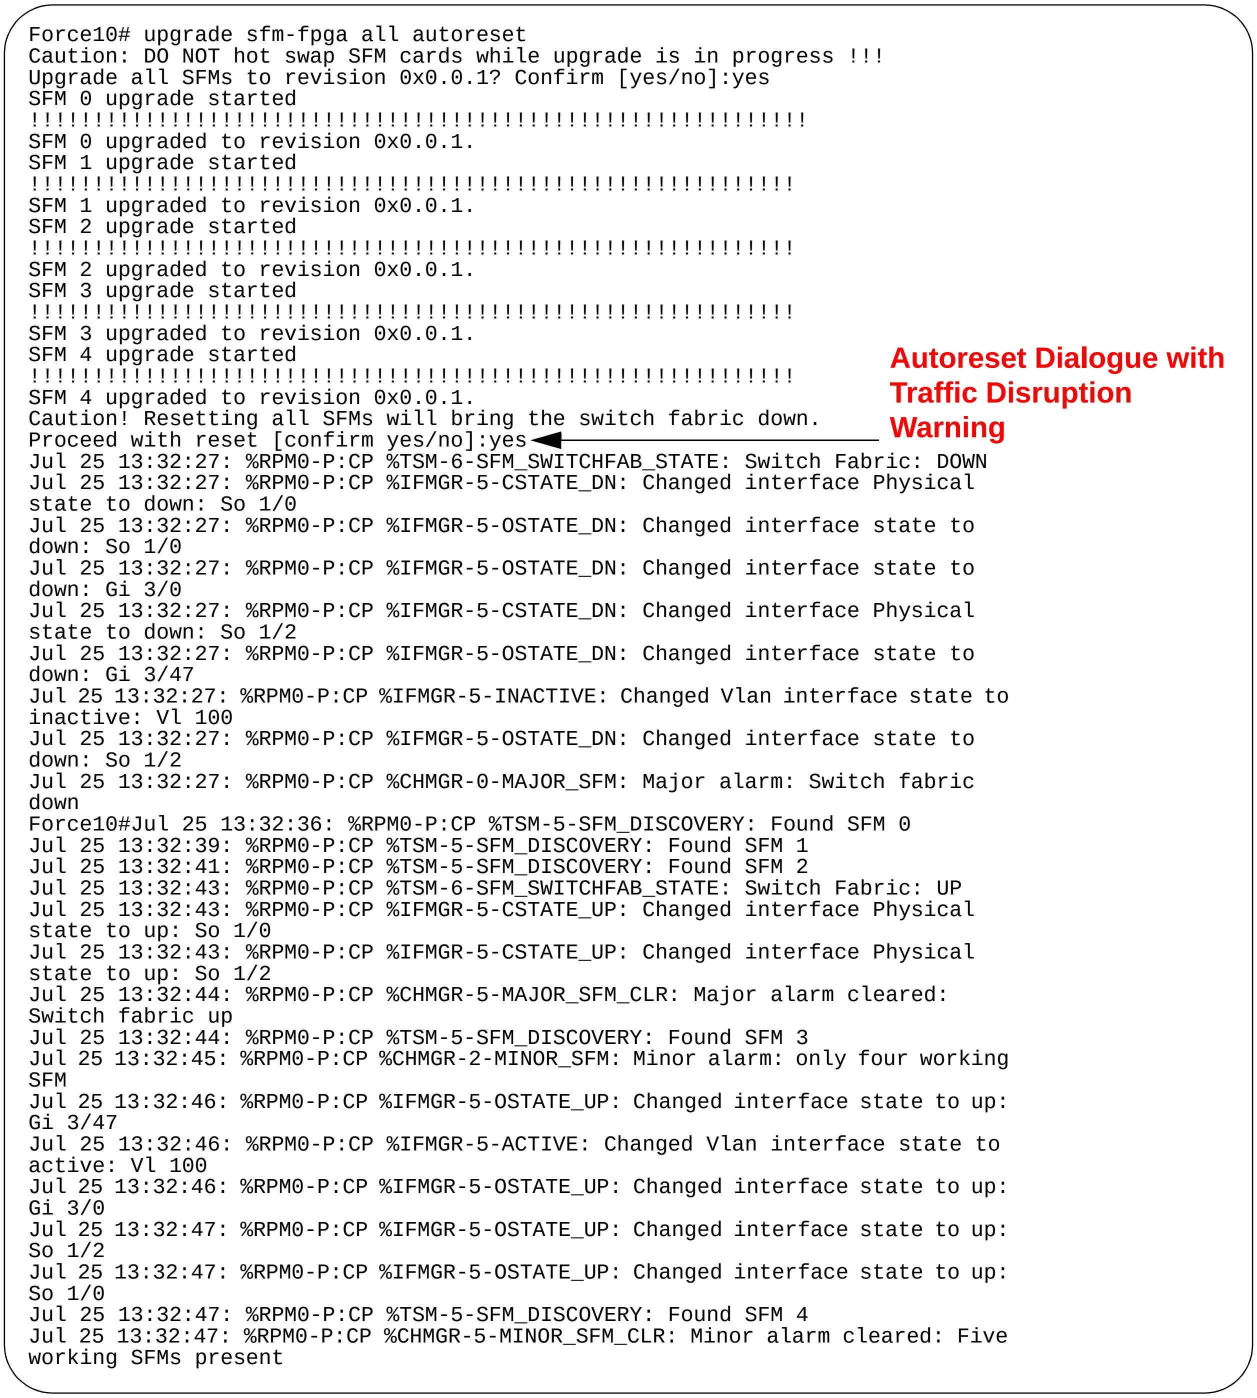

```
Force10# upgrade sfm-fpga all autoreset
Caution: DO NOT hot swap SFM cards while upgrade is in progress !!!
Upgrade all SFMs to revision 0x0.0.1? Confirm [yes/no]:yes
SFM 0 upgrade started
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
SFM 0 upgraded to revision 0x0.0.1.
SFM 1 upgrade started
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
SFM 1 upgraded to revision 0x0.0.1.
SFM 2 upgrade started
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
SFM 2 upgraded to revision 0x0.0.1.
SFM 3 upgrade started
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
SFM 3 upgraded to revision 0x0.0.1.
SFM 4 upgrade started
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
SFM 4 upgraded to revision 0x0.0.1.
Caution! Resetting all SFMs will bring the switch fabric down.
Proceed with reset [confirm yes/no]:yes
Jul 25 13:32:27: %RPM0-P:CP %TSM-6-SFM_SWITCHFAB_STATE: Switch Fabric: DOWN
Jul 25 13:32:27: %RPM0-P:CP %IFMGR-5-CSTATE_DN: Changed interface Physical
state to down: So 1/0
Jul 25 13:32:27: %RPM0-P:CP %IFMGR-5-OSTATE_DN: Changed interface state to
down: So 1/0
Jul 25 13:32:27: %RPM0-P:CP %IFMGR-5-OSTATE_DN: Changed interface state to
down: Gi 3/0
Jul 25 13:32:27: %RPM0-P:CP %IFMGR-5-CSTATE_DN: Changed interface Physical
state to down: So 1/2
Jul 25 13:32:27: %RPM0-P:CP %IFMGR-5-OSTATE_DN: Changed interface state to
down: Gi 3/47
Jul 25 13:32:27: %RPM0-P:CP %IFMGR-5-INACTIVE: Changed Vlan interface state to
inactive: Vl 100
Jul 25 13:32:27: %RPM0-P:CP %IFMGR-5-OSTATE_DN: Changed interface state to
down: So 1/2
Jul 25 13:32:27: %RPM0-P:CP %CHMGR-0-MAJOR_SFM: Major alarm: Switch fabric
down
Force10#Jul 25 13:32:36: %RPM0-P:CP %TSM-5-SFM_DISCOVERY: Found SFM 0
Jul 25 13:32:39: %RPM0-P:CP %TSM-5-SFM_DISCOVERY: Found SFM 1
Jul 25 13:32:41: %RPM0-P:CP %TSM-5-SFM_DISCOVERY: Found SFM 2
Jul 25 13:32:43: %RPM0-P:CP %TSM-6-SFM_SWITCHFAB_STATE: Switch Fabric: UP
Jul 25 13:32:43: %RPM0-P:CP %IFMGR-5-CSTATE_UP: Changed interface Physical
state to up: So 1/0
Jul 25 13:32:43: %RPM0-P:CP %IFMGR-5-CSTATE_UP: Changed interface Physical
state to up: So 1/2
Jul 25 13:32:44: %RPM0-P:CP %CHMGR-5-MAJOR_SFM_CLR: Major alarm cleared:
Switch fabric up
Jul 25 13:32:44: %RPM0-P:CP %TSM-5-SFM_DISCOVERY: Found SFM 3
Jul 25 13:32:45: %RPM0-P:CP %CHMGR-2-MINOR_SFM: Minor alarm: only four working
SFM
Jul 25 13:32:46: %RPM0-P:CP %IFMGR-5-OSTATE_UP: Changed interface state to up:
Gi 3/47
Jul 25 13:32:46: %RPM0-P:CP %IFMGR-5-ACTIVE: Changed Vlan interface state to
active: Vl 100
Jul 25 13:32:46: %RPM0-P:CP %IFMGR-5-OSTATE_UP: Changed interface state to up:
Gi 3/0
Jul 25 13:32:47: %RPM0-P:CP %IFMGR-5-OSTATE_UP: Changed interface state to up:
So 1/2
Jul 25 13:32:47: %RPM0-P:CP %IFMGR-5-OSTATE_UP: Changed interface state to up:
So 1/0
Jul 25 13:32:47: %RPM0-P:CP %TSM-5-SFM_DISCOVERY: Found SFM 4
Jul 25 13:32:47: %RPM0-P:CP %CHMGR-5-MINOR_SFM_CLR: Minor alarm cleared: Five
working SFMs present
```

Autoreset Dialogue with Traffic Disruption Warning

The upgrade process is reflected in the command **show sfm**, as shown in Figure 445.

Figure 445 Viewing the SFM Upgrade Status on E-Series

```
core1#show sfm 0

Switch Fabric State:  up
Switch Mode: SFM3

-- SFM card 0 --
Status         : active     (upgrade in progress)
Card Type      : SFM3 - Switch Fabric Module
Up Time        : 1 min, 7 sec
Last Restart   : user-off
Temperature    : 32C
Power Status   : AC
Serial Number  : 0068148
Part Number    : 7520020001 Rev 03
Vendor Id      : 04
Date Code      : 01402006
Country Code   : 01
FPGA           : 0x0.0.0
Booting from   : EEPROM0
```

Upgrade in Progress

Appendix A IS-IS Metric Styles

C-Series	NO
E-Series	✓

Platform Specific Feature: IS-IS is supported on E-Series only.

The following sections provide additional information on IS-IS Metric Styles:

- IS-IS Metric Styles on page 709
- Configuring Metric Values on page 709

IS-IS Metric Styles

FTOS supports the following IS-IS metric styles:

- narrow (supports only type, length, and value (TLV) up to 63)
- wide (supports TLV up to 16777215)
- transition (supports both narrow and wide and uses a TLV up to 63)
- narrow transition (accepts both narrow and wide and sends only narrow or old-style TLV)
- wide transition (accepts both narrow and wide and sends only wide or new-style TLV)

Configuring Metric Values

The following topics are covered in this section:

- Maximum Values in the Routing Table on page 710
- Changing the IS-IS Metric Style in One Level Only on page 710
- Leaking from One Level to Another on page 712

For any level (Level-1, Level-2, or Level-1-2), the value range possible in the **isis metric** command in the INTERFACE mode changes depending on the metric style.

Note: In the E-Series, the CLI help always states the value range (0-16777215) for the metric style. Refer to Table 64 for the correct value range.

Table 64 Correct Value Range for the isis metric Command

Metric Style	Correct Value Range for the isis metric Command
wide	0 to 16777215
narrow	0 to 63
wide transition	0 to 16777215
narrow transition	0 to 63
transition	0 to 63

Maximum Values in the Routing Table

In the E-Series, the IS-IS metric styles support different cost ranges for the route. The cost range for the narrow metric style is 0 to 1023, while all other metric styles support a range of 0 to 0xFE000000.

Changing the IS-IS Metric Style in One Level Only

By default, the IS-IS metric style is narrow. When you change from one IS-IS metric style to another, the IS-IS metric value (configured with the **isis metric** command) could be affected.

In the following scenarios, the is-type is either Level-1 or Level-2 or Level-1-2 and the metric style changes.

Table 65 Metric Value when Metric Style Changes

Beginning metric style	Final metric style	Resulting IS-IS metric value
wide	narrow	default value (10) if the original value is greater than 63. A message is sent to the console.
wide	transition	truncated value[1] (the truncated value appears in the LSP only.) The original **isis metric** value is displayed in the **show config** and **show running-config** commands and is used if you change back to transition metric style.
wide	narrow transition	default value (10) if the original value is greater than 63. A message is sent to the console.
wide	wide transition	original value
narrow	wide	original value
narrow	transition	original value
narrow	narrow transition	original value
narrow	wide transition	original value
transition	wide	original value
transition	narrow	original value

Table 65 Metric Value when Metric Style Changes

Beginning metric style	Final metric style	Resulting IS-IS metric value
transition	narrow transition	original value
transition	wide transition	original value
narrow transition	wide	original value
narrow transition	narrow	original value
narrow transition	wide transition	original value
narrow transition	transition	original value
wide transition	wide	original value
wide transition	narrow	default value (10) if the original value is greater than 63. A message is sent to the console.
wide transition	narrow transition	default value (10) if the original value is greater than 63. A message is sent to the console.
wide transition	transition	truncated value (the truncated value appears in the LSP only.) The original **isis metric** value is displayed in the **show config** and **show running-config** commands and is used if you change back to transition metric style.

[1] a truncated value is a value that is higher than 63, but set back to 63 because the higher value is not supported.

Moving to transition and then to another metric style produces different results (see Table 66).

Table 66 Metric Value when Metric Style Changes Multiple Times

Beginning metric style	next isis metric style	resulting isis metric value	Next metric style	final isis metric value
wide	transition	truncated value	wide	original value is recovered
wide transition	transition	truncated value	wide transition	original value is recovered
wide	transition	truncated value	narrow	default value (10) A message is sent to the logging buffer
wide transition	transition	truncated value	narrow transition	default value (10) A message is sent to the logging buffer

Leaking from One Level to Another

In the following scenarios, each IS-IS level is configured with a different metric style.

Table 67 Metric Value with Different Levels Configured with Different Metric Styles

Level-1 metric style	Level-2 metric style	Resulting isis metric value
narrow	wide	original value
narrow	wide transition	original value
narrow	narrow transition	original value
narrow	transition	original value
wide	narrow	truncated value
wide	narrow transition	truncated value
wide	wide transition	original value
wide	transition	truncated value
narrow transition	wide	original value
narrow transition	narrow	original value
narrow transition	wide transition	original value
narrow transition	transition	original value
transition	wide	original value
transition	narrow	original value
transition	wide transition	original value
transition	narrow transition	original value
wide transition	wide	original value
wide transition	narrow	truncated value
wide transition	narrow transition	truncated value
wide transition	transition	truncated value

Appendix B Configuring MTU Size

The E-Series supports a link Maximum Transmission Unit (MTU) of 9252 bytes and maximum IP MTU of 9234 bytes. The link MTU is the frame size of a packet, and the IP MTU size is used for IP fragmentation. If the system determines that the IP packet must be fragmented as it leaves the interface, FTOS divides the packet into fragments no bigger than the size set in the **ip mtu** command.

In FTOS:

MTU = Entire Ethernet packet (Ethernet header + FCS + payload)

Since different networking vendors define MTU differently, check their documentation when planing MTU sizes across a network.

Table 68 lists the range for each transmission media.

Table 68 MTU Range

Transmission Media	MTU Range (in bytes)
Ethernet	594-9252 = link MTU 576-9234 = IP MTU

Configuring MTU Size on an Interface

You must compensate for Layer-2 header when configuring IP MTU. If the packet includes a Layer-2 header, the difference between the link MTU and IP MTU must be enough bytes to include for the Layer-2 header. For example, for VLAN packets, if the IP MTU is 1400, the Link MTU must be no less than 1422:

1400 IP MTU + 22 VLAN Tag = 1422 bytes Link MTU

Table 69 lists the various Layer 2 overheads found in FTOS and the number of bytes.

Table 69 Difference between Link MTU and IP MTU

Layer-2 Overhead	Difference between Link MTU and IP MTU
Ethernet (untagged)	18 bytes
VLAN Tag	22 bytes
Untagged Packet with VLAN-Stack Header	22 bytes
Tagged Packet with VLAN-Stack Header	26 bytes

Link MTU and IP MTU considerations for port channels and VLANs are as follows.

Port Channels:

- All members must have the same link MTU value and the same IP MTU value.
- The port channel link MTU and IP MTU must be less than or equal to the link MTU and IP MTU values configured on the channel members.

Example: if the members have a link MTU of 2100 and an IP MTU 2000, the port channel's MTU values cannot be higher than 2100 for link MTU or 2000 bytes for IP MTU.

VLANs:

- All members of a VLAN must have same IP MTU value.
- Members can have different Link MTU values. Tagged members must have a link MTU 4 bytes higher than untagged members to account for the packet tag.
- The VLAN link MTU and IP MTU must be less than or equal to the link MTU and IP MTU values configured on the VLAN members.

Example: The VLAN contains tagged members with Link MTU of 1522 and IP MTU of 1500 and untagged members with Link MTU of 1518 and IP MTU of 1500. The VLAN's Link MTU cannot be higher than 1518 bytes and its IP MTU cannot be higher than 1500 bytes.

Appendix C SNMP Traps

SNMP

SNMP is used to communicate management information between the network management stations and the agents in the network elements. FTOS supports SNMP versions 1, 2c, and 3, for both read-only and read-write modes. In addition, FTOS supports SNMP traps, which are messages informing the SNMP manager about the network, and up to 16 SNMP trap receivers.

Table 70 contains SNMP traps and OIDs that provide information on hardware components.

Table 70 SNMP Traps and OIDs

OID String	OID Name	Description
RPM		
.1.3.6.1.4.1.6027.3.1.1.3.8	chRPMMajorAlarmStatus	Fault status of the major alarm LED on the RPM.
.1.3.6.1.4.1.6027.3.1.1.3.9	chRPMMinorAlarmStatus	Fault status of the minor alarm LED on the RPM.
.1.3.6.1.4.1.6027.3.1.1.4.0.11	chAlarmRpmUp	Trap generated when the status of primary or secondary RPM changes to up and running.
.1.3.6.1.4.1.6027.3.1.1.4.0.12	chAlarmRpmDown	Trap generated when the status of primary or secondary RPM changes to down, either by software reset or by being physically removed from the chassis.
Line Card		
.1.3.6.1.4.1.6027.3.1.1.2.3.1.15	chSysCardOperStatus	Operational status of a line card. If the chSysCardAdminStatus is “up”, the valid state is “ready”, that is, the card is present and ready, and the chSysCardOperStatus status is “up.” If the chSysCardAdminStatus is “down”, the service states can be: • **offline**—the card is not used. • **cardNotmatch**—the card does not matched what is configured. • **cardProblem**—a hardware problem is detected on the card. • **diagMode**—the card is in diagnostic mode.
.1.3.6.1.4.1.6027.3.1.1.2.3.1.16	chSysCardFaultStatus	Fault status of the card: • **on**—the system fault light is on • **off**—the system fault light is off
.1.3.6.1.4.1.6027.3.1.1.4.0.1	chAlarmCardDown	Trap reported when the card operational status changes to down.
.1.3.6.1.4.1.6027.3.1.1.4.0.2	chAlarmCardUp	Trap reported when the card operational status changes to up.
.1.3.6.1.4.1.6027.3.1.1.4.0.3	chAlarmCardReset	Trap reported when a card is reset.
.1.3.6.1.4.1.6027.3.1.1.4.0.7	chAlarmCardProblem	Trap reported when the card operational status changes to "card problem."
SFM		

Table 70 SNMP Traps and OIDs

OID String	OID Name	Description
.1.3.6.1.4.1.6027.3.1.1.2.8.1.7	chSysSfmOperStatus	Operational status of an SFM. If the chSysCardOperStatus is "down", the service states can be: • **absent**—the card is not present. • **standby**—the card is in standby • mode.
.1.3.6.1.4.1.6027.3.1.1.2.8.1.8	chSysSfmErrorStatus	Provides further information on the status of an SFM. If the chSysSfmAdminStatus is up, the valid state is "ok", that is, the card is present and ready. If the chSysSfmAdminStatus is "down", the error status can be: • **not-available**—status not available. • **error**—the card is in error state.
.1.3.6.1.4.1.6027.3.1.1.4.0.9	chAlarmSfmUp	Trap reported when SFM operational status changes to up.
.1.3.6.1.4.1.6027.3.1.1.4.0.10	chAlarmSfmDown	Trap reported when SFM operational status changes to down.
.1.3.6.1.4.1.6027.3.1.1.4.0.22	chAlarmMajorSfmDown	Trap reported when the majority of SFMs are down.
.1.3.6.1.4.1.6027.3.1.1.4.0.24	chAlarmMinorSfmDown	Trap reported when several SFMs are down.
Power Supply		
.1.3.6.1.4.1.6027.3.1.1.2.1.1.2	chSysPowerSupplyOperStatus	Each entry in the chSysPowerSupplyTable includes a set of objects which describe the status of a particular power supply.
.1.3.6.1.4.1.6027.3.1.1.4.0.13	chAlarmPowerSupplyDown	Trap generated when the power supply status changes to non-operational.
.1.3.6.1.4.1.6027.3.1.1.4.0.17	chAlarmPowerSupplyClear	Trap generated when the power supply status changes to operational.
.1.3.6.1.4.1.6027.3.1.1.4.0.32	chAlarmMajorPS	Trap generated when a power supply major alarm is issued.
.1.3.6.1.4.1.6027.3.1.1.4.0.33	chAlarmMajorPSClr	Trap generated when a power supply major alarm is cleared.
.1.3.6.1.4.1.6027.3.1.1.4.0.34	chAlarmMinorPS	Trap generated when a power supply minor alarm is issued.
.1.3.6.1.4.1.6027.3.1.1.4.0.35	chAlarmMinorPSClr	Trap generated when a power supply minor alarm is cleared.

Table 70 SNMP Traps and OIDs

OID String	OID Name	Description
Fan		
.1.3.6.1.4.1.6027.3.1.1.2.2.1.2	chSysFanTrayOperStatus	Each entry in the chSysFanTrayTable includes a set of objects which describe the status of a particular fan tray, as identified by the chSysFanTrayIndex.
.1.3.6.1.4.1.6027.3.1.1.4.0.36	chAlarmMinorFanBad	Trap generated on the E600 or E1200 when the status of one or more fans changes to down and generates a minor alarm. An E300 does not generate a minor alarm as the system monitors only the overall status of the fan tray.
.1.3.6.1.4.1.6027.3.1.1.4.0.21	chAlarmMinorFanBadClear	Trap generated when a minor alarm on one or more fans is cleared.
.1.3.6.1.4.1.6027.3.1.1.4.0.16	chAlarmFanTrayDown	Trap generated when all fans are down and/or when the fan tray status changes to missing or down.
.1.3.6.1.4.1.6027.3.1.1.4.0.20	chAlarmFanTrayClear	Trap generated when the fan tray status changes to operational.

Table 446 lists the traps sent by FTOS. Each trap is listed by Message ID, Trap Type,Trap Option, and followed by the error message(s) associated with the trap.

Figure 446 SNMP Traps and Error Messages

Message ID	Trap Type	Trap Option
COLD_START	**SNMP**	**COLDSTART**
%SNMP-5-SNMP_COLD_START: SNMP COLD_START trap sent.		
COPY_CONFIG_COMPLETE	**SNMP**	**NONE**
SNMP Copy Config Command Completed		
LINK_DOWN	**SNMP**	**LINKDOWN**
%IFA-1-PORT_LINKDN: changed interface state to down:%d		
LINK_UP	**SNMP**	**LINKUP**
%IFA-1-PORT_LINKUP: changed interface state to up:%d		
AUTHENTICATION_FAIL	**SNMP**	**AUTH**

Figure 446 SNMP Traps and Error Messages (continued)

Message ID	Trap Type	Trap Option
%SNMP-3-SNMP_AUTH_FAIL: SNMP Authentication failed.Request with invalid community string.		
RMON_RISING_THRESHOLD	**SNMP**	**NONE**
%RPM0-P:CP %SNMP-4-RMON_RISING_THRESHOLD: RMON rising threshold alarm from SNMP OID <oid>		
RMON_FALLING_THRESHOLD	**SNMP**	**NONE**
%RPM0-P:CP %SNMP-4-RMON_FALLING_THRESHOLD: RMON falling threshold alarm from SNMP OID <oid>		
RMON_HC_RISHING_THRESHOLD	**SNMP**	**NONE**
%RPM0-P:CP %SNMP-4-RMON_HC_RISING_THRESHOLD: RMON high-capacity rising threshold alarm from SNMP OID <oid>		
RMON_HC_FALLING_THRESHOLD	**SNMP**	**NONE**
%RPM0-P:CP %SNMP-4-RMON_HC_FALLING_THRESHOLD: RMON high-capacity falling threshold alarm from SNMP OID <oid>		
RESV	**NONE**	**NONE**
N/A		
CHM_CARD_DOWN	**ENVMON**	**NONE**
%CHMGR-1-CARD_SHUTDOWN: %sLine card %d down - %s %CHMGR-2-CARD_DOWN: %sLine card %d down - %s		
CHM_CARD_UP	**ENVMON**	**NONE**
%CHMGR-5-LINECARDUP: %sLine card %d is up		
CHM_CARD_MISMATCH	**ENVMON**	**NONE**
%CHMGR-3-CARD_MISMATCH: Mismatch: line card %d is type %s - type %s required.		
CHM_RPM_UP	**ENVMON**	**NONE**
%RAM-6-RPM_STATE: RPM1 is in Active State %RAM-6-RPM_STATE: RPM0 is in Standby State		
CHM_RPM_DOWN	**ENVMON**	**NONE**
%CHMGR-2-RPM_DOWN: RPM 0 down - hard reset %CHMGR-2-RPM_DOWN: RPM 0 down - card removed		
CHM_RPM_PRIMARY	**ENVMON**	**NONE**
%RAM-5-COLD_FAILOVER: RPM Failover Completed %RAM-5-HOT_FAILOVER: RPM Failover Completed %RAM-5-FAST_FAILOVER: RPM Failover Completed		

Figure 446 SNMP Traps and Error Messages (continued)

Message ID	Trap Type	Trap Option
CHM_SFM_ADD	**ENVMON**	**NONE**
%TSM-5-SFM_DISCOVERY: Found SFM 1		
CHM_SFM_REMOVE	**ENVMON**	**NONE**
%TSM-5-SFM_REMOVE: Removed SFM 1		
CHM_MAJ_SFM_DOWN	**ENVMON**	**NONE**
%CHMGR-0-MAJOR_SFM: Major alarm: Switch fabric down		
CHM_MAJ_SFM_DOWN_CLR	**ENVMON**	**NONE**
%CHMGR-5-MAJOR_SFM_CLR: Major alarm cleared: Switch fabric up		
CHM_MIN_SFM_DOWN	**ENVMON**	**NONE**
%CHMGR-2-MINOR_SFM: MInor alarm: No working standby SFM		

Figure 446 SNMP Traps and Error Messages (continued)

Message ID	Trap Type	Trap Option
CHM_MIN_SFM_DOWN_CLR	**ENVMON**	**NONE**
%CHMGR-5-MINOR_SFM_CLR: Minor alarm cleared: Working standby SFM present		
CHM_PWRSRC_DOWN	**ENVMON**	**SUPPLY**
%CHMGR-2-PEM_PRBLM: Major alarm: problem with power entry module %s		
CHM_PWRSRC_CLR	**ENVMON**	**SUPPLY**
%CHMGR-5-PEM_OK: Major alarm cleared: power entry module %s is good		
CHM_MAJ_ALARM_PS	**ENVMON**	**SUPPLY**
%CHMGR-0-MAJOR_PS: Major alarm: insufficient power %s		
CHM_MAJ_ALARM_PS_CLR	**ENVMON**	**SUPPLY**
%CHMGR-5-MAJOR_PS_CLR: major alarm cleared: sufficient power		
CHM_MIN_ALARM_PS	**ENVMON**	**SUPPLY**
%CHMGR-1-MINOR_PS: Minor alarm: power supply non-redundant		
CHM_MIN_ALARM_PS_CLR	**ENVMON**	**SUPPLY**
%CHMGR-5-MINOR_PS_CLR: Minor alarm cleared: power supply redundant		
CHM_MIN_ALRM_TEMP	**ENVMON**	**TEMP**
%CHMGR-2-MINOR_TEMP: Minor alarm: chassis temperature		
CHM_MIN_ALRM_TEMP_CLR	**ENVMON**	**TEMP**
%CHMRG-5-MINOR_TEMP_CLR: Minor alarm cleared: chassis temperature normal (%s %d temperature is within threshold of %dC)		
CHM_MAJ_ALRM_TEMP	**ENVMON**	**TEMP**
%CHMGR-2-MAJOR_TEMP: Major alarm: chassis temperature high (%s temperature reaches or exceeds threshold of %dC)		
CHM_MAJ_ALRM_TEMP_CLR	**ENVMON**	**TEMP**
%CHMGR-2-MAJOR_TEMP_CLR: Major alarm cleared: chassis temperature lower (%s %d temperature is within threshold of %dC)		
CHM_FANTRAY_BAD	**ENVMON**	**FAN**
For E1200: %CHMGR-2-FAN_TRAY_BAD: Major alarm: fantray %d is missing or down %CHMGR-2-ALL_FAN_BAD: Major alarm: all fans in fan tray %d are down. For E600 and E300: %CHMGR-2-FANTRAYBAD: Major alarm: fan tray is missing %CHMGR-2-FANSBAD: Major alarm: most or all fans in fan tray are down		

Figure 446 SNMP Traps and Error Messages (continued)

Message ID	Trap Type	Trap Option
CHM_FANTRAY_BAD_CLR	**ENVMON**	**FAN**
For the E1200: %CHMGR-5-FAN_TRAY_OK: Major alarm cleared: fan tray %d present For the E600 and E300: %CHMGR-5-FANTRAYOK: Major alarm cleared: fan tray present		
CHM_MIN_FANBAD	**ENVMON**	**FAN**
For the E1200: %CHMGR-2-FAN_BAD: Minor alarm: some fans in fan tray %d are down For the E600 and E300: %CHMGR- 2-1FANBAD: Minor alarm: fan in fan tray is down		
CHM_MIN_FANBAD_CLR	**ENVMON**	**FAN**
For E1200: %CHMGR-2-FAN_OK: Minor alarm cleared: all fans in fan tray %d are good For E600 and E300: %CHMGR-5-FANOK: Minor alarm cleared: all fans in fan tray are good		
TME_TASK_SUSPEND	**ENVMON**	**NONE**
%TME-2-TASK SUSPENDED: SUSPENDED - svce:%d - inst:%d - task:%s		
TME_TASK_TERM	**ENVMON**	**NONE**
%TME-2-ABNORMAL_TASK_TERMINATION: CRASH - task:%s %s		
CHM_CPU_THRESHOLD	**ENVMON**	**NONE**
%CHMGR-5-CPU_THRESHOLD: Cpu %s usage above threshold. Cpu5SecUsage (%d)		
CHM_CPU_THRESHOLD_CLR	**ENVMON**	**NONE**
%CHMGR-5-CPU_THRESHOLD_CLR: Cpu %s usage drops below threshold. Cpu5SecUsage (%d)		
CHM_MEM_THRESHOLD	**ENVMON**	**NONE**
%CHMGR-5-MEM_THRESHOLD: Memory %s usage above threshold. MemUsage (%d)		
CHM_MEM_THRESHOLD_CLR	**ENVMON**	**NONE**
%CHMGR-5-MEM_THRESHOLD_CLR: Memory %s usage drops below threshold. MemUsage (%d)		
MACMGR_STN_MOVE	**ENVMON**	**NONE**
%MACMGR-5-DETECT_STN_MOVE: Station Move threshold exceeded for Mac %s in vlan %d		
VRRP_BADAUTH	**PROTO**	**NONE**
%RPM1-P:RP2 %VRRP-3-VRRP_BAD_AUTH: vrid-1 on Gi 11/12 rcvd pkt with authentication type mismatch. %RPM1-P:RP2 %VRRP-3-VRRP_BAD_AUTH: vrid-1 on Gi 11/12 rcvd pkt with authentication failure.		
VRRP_GO_MASTER	**PROTO**	**NONE**
%VRRP-6-VRRP_MASTER: vrid-%d on %s entering MASTER		

Figure 446 SNMP Traps and Error Messages (continued)

Message ID	Trap Type	Trap Option
BGP4_ESTABLISHED	**PROTO**	**NONE**
%TRAP-5-PEER_ESTABLISHED: Neighbor %a, state %s		
BGP4_BACKW_XSITION	**PROTO**	**NONE**
%TRAP-5-BACKWARD_STATE_TRANS: Neighbor %a, state %s		

Appendix D RFCs and MIBs

This appendix contains the following sections:

IEEE Compliance

- 802.1AB—LLDP
- 802.3ae—10 Gigabit Ethernet
- 802.3ab—1000Base-T
- 802.1p/Q—VLAN Tagging
- 802.1s—Multiple Spanning Tree Protocol
- 802.1w—Rapid Spanning Tree Protocol
- 802.3ad—Link Aggregation
- 802.1D—Bridging
- 802.1X—Port Authentication
- 802.3x—Flow Control

RFC Compliance

The following is a list of the RFCs supported by FTOS, listed by related protocol:

BGP

- RFC 1657—Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIv2
- RFC 1771—A Border Gateway Protocol 4 (BGP-4), ID draft-ietf-idr-bgp4-20.txt (revision to BGPv4)
- RFC 1772—Application of the Border Gateway Protocol in the Internet
- RFC 1997—BGP Communities Attribute

- RFC 1998—An Application of the BGP Community Attribute in Multi-home Routing
- RFC 2385—Protection of BGP Sessions via the TCP MD5 Signature Option, MD5 encryption ID draft-ietf-idr-restart-06.txt (BGP Graceful Restart) ID draft-ietf-idr-bgp4-mib-05.txt (BGP MIB)
- RFC 2439—BGP Route Flap Dampening
- RFC 2519—A Framework for Inter-Domain Route Aggregation
- RFC 2796—BGP Route Reflection—An Alternative to Full Mesh IBGP
- RFC 2842—Capabilities Advertisement with BGP-4
- RFC 2858—Multi-protocol Extensions for BGP-4
- RFC 2918—Route Refresh Capability for BGP-4
- RFC 3065—Autonomous System Confederations for BGP
- ietf draft—Graceful BGP restart

General Routing Protocols

- RFC 768—UDP
- RFC 783—TFTP
- RFC 791—IP
- RFC 792—ICMP
- RFC 793—TCP
- RFC 826—ARP
- RFC 854—Telnet
- RFC 959—FTP
- RFC 1027—Proxy ARP
- RFC 1157—SNMP v1/v2/v3
- RFC 1213—Management Information Base for Network Management of TCP/IP-based internets: MIB-II [except the EGP Group]
- RFC 1215—A Convention for Defining Traps for use with the SNMP
- RFC 1305—NTP v3
- RFC 1493—Definitions of Managed Objects for Bridges [except for the dot1dTpLearnedEntryDiscards object]
- RFC 1542—BootP (relay)
- RFC 1573—Interfaces Group MIB
- RFC 1591—DNS client
- RFC 1619—PPP over SONET
- RFC 1812—IP v4 routers
- RFC 2096—IP Forwarding Table MIB
- RFC 2131—BootP/DHCP helper
- RFC 2236—IGMP v1 and v2
- RFC 2338 and RFC 3768—VRRP
- RFC 2270—Using a Dedicated AS for Sites Homed to a Single Provider
- RFC 2571—An Architecture for Describing SNMP Management Frameworks

- RFC 2572—Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
- RFC 2574—User-based Security Model (USM) for Version 3 of the Simple Network Management Protocol (SNMPv3)
- RFC 2575—View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)
- RFC 2576—Co-existence between Version 1, Version 2, and Version 3 of the Internet-Standard Network Management Framework
- RFC 2665—Definitions of Managed Objects for the Ethernet-like Interface Types
- RFC 2674—The Q-BRIDGE of Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering, and Virtual LAN Extensions
- RFC 2787—Definitions of Managed Objects for the Virtual Router Redundancy Protocol (VRRP)
- draft-ietf-bfd-base-05—Bidirectional Forwarding Detection

IP Multicast

- RFC 1112—IGMP
- RFC 2236—IGMP v2
- RFC 2858—Multi-protocol Extensions for BGP4 (MBGP)
- RFC 3618—Multicast Source Discovery Protocol (MSDP)
- RFC 3810 - Multicast Listener Discovery Version 2 (MLDv2) for IPv6
- RFC 4541 - Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches
- ietf-drafts:
    - PIM — SM v2
    - PIM BSR
    - IGMP Snooping

IS-IS

- RFC 1142—Intra-domain Routing Protocol
- RFC 1195—TCP
- RFC 2763—Dynamic Hostname Exchange
- RFC 2966—Domain-wide Prefixes
- RFC 3373—Three-Way Handshake
- Routing IPv6 with IS-IS
    - draft-ietf-isis-ipv6-06
- ietf-drafts IPv4:
    - Point-to-point operation over LAN
    - Cryptographic Authentication
    - Maintaining more than 255 circuits in IS-IS
    - Extended Ethernet Frame Size support
    - Extensions for Traffic Engineering (wide metrics)

OSPF

- RFC 1587—NSSA option
- RFC 1850—OSPF Version 2 Management Information Base
- RFC 2154—OSPF MD5
- RFC 2328—OSPF v2
- RFC 2370—Opaque LSA option
- RFC 3623—Graceful OSPF Restart

RIP

- RFC 1058—RIP v1
- RFC 1724—RIP Version 2 MIB Extension
- RFC 2453—RIP v2

RMON

- RFC 1757—RMON
- RFC 2819—Remote Network Monitoring MIB: Ethernet Statistics Table, Ethernet History Control Table, Ethernet History Table, Alarm Table, Event Table, Log Table
- RFC 3273—Remote Network Monitoring MIB for High-Capacity Networks (64 bits): Ethernet Statistics High-Capacity Table, Ethernet History High-Capacity Table
- RFC 3434—Remote Monitoring MIB Extensions for High-Capacity Alarms, High-Capacity Alarm Table (64 bits)

Security

- RFC 1492—TACACS+
- RFC 2865—RADIUS
- RFC 3128—Protection Against a Variant of the Tiny Fragment Attack
- RFC 3580—Remote Authentication Dial In User Service (RADIUS)
- Secure Copy (SCP)
- SSH v1/v2

SONET

- RFC 2558—Definitions of Managed Objects for the SONET/SDH Interface Type
- RFC 2615—PPP-over-SONET/SDH

MIBs

The following is a list of the Management Information Bases (MIBs) supported by FTOS:

- Management Information Base for Network Management of TCP/IP-based internets: MIB-II [except the EGP Group] (RFC 1213)
- A Convention for Defining Traps for use with the SNMP (RFC 1215)
- Definitions of Managed Objects for Bridges [except for the dot1dTpLearnedEntryDiscards object] (RFC 1493)
- Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIv2 (RFC 1657)
- RIP Version 2 MIB Extension (RFC 1724)
- OSPF Version 2 Management Information Base (RFC 1850)
- Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2) (RFC 1907)
- SNMPv2 Management Information Base for the Internet Protocol using SMIv2 (RFC 2011)
- SNMPv2 Management Information Base for the Transmission Control Protocol using SMIv2 (RFC 2012)
- SNMPv2 Management Information Base for the User Datagram Protocol using SMIv2 (RFC 2013)
- IP Forwarding Table MIB (RFC 2096)
- Definitions of Managed Objects for the SONET/SDH Interface Type (RFC 2558)
- Definitions of Managed Objects for the Ethernet-like Interface Types (RFC 2665)
- An Architecture for Describing SNMP Management Frameworks (RFC 2571)
- Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) (RFC 2572)
- User-based Security Model (USM) for Version 3 of the Simple Network Management Protocol (SNMPv3) (RFC 2574)
- View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) (RFC 2575)
- Co-existence between Version 1, Version 2, and Version 3 of the Internet-Standard Network Management Framework (RFC 2576)
- RADIUS Authentication Client (RFC 2618), except the following four counters:
  - radiusAuthClientInvalidServerAddresses
  - radiusAuthClientMalformedAccessResponses
  - radiusAuthClientUnknownTypes
  - radiusAuthClientPacketsDropped
- The Q-BRIDGE of Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering, and Virtual LAN Extensions (RFC 2674)
- Definitions of Managed Objects for the Virtual Router Redundancy Protocol (RFC 2787)
- Remote Network Monitoring MIB: Ethernet Statistics Table, Ethernet History Control Table, Ethernet History Table, Alarm Table, Event Table, Log Table (RFC 2819)
- Interfaces Group MIB (RFC 2863)

- Remote Network Monitoring MIB for High-Capacity Networks (64 bits): Ethernet Statistics High-Capacity Table, Ethernet History High-Capacity Table (RFC 3273)
- Remote Monitoring MIB Extensions for High-Capacity Alarms, High-Capacity Alarm Table (64 bits) (RFC 3434)

Force10 Specific MIBs

- Force10 Enterprise Link Aggregation MIB (LAGs)
- Force10 Chassis MIB
- Force10-COPY-CONFIG-MIB—supporting SNMP SET requests (SNMP Copy Config command completed)
- Force10 Fault Management MIB—alarms and status reporting
- Force10 Monitor MIB
- Force10 Product MIB—(E1200, E600, E600i, E300) defines system Object Identifier values.
- Force10 SMI MIB
- Force10 System Component MIB—allows the user to view CAM usage information.
- Force10 Textual Convention MIB
- Internet Draft—Management Information Base for IS-IS (draft-ietf-isis-wg-mib-16.txt):

  isisSysObject (top level scalar objects)

  isisISAdjTable

  isisISAdjAreaAddrTable

  isisISAdjIPAddrTable

  isisISAdjProtSuppTable
- Force10 BGP MIB—Force10-BGP4-V2-MIB

MIB Location

Force10 MIBs are under the **Force10 MIBs** subhead on the **Documentation** page of iSupport:

https://www.force10networks.com/csportal20/KnowledgeBase/Documentation.aspx

You also can obtain a list of selected MIBs and their OIDs at the following URL:

https://www.force10networks.com/csportal20/MIBs/MIB_OIDs.aspx

Some pages of iSupport require a login. To request an iSupport account, go to:

https://www.force10networks.com/CSPortal20/Support/AccountRequest.aspx

If you have forgotten or lost your account information, contace Force10 TAC for assistance.

Appendix E Managing VLANs using SNMP Set in Q-Bridge MIB

The Q-Bridge MIB Set allows you to use the SNMP Set feature to maintain VLANs in an E-Series chassis. You can:

- Assigning access and configuring Layer 2 ports on page 733
- Add tagged ports using SNMP Set on page 735
- Changing Administrative Status with SNMP Set on page 739

To manage a VLAN using the VLAN Static table in the Q-Bridge MIB, use the following procedures.

Assigning access and configuring Layer 2 ports

Before using the SNMP Set to manage the FTOS VLANs, use the FTOS commands below to ensure all interfaces are switch port..

Step	Task
1	Execute the FTOS command to assign read/write access from the global CONFIGURATION mode: `Force10(conf)# snmp-server community public rw`
2	Set the port range for your Layer 2 ports: `Force10(conf)# interface range gi 0/0 - 48`
3	Set the ports to Layer 2 ports: `Force10(conf-if-gi-0/0- 48)# switchport`

Example of SNMP Set

The example below creates vlan 10.

```
abcServer > snmpset -c public 123.45.6.78 .1.3.6.1.2.1.17.7.1.4.3.1.5.10 i 4
SNMPv2-SMI::mib-2.17.7.1.4.3.1.5.10 = INTEGER: 4
```

Figure 447 displays the **show vlan** output.

```
Force10# show vlan
Codes: * - Default VLAN, G - GVRP VLANs
    NUM    Status    Q Ports
*   1      Inactive
    10     Inactive
Force10#
```

Figure 447 show vlan command

SNMP output

```
abcServer > snmpwalk -c public 123.34.5.67 .1.3.6.1.2.1.17.7.1.4.3.1.1
SNMPv2-SMI::mib-2.17.7.1.4.3.1.1.1107787777 = STRING: "Vlan 1"
SNMPv2-SMI::mib-2.17.7.1.4.3.1.1.1107787786 = STRING: "Vlan 10"
```

Assign line card ports

Use the FTOS **interface range** command to assign ports in line card 0, 1, and 2 as Layer 2 ports.

```
Force10(conf)#inteface range gi 0/0 - 23
Force10(conf-if-range-gi-0/0-23)#switchport
Force10(conf-if-range-gi-0/0-23)#interface range gi 1/0 - 23
Force10(conf-if-range-gi-1/0-23)#switchport
Force10(conf-if-range-gi-1/0-23)#interface range gi 2/0 - 23
Force10(conf-if-range-gi-2/0-23)#switchport
Force10(conf-if-range-gi-2/0-23)#end
Mar 17 11:14:01: %RPM0-P:CP %SYS-5-CONFIG_I: Configured from console by
console  - repeated 2 times
Force10#show vlan
Codes: * - Default VLAN, G - GVRP VLANs
    NUM    Status    Q Ports
*   1      Inactive  U Gi 0/0-23
                     U Gi 1/0-23
                     U Gi 2/0-23

    10     Inactive

Force10#
```

Figure 448 interface range command

Query q-bride MIB object, dot1qVlanStaticEgressPorts

In the SNMP query below, all the Layer 2 ports in line card 0, 1, and 2 belonged to the default VLAN, VLAN 1. The .1107787777 is for VLAN 10 with no ports assigned.

```
abcServer > snmpwalk -c public 172.16.1.43 .1.3.6.1.2.1.17.7.1.4.3.1.2
SNMPv2-SMI::mib-2.17.7.1.4.3.1.2.1107787777 = Hex-STRING:

FF FF FF 00 00 00 00 00 00 00 00 00 FF FF FF 00
00 00 00 00 00 00 00 00 FF FF FF 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 01 08 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
SNMPv2-SMI::mib-2.17.7.1.4.3.1.2.1107787786 = Hex-STRING:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
```

Add tagged ports using SNMP Set

In this example, port 1-24 (line card 0) are being assigned to VLAN 10 as tagged ports.

Note: You must always specify dot1qVlanStaticEgressPorts and dot1qVlanStaticUntaggedPorts in a single SNMP Set operation. The SNMP Set operation fails if only one OIDC is specified.

```
Snmpset -c public 123.45.6.78 .1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786 x "FF FF FF 00 00
00 00 00 00 00 00 00 FF FF FF 00 00 00 00 00 00 00 00 00 FF FF FF"
.1.3.6.1.2.1.17.7.1.4.3.1.4.1107787786 x "00"
```

Figure 449 displays the **show vlan** output of the tagged ports.

```
Force10# show vlan
Codes: * - Default VLAN, G - GVRP VLANs
    NUM    Status    Q Ports
*   1      Inactive
    10     Inactive T Gi 0/0-23
                    T Gi 1/0-23
                    T Gi 2/0-23
Force10#
```

Figure 449 show vlan command with tagged ports

The output from SNMP query of VLAN 10, dot1qVlanStaticEgressPorts and dot1qVlanStaticUntaggedPorts is shown below.

```
abcServer > snmpwalk -c public 123.45.6.78 .1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786
SNMPv2-SMI::mib-2.17.7.1.4.3.1.2.1107787786 = Hex-STRING:
FF FF FF 00 00 00 00 00 00 00 00 00 FF FF FF 00
00 00 00 00 00 00 00 00 FF FF FF 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
abcServer > snmpwalk -c public 123.45.6.78 .1.3.6.1.2.1.17.7.1.4.3.1.4.1107787786
SNMPv2-SMI::mib-2.17.7.1.4.3.1.4.1107787786 = Hex-STRING:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
```

To assign ports 1-23 (line card 0) on VLAN 10 as untagged ports, do the following.

Step **Task**

1 Remove ports 1-23 from VLAN 10 (currently tagged ports) by changing dot1qVlanStaticEgressPorts, port 1-23 as zero.

```
snmpset -c public 172.16.1.43 .1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786 x "00
00 00 00 00 00 00 00 00 00 00 00 FF FF FF 00 00 00 00 00 00 00 00 FF FF
FF" .1.3.6.1.2.1.17.7.1.4.3.1.4.1107787786 x "00"
```

Note: Remember to maintain the correct dot1qVlanStaticUntaggedPorts which is currently all zeros as tagged ports.

2 The result of SNMP query for dot1qVlanStaticEgressPorts and dot1qVlanStaticEgressPorts on VLAN 10 is:

```
abcServer > snmpwalk -c public 123.45.6.78
.1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786

SNMPv2-SMI::mib-2.17.7.1.4.3.1.2.1107787786 = Hex-STRING:
00 00 00 00 00 00 00 00 00 00 00 00 FF FF FF 00? Port 1-23 is removed
00 00 00 00 00 00 00 00 FF FF FF 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
abcServer > snmpwalk -c public 123.45.6.78
.1.3.6.1.2.1.17.7.1.4.3.1.4.1107787786

SNMPv2-SMI::mib-2.17.7.1.4.3.1.4.1107787786 = Hex-STRING:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
```

3 Execute the FTOS **show vlan** command:

```
Force10# show vlan
Codes: * - Default VLAN, G - GVRP VLANs
    NUM    Status    Q Ports
*   1      Inactive  U Gi 0/0-23
    10     Inactive  T Gi 1/0-23
                     T Gi 2/0-23
```

Step **Task**

4 Assign port 0/0-23 as untagged and the rest of the ports 1/0-23, 2/0-23 are tagged to VLAN10.

```
snmpset -c public 123.45.6.78 .1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786 x "FF
FF FF 00 00 00 00 00 00 00 00 00 FF FF FF 00 00 00 00 00 00 00 00 00 FF FF
FF" .1.3.6.1.2.1.17.7.1.4.3.1.4.1107787786 x "FF FF FF"
```

5 The SNMP query now displays:

```
snmpwalk -c public 123.45.6.78 .1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786

SNMPv2-SMI::mib-2.17.7.1.4.3.1.2.1107787786 = Hex-STRING:
FF FF FF 00 00 00 00 00 00 00 00 00 FF FF FF 00
00 00 00 00 00 00 00 00 FF FF FF 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00

snmpwalk -c public 123.45.6.78 .1.3.6.1.2.1.17.7.1.4.3.1.4.1107787786

SNMPv2-SMI::mib-2.17.7.1.4.3.1.4.1107787786 = Hex-STRING:
FF FF FF 00 00 00 00 00 00 00 00 00 00 00 00 00  ? Port 1 – 24 are untagged
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
```

6 Execute the FTOS **show vlan** command:

```
Force10# show vlan
Codes: * - Default VLAN, G - GVRP VLANs
    NUM    Status    Q Ports
*   1      Inactive
    10     Inactive  U Gi 0/0-23? port 0-23 as untagged
                     T Gi 1/0-23
                     T Gi 2/0-23
```

Changing Administrative Status with SNMP Set

Follow the steps below to change administrative status using SNMP Set.

Step	Task

1	Get the interface index of the specific slot/port that you want to change:
	`snmpwalk -c public 123.45.6.78 .1.3.6.1.2.1.2.2.1.2`
	
	`IF-MIB::ifDescr.33865767 = STRING: GigabitEthernet 0/0`
	`IF-MIB::ifDescr.34127911 = STRING: GigabitEthernet 0/1`
	`IF-MIB::ifDescr.34390055 = STRING: GigabitEthernet 0/2`
	`IF-MIB::ifDescr.34652199 = STRING: GigabitEthernet 0/3`

2	Bring up the interface administrative status of slot 0/0:
	`snmpset -c public 123.45.6.78 .1.3.6.1.2.1.2.2.1.7. 33865767 i 1`
	
	`IF-MIB::ifAdminStatus. 33865767 = INTEGER: up(1)`

3	Shut down the interface slot 0/0 administrative status:
	`snmpset -c public 123.45.6.78 .1.3.6.1.2.1.2.2.1.7. 33865767 i 2`
	
	`IF-MIB::ifAdminStatus. 33865767 = INTEGER: down(2)`

4	Get the interface ifAdminStatus:
	`snmpwalk -c public 123.45.6.78 .1.3.6.1.2.1.2.2.1.7`
	
	`IF-MIB::ifAdminStatus.33865767 = INTEGER: down(2)`
	`IF-MIB::ifAdminStatus.34127911 = INTEGER: down(2)`
	`IF-MIB::ifAdminStatus.34390055 = INTEGER: down(2)`
	`IF-MIB::ifAdminStatus.34652199 = INTEGER: down(2)`

Assigning VLAN Alias Name using SNMP SET

You can assign an alias name for a VLAN using the **snmpset** command.

```
abcServer > snmpset -c public -v1 <ip-address of chassis> .1.3.6.1.2.1.17.7.1.4.3.1.1.
<if index of vlan>s ‘testvlan1’
```

You can view the if-index using the **show interface vlan** command.

```
Force10(conf-if-vl-2)#show interface vlan 2
Vlan 2 is down, line protocol is down
Vlan alias name is: testvlan1
Address is 00:01:e8:01:f6:d9, Current address is 00:01:e8:01:f6:d9
Interface index is 1107787778 -> if-index
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 21:41:10
Queueing strategy:fifo
InputStatistics:
  0 packets, 0 bytes
Time since last interface status change: 21:41:10
Force10#
```

Example of an **snmpset** for VLAN 2:

```
abcServer > snmpset -c public -v1 10.11.198.75 .1.3.6.1.2.1.17.7.1.4.3.1.1.
<1107787778s ‘testvlan’
```

Then issue the **show interface vlan** command again

```
Force10(conf-if-vl-2)#show interface vlan 2
Vlan 2 is down, line protocol is down
Vlan alias name is: testvlan1
Address is 00:01:e8:01:f6:d9, Current address is 00:01:e8:01:f6:d9
Interface index is 1107787778
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 21:41:10
Queueing strategy:fifo
InputStatistics:
  0 packets, 0 bytes
Time since last interface status change: 21:41:10
Force10(conf-if-vl-2)#
```

In the VLAN interface context, issue a **show config** command:

```
Force10(conf-if-vl-2)#show conf
!
interface Vlan 2
name testvlan1
no ip address
shutdown
Force10(conf-if-vl-2)#
```

List of Commands

Index