




Dell Force10 Configuration Guide for the MXL 10/40GbE Switch IO Module

Publication Date: July 2012



Force10

Notes, Cautions, and Warnings

-  **NOTE:** A NOTE indicates important information that helps you make better use of your computer.
-  **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.
-  **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Information in this publication is subject to change without notice.

© 2012 Dell Force10. All rights reserved.

Reproduction of these materials in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Trademarks used in this text: Dell™, the DELL logo, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, KACE™, FlexAddress™ and Vostro™ are trademarks of Dell Inc. Intel®, Pentium®, Xeon®, Core™ and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™, and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, MS-DOS® and Windows Vista® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat Enterprise Linux® and Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® is a registered trademark and SUSE™ is a trademark of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter®, and vSphere® are registered trademarks or trademarks of VMWare, Inc. in the United States or other countries.

Other trademarks and trade names may be used in this publication to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

1	About this Guide	23
	Objectives	23
	Audience	23
	Conventions	24
	Information Symbols	24
	Related Documents	24
2	Configuration Fundamentals	25
	Accessing the Command Line	25
	CLI Modes	26
	Navigating CLI Modes	27
	The do Command	29
	Undoing Commands	30
	Obtaining Help	31
	Entering and Editing Commands	32
	Command History	33
	Filtering show Command Outputs	33
	Multiple Users in Configuration Mode	35
3	Getting Started	37
	Console access	37
	Serial Console	37
	External Serial Port with a USB Connector	39
	Boot Process	39
	Default Configuration	42
	Configure a Host Name	42
	Access the System Remotely	42
	Access the MXL Switch Remotely	42
	Configure the Management Port IP Address	43
	Configure a Management Route	43
	Configure a Username and Password	44
	Configure the Enable Password	44
	Configuration File Management	45
	Copy Files to and from the System	45
	Important Points to Remember	45
	Save the Running-Configuration	46
	View Files	47
	View Configuration Files	48
	File System Management	49
	View the Command History	50
	Upgrading and Downgrading FTOS	50

4	Management	51
	Configure Privilege Levels	51
	Create a Custom Privilege Level	51
	Removing a Command from EXEC Mode	52
	Move a Command from EXEC Privilege Mode to EXEC Mode	52
	Allow Access to CONFIGURATION Mode Commands	52
	Allow Access to INTERFACE, LINE, ROUTE-MAP, and ROUTER Mode	52
	Apply a Privilege Level to a Terminal Line	55
	Configure Logging	55
	Log Messages in the Internal Buffer	55
	Configuration Task List for System Log Management	55
	Disable System Logging	56
	Send System Messages to a Syslog Server	56
	Configure a Unix System as a Syslog Server	56
	Change System Logging Settings	56
	Display the Logging Buffer and the Logging Configuration	57
	Configure a UNIX Logging Facility Level	59
	Synchronize log messages	60
	Enable timestamp on Syslog Messages	61
	File Transfer Services	61
	Configuration Task List for File Transfer Services	61
	Enable the FTP Server	62
	Configure the FTP Server Parameters	62
	Configure FTP Client Parameters	63
	Terminal Lines	63
	Configure Login Authentication for Terminal Lines	64
	Time Out of EXEC Privilege Mode	65
	Telnet to Another Network Device	66
	Lock CONFIGURATION Mode	66
	Viewing the Configuration Lock Status	67
	Recovering from a Forgotten Password	68
	Recovering from a Forgotten Enable Password	68
	Recovering from a Failed Start	69
5	Access Control Lists (ACLs)	71
	Overview	71
	IP Access Control Lists (ACLs)	72
	Implementing ACLs on FTOS	72
	ACLs and VLANs	73
	ACL Optimization	73
	Determine the Order in Which ACLs are Used to Classify Traffic	73
	IP Fragment Handling	74
	IP Fragments ACL Examples	74

Layer 4 ACL Rules Examples	75
Configure a Standard IP ACL	76
Configure an Extended IP ACL	78
Configure Filters with a Sequence Number	79
Configure Filters Without a Sequence Number	79
Established Flag	80
Configuring Layer 2 and Layer 3 ACLs on an Interface	81
Assign an IP ACL to an Interface	81
Counting ACL Hits	82
Configuring Ingress ACLs	83
Configuring Egress ACLs	84
Egress Layer 3 ACL Lookup for Control-Plane IP Traffic	84
IP Prefix Lists	85
Implementation Information	86
Configuration Task List for Prefix Lists	86
Configure a Prefix List	86
Use a Prefix List for Route Redistribution	89
ACL Resequencing	90
Resequencing an ACL or Prefix List	91
Route Maps	92
Implementation Information	92
Important Points to Remember	92
Configuration Task List for Route Maps	93
Create a Route Map	93
Configure Route Map Filters	95
Configure a Route Map for Route Redistribution	97
Configure a Route Map for Route Tagging	97
Continue Clause	98
6 Bare Metal Provisioning (BMP)	101
Overview	101
Auto-Configuration	103
BMP Mode	103
MAC-Based IP Assignment	103
DHCP Configuration	104
IP Server	105
Domain Name Server	105
Boot Commands	106
System Boot and Set-Up Behavior	106
7 Content Addressable Memory (CAM)	109
CAM Allocation	109
Test CAM Usage	110

View CAM-ACL Settings	111
CAM Optimization	112
8 Data Center Bridging (DCB)	113
Ethernet Enhancements in Data Center Bridging	113
Priority-Based Flow Control	114
Enhanced Transmission Selection	115
Data Center Bridging Exchange Protocol (DCBX)	117
Data Center Bridging in a Traffic Flow	117
Enabling Data Center Bridging	118
QoS dot1p Traffic Classification and Queue Assignment	119
Configuring Priority-Based Flow Control	120
Configuring Lossless Queues	122
Configuring the PFC Buffer in a Switch Stack	123
Configuring Enhanced Transmission Selection	124
ETS Prerequisites and Restrictions	124
Creating a QoS ETS Output Policy	125
Creating an ETS Priority Group	127
Applying an ETS Output Policy for a Priority Group to an Interface	128
ETS Operation with DCBX	129
Configuring Bandwidth Allocation for DCBX CIN	129
Applying DCB Policies in a Switch Stack	131
Configuring DCBX Operation	132
DCBX Operation	132
DCBX Port Roles	133
DCB Configuration Exchange	134
Configuration Source Election	135
Propagation of DCB Information	135
Auto-Detection and Manual Configuration of the DCBX Version	136
DCBX Example	136
DCBX Prerequisites and Restrictions	137
DCBX Configuration Procedure	138
Configuring DCBX on an Interface	138
Configuring DCBX Globally on the Switch	140
DCBX Error Messages	141
An error in DCBX operation is displayed using the following syslog messages:	141
Debugging DCBX on an Interface	142
Verifying DCB Configuration	143
PFC and ETS Configuration Examples	153
Using PFC and ETS to Manage Data Center Traffic	153
Using PFC and ETS to Manage Converged Ethernet Traffic in a Switch Stack	156
Hierarchical Scheduling in ETS Output Policies	157

9	Dynamic Host Configuration Protocol (DHCP)	159
	Overview	159
	DHCP Packet Format and Options	160
	Assigning an IP Address Using DHCP	161
	Implementation Information	162
	Configuration Tasks	162
	Configure the System to be a DHCP Server	162
	Configuration Tasks	163
	Related Configuration Tasks	163
	Configure the Server for Automatic Address Allocation	163
	Create an IP Address Pool	163
	Exclude Addresses from the Address Pool	164
	Specify an Address Lease Time	164
	Specify a Default Gateway	164
	Enable DHCP Server	165
	Configure a Method of Hostname Resolution	165
	Address Resolution using DNS	165
	Address Resolution using NetBIOS WINS	166
	Create Manual Binding Entries	166
	Debug DHCP Server	167
	DHCP Clear Commands	167
	Configure the System to be a Relay Agent	167
	Configure the System to be a DHCP Client	169
	DHCP Client on a Management Interface	175
	DHCP Client Operation with other Features	176
	Stacking	176
	VLT	176
	VLAN and Port Channels	176
	DHCP Snooping	176
	DHCP Server	176
	VRRP	177
	Configure Secure DHCP	178
	Option 82	178
	DHCP Snooping	179
	Enable DHCP Snooping	179
	Add a Static Entry in the Binding Table	180
	Clear the Binding Table	180
	Display the Contents of the Binding Table	180
	Drop DHCP Packets on Snooped VLANs Only	181
	Dynamic ARP Inspection	182
	Bypass the ARP Inspection	184
	Source Address Validation	184
	IP Source Address Validation	185

	DHCP MAC Source Address Validation	185
	IP+MAC Source Address Validation	185
10	FIP Snooping.	187
	Fibre Channel over Ethernet	187
	Ensuring Robustness in a Converged Ethernet Network	187
	FIP Snooping on Ethernet Bridges	189
	FIP Snooping in a Switch Stack	191
	Configuring FIP Snooping	191
	Enabling the FIP Snooping Feature	192
	Enabling FIP Snooping on VLANs	192
	Configuring the FC-MAP Value	192
	Configuring a Port for a Bridge-to-FCF Link	193
	Impact on other Software Features	193
	FIP Snooping Prerequisites	193
	FIP Snooping Restrictions	194
	Configuration Procedure	194
	Displaying FIP Snooping Information	195
	FIP Snooping Configuration Example	202
11	GARP VLAN Registration Protocol (GVRP)	205
	Overview	205
	Important Points to Remember	205
	Configuring GVRP	206
	Related Configuration Tasks	207
	Enabling GVRP Globally	208
	Enabling GVRP on a Layer 2 Interface	208
	Configuring GVRP Registration	208
	Configuring a GARP Timer	209
12	Internet Group Management Protocol (IGMP).	211
	Overview	211
	IGMP Version 2	211
	Joining a Multicast Group	212
	Leaving a Multicast Group	212
	IGMP Version 3	213
	Joining and Filtering Groups and Sources.	214
	Leaving and Staying in Groups	215
	IGMP Snooping	215
	IGMP Snooping Implementation Information	216
	Configuring IGMP Snooping	216
	Related Configuration Tasks	216

Enabling IGMP Immediate-leave	216
Disabling Multicast Flooding	217
Specifying a Port as Connected to a Multicast Router	217
Configuring the Switch as Querier	217
Adjusting the Last Member Query Interval	217
Fast Convergence after MSTP Topology Changes	218
Designating a Multicast Router Interface	218
13 Interfaces	219
Basic Interface Configuration:	219
Advanced Interface Configuration:	219
Interface Types	220
View Basic Interface Information	220
Enable a Physical Interface	223
Physical Interfaces	223
Configuration Task List for Physical Interfaces	223
Overview of Layer Modes	224
Configure Layer 2 (Data Link) Mode	224
Configure Layer 3 (Network) Mode	225
Management Interfaces	226
Configure Management Interfaces on the MXL Switch	227
VLAN Interfaces	229
Loopback Interfaces	230
Null Interfaces	230
Port Channel Interfaces	231
Port Channel Definition and Standards	231
Port Channel Benefits	231
Port Channel Implementation	231
100/1000/10000 Mbps Interfaces in Port Channels	232
Configuration Task List for Port Channel Interfaces	233
Create a Port Channel	233
Add a Physical Interface to a Port Channel	233
Reassign an Interface to a New Port Channel	236
Configure the Minimum oper up Links in a Port Channel (LAG)	237
Add or Remove a Port Channel from a VLAN	237
Assign an IP Address to a Port Channel	238
Delete or Disable a Port Channel	238
Bulk Configuration	238
Interface Range	238
Bulk Configuration Examples	239
Create a Single-Range	239
Create a Multiple-Range	239
Exclude Duplicate Entries	239

Exclude a Smaller Port Range	240
Overlap Port Ranges	240
Commas	240
Add Ranges	240
Interface Range Macros	241
Define the Interface Range	241
Choose an Interface-range Macro	241
Monitor and Maintain Interfaces	243
Maintenance Using TDR	244
Splitting QSFP Ports to SFP+ Ports	245
Important Points	246
MTU Size on an Interface	246
Layer 2 Flow Control Using Ethernet Pause Frames	247
Enable Pause Frames	247
Configure MTU Size on an Interface	248
Port-Pipes	249
Auto-Negotiation on Ethernet Interfaces	250
Setting Speed and Duplex Mode of Ethernet Interfaces	250
Setting Auto-Negotiation Options	252
Adjust the Keepalive Timer	253
View Advanced Interface Information	253
Display Only Configured Interfaces	253
Configure Interface Sampling Size	255
Dynamic Counters	257
Clear Interface Counters	257
14 IPv4 Routing	259
IP Addresses	259
Implementation Information	260
Configuration Task List for IP Addresses	260
Assign IP Addresses to an Interface	260
Configure Static Routes	261
Configure Static Routes for the Management Interface	263
Directed Broadcast	263
Resolution of Host Names	264
Enable Dynamic Resolution of Host Names	264
Specify Local System Domain and a List of Domains	265
DNS with Traceroute	265
Address Resolution Protocol (ARP)	266
Configuration Task List for ARP	267
Configure Static ARP Entries	267
Enable Proxy ARP	268
Clear ARP Cache	269

ARP Learning via Gratuitous ARP	269
ARP Learning via ARP Request	270
Configurable ARP Retries	270
Internet Control Message Protocol (ICMP)	271
Configuration Task List for ICMP	271
Enable ICMP Unreachable Messages	271
UDP Helper	272
Configuring UDP Helper	272
Important Points to Remember	272
Enabling UDP Helper	272
Configurations Using UDP Helper	273
UDP Helper with Broadcast-All Addresses	273
UDP Helper with Subnet Broadcast Addresses	274
UDP Helper with Configured Broadcast Addresses	274
UDP Helper with No Configured Broadcast Addresses	275
Troubleshooting UDP Helper	275
15 iSCSI Optimization	277
iSCSI Optimization Overview	277
Monitoring iSCSI Traffic Flows	279
Application of Quality of Service to iSCSI Traffic Flows	279
Information Monitored in iSCSI Traffic Flows	279
Detection and Autoconfiguration for Dell EqualLogic Arrays	280
Detection and Port Configuration for Dell Compellent Arrays	280
Enabling and Disabling iSCSI Optimization	281
Default iSCSI Optimization Values	282
iSCSI Optimization Prerequisites	282
Configuring iSCSI Optimization	283
Displaying iSCSI Optimization Information	284
16 Link Aggregation Control Protocol (LACP)	287
Introduction to Dynamic LAGs and LACP	287
Important Points to Remember	288
LACP Modes	288
LACP Configuration Commands	289
LACP Configuration Tasks	289
Create a LAG	289
Configure the LAG Interfaces as Dynamic	290
Set the LACP Long Timeout	290
Monitor and Debugging LACP	291
Shared LAG State Tracking	292
Configure Shared LAG State Tracking	292
Important Points about Shared LAG State Tracking	294

LACP Basic Configuration Example	294
Configuring a LAG on ALPHA	295
Summary of the Configuration on ALPHA	299
Summary of the Configuration on BRAVO	300
17 Layer 2	305
Managing the MAC Address Table	305
Clear the MAC Address Table	305
Set the Aging Time for Dynamic Entries	305
Configure a Static MAC Address	306
Display the MAC Address Table	306
MAC Learning Limit	307
MAC Learning Limit Dynamic	308
MAC Learning Limit Station-Move	308
Learning Limit Violation Actions	308
Station Move Violation Actions	308
Recovering from Learning Limit and Station Move Violations	309
Network Interface Controller (NIC) Teaming	309
MAC Move Optimization	310
18 Link Layer Discovery Protocol (LLDP)	313
Overview	313
Protocol Data Units	313
Optional TLVs	314
Management TLVs	315
Organizationally Specific TLVs	315
IEEE Organizationally Specific TLVs	315
TIA-1057 (LLDP-MED) Overview	316
TIA Organizationally Specific TLVs	317
LLDP-MED Capabilities TLV	318
LLDP-MED Network Policies TLV	319
Extended Power via MDI TLV	320
Configuring LLDP	320
Related Configuration Tasks	321
Important Points to Remember	321
LLDP Compatibility	321
CONFIGURATION versus INTERFACE Configurations	321
Enabling LLDP	322
Disabling and Undoing LLDP	322
Advertising TLVs	323
Viewing the LLDP Configuration	324
Viewing Information Advertised by Adjacent LLDP Agents	325
Configuring LLDPDU Intervals	327

Configuring Transmit and Receive Mode	328
Configuring a Time to Live	329
Debugging LLDP	330
Relevant Management Objects	331
19 Multiple Spanning Tree Protocol (MSTP)	337
Overview	337
Implementation Information	338
Configure Multiple Spanning Tree Protocol	338
Related Configuration Tasks	338
Enable Multiple Spanning Tree Globally	339
Create Multiple Spanning Tree Instances	339
Influence MSTP Root Selection	340
Interoperate with Non-FTOS Bridges	341
Modify Global Parameters	342
Enable BPDU Filtering globally	343
Modify Interface Parameters	344
Configure an EdgePort	345
Flush MAC Addresses after a Topology Change	346
MSTP Sample Configurations	346
Debugging and Verifying an MSTP Configuration	351
20 Open Shortest Path First (OSPFv2)	355
Overview	355
Autonomous System (AS) Areas	356
Area Types	357
Networks and Neighbors	357
Router Types	357
Backbone Router (BR)	359
Area Border Router (ABR)	359
Autonomous System Border Router (ASBR)	359
Internal Router (IR)	359
Designated and Backup Designated Routers	359
Link-State Advertisements (LSAs)	360
LSA Throttling	361
Router Priority and Cost	361
Implementing OSPF with FTOS	362
Fast Convergence (OSPFv2, IPv4 only)	363
Multi-Process OSPF (OSPFv2, IPv4 only)	363
Processing SNMP and Sending SNMP Traps	363
RFC-2328 Compliant OSPF Flooding	363
OSPF ACK Packing	364
OSPF Adjacency with Cisco Routers	365

Configuration Information	365
Configuration Task List for OSPFv2 (OSPF for IPv4)	366
Enable OSPFv2	366
Enable Multi-Process OSPF	368
Assign an OSPFv2 area	369
Enable OSPFv2 on Interfaces	369
Configure Stub Areas	371
Configure LSA Throttling Timers	372
Enable Passive Interfaces	372
Enable Fast-Convergence	374
Change OSPFv2 Parameters on Interfaces	375
Enable OSPFv2 Authentication	377
Filter Routes	378
Redistribute Routes	379
Troubleshooting OSPFv2	380
Sample Configurations for OSPFv2	382
Basic OSPFv2 Router Topology	382
21 Port Monitoring	385
Important Points to Remember	385
Port Monitoring	386
Configuring Port Monitoring	388
22 Private VLANs (PVLAN)	391
Private VLAN Concepts	392
Private VLAN Commands	393
Private VLAN Configuration Task List	394
Creating PVLAN Ports	394
Creating a Primary VLAN	395
Creating a Community VLAN	396
Creating an Isolated VLAN	396
Private VLAN Configuration Example	397
Inspecting the Private VLAN Configuration	398
23 Per-VLAN Spanning Tree Plus (PVST+)	403
Overview	403
Implementation Information	404
Configure Per-VLAN Spanning Tree Plus	404
Related Configuration Tasks	404
Enable PVST+	405
Disable PVST+	405
Influence PVST+ Root Selection	405

Modify Global PVST+ Parameters	407
Enable BPDU Filtering globally	408
Modify Interface PVST+ Parameters	409
Configure an EdgePort	410
PVST+ in Multi-vendor Networks	411
PVST+ Extended System ID	411
PVST+ Sample Configurations	413
24 Quality of Service (QoS)	415
Overview	415
Implementation Information	417
Port-Based QoS Configurations	417
Set dot1p Priorities for Incoming Traffic	418
Honor dot1p Priorities on Ingress Traffic	418
Priority-Tagged Frames on the Default VLAN	419
Configure Port-based Rate Policing	419
Configure Port-based Rate Shaping	420
Policy-Based QoS Configurations	420
Classify Traffic	421
Create a Layer 3 Class Map	421
Create a Layer 2 Class Map	422
Determine the Order in Which You Use ACLs to Classify Traffic	422
Set DSCP Values for Egress Packets Based on Flow	422
Display Configured Class Maps and Match Criteria	423
Create a QoS Policy	423
Create an Input QoS Policy	423
Create an Output QoS Policy	424
Create Policy Maps	425
Create Input Policy Maps	425
Apply an Input Policy Map to an Interface	429
Create Output Policy Maps	429
QoS Rate Adjustment	430
Strict-Priority Queueing	431
Weighted Random Early Detection	431
Create WRED Profiles	432
Apply a WRED Profile to Traffic	432
Display Default and Configured WRED Profiles	432
Display WRED Drop Statistics	433
25 Routing Information Protocol (RIP)	435
Overview	435
RIPv1	435
RIPv2	436

Implementation Information	436
Configuration Information	436
Configuration Task List for RIP	436
Enable RIP Globally	437
Configure RIP on Interfaces	438
Control RIP Routing Updates	439
Set the Send and Receive Version	440
Generate a Default Route	442
Summarize Routes	443
Control Route Metrics	443
Debug RIP	444
RIP Configuration Example	444
Configuring RIPv2 on Core 2	445
Core 2 Output	445
RIP Configuration on Core 3	447
Core 3 RIP Output	448
RIP Configuration Summary	450
26 Remote Monitoring (RMON)	453
Overview	453
Implementation	453
Fault Recovery	454
Set the RMON Alarm	455
Configure an RMON Event	456
Configure RMON Collection Statistics	457
Configure RMON Collection History	458
Enable an RMON MIB Collection History Group	458
27 Rapid Spanning Tree Protocol (RSTP)	461
Overview	461
Configuring Rapid Spanning Tree	461
Related Configuration Tasks	461
Important Points to Remember	462
Configure Interfaces for Layer 2 Mode	462
Enable Rapid Spanning Tree Protocol Globally	463
Add and Remove Interfaces	466
Modify Global Parameters	466
Enable BPDU Filtering globally	468
Modify Interface Parameters	468
Configure an EdgePort	469
Influence RSTP Root Selection	470
SNMP Traps for Root Elections and Topology Changes	471
Fast Hellos for Link State Detection	471

28 Security	473
AAA Accounting	473
Configuration Task List for AAA Accounting	473
Enable AAA Accounting	474
Suppress AAA Accounting for Null Username Sessions	474
Configure Accounting of EXEC and Privilege-Level Command Usage	475
Configure AAA Accounting for Terminal Lines	475
Monitor AAA Accounting	475
AAA Authentication	476
Configuration Task List for AAA Authentication	476
Configure Login Authentication for Terminal Lines	476
Configure AAA Authentication Login Methods	477
Enable AAA Authentication	478
AAA Authentication—RADIUS	478
Server-Side Configuration	479
AAA Authorization	479
Privilege Levels Overview	479
Configuration Task List for Privilege Levels	480
Configure a Username and Password	480
Configure the Enable Password Command	481
Configure Custom Privilege Levels	482
Specify the LINE Mode Password and Privilege	484
Enable and Disable Privilege Levels	484
RADIUS	484
RADIUS Authentication and Authorization	485
Idle Time	485
ACL	485
Auto-Command	486
Set Access to Privilege Levels through RADIUS	486
Configuration Task List for RADIUS	486
Define an aaa Method List to be Used for RADIUS	487
Apply the Method List to Terminal Lines	487
Specify a RADIUS Server Host	487
Set the Global Communication Parameters for all RADIUS Server Hosts	488
Monitor RADIUS	489
TACACS+	489
Configuration Task List for TACACS+	489
Choose TACACS+ as the Authentication Method	490
Monitor TACACS+	491
TACACS+ Remote Authentication and Authorization	491
Command Authorization	493
Protection from TCP Tiny and Overlapping Fragment Attacks	493
SCP and SSH	494

Using SCP with SSH to Copy a Software Image	495
Secure Shell Authentication	496
Important Points to Remember for SSH Authentication	496
SSH Authentication by Password	496
RSA Authentication of SSH	497
Host-Based SSH Authentication	497
Client-based SSH Authentication	499
Troubleshooting SSH	499
Telnet	500
VTY Line and Access-Class Configuration	500
VTY Line Local Authentication and Authorization	500
VTY Line Remote Authentication and Authorization	501
VTY MAC-SA Filter Support	502
29 sFlow	505
Overview	505
Implementation Information	506
Important Points to Remember	506
Enable and Disable sFlow	507
Enable and Disable on an Interface	507
sFlow Show Commands	507
Show sFlow Globally	508
Show sFlow on an Interface	508
Show sFlow on a Stack Unit	509
Specify Collectors	509
Polling Intervals	509
Sampling Rate	510
Sub-Sampling	510
Back-Off Mechanism	511
sFlow on LAG ports	511
Extended sFlow	511
30 Simple Network Management Protocol (SNMP)	515
Protocol Overview	515
Implementation Information	515
Configure Simple Network Management Protocol	515
Related Configuration Tasks	517
Important Points to Remember	517
Setting up SNMP	517
Create a Community	518
Setting Up User-based Security (SNMPv3)	518
Read Managed Object Values	520
Write Managed Object Values	521

Configure Contact and Location Information Using SNMP	521
Subscribe to Managed Object Value Updates using SNMP	522
Copy Configuration Files Using SNMP	525
Manage VLANs Using SNMP	531
Create a VLAN	532
Assign a VLAN Alias	532
Display the Ports in a VLAN	533
Add Tagged and Untagged Ports to a VLAN	534
Enable and Disable a Port Using SNMP	536
Fetch Dynamic MAC Entries Using SNMP	536
Deriving Interface Indices	538
Monitor Port-channels	539
BMP functionality using SNMP SET	540
Entity MIBS	541
Troubleshooting SNMP Operations	544
31 Stacking.....	545
Overview	545
Stacking MXL 10/40GbE Switches	545
Stack Management Roles	546
Stack Master Election	547
Failover Roles	548
MAC Addressing	548
Stacking LAG	548
Supported Stacking Topologies	549
Example 1: Dual-Ring Stack Across Multiple Chassis.....	549
Example 2: Dual Daisy-Chain Stack Across Multiple Chassis.....	550
Stack Group/Port Numbers	551
Configuring a Switch Stack	551
Stacking Prerequisites	552
Cabling Stacked Switches	552
Cabling Restrictions.....	552
Cabling Redundancy	552
Cabling Procedure.....	553
Accessing the CLI	553
Configuring and Bringing Up a Stack	553
Assigning a Priority to Stacked Switches.....	554
Renumbering a Stack Unit.....	555
Provisioning a Stack Unit.....	555
Converting 4x10GbE Ports to 40GbE for Stacking	556
Removing a Port from the Stacking Mode.....	556
Removing a Switch from a Stack	557
Adding a Stack Unit	557

Merging Two Stacks	558
Splitting a Stack	559
Managing Redundant Stack Management	559
Reset a Unit on a Stack	560
Verifying a Stack Configuration	560
Using LEDs	560
Using Show Commands	560
Troubleshooting a Switch Stack	565
Troubleshooting Commands	565
Failure Scenarios	567
Stack Member Fails	567
Unplugged Stacking Cable	567
Master Switch Fails	568
Stack-Link Flapping Error	568
Master Switch Recovers from Failure	569
Stack Unit in Card-Problem State Due to Incorrect FTOS Version	569
Stack Unit in Card-Problem State Due to Configuration Mismatch	570
Upgrading a Switch Stack	571
Upgrading a Single Stack Unit	572
32 Storm Control	575
Overview	575
Configure Storm Control	575
Configure Storm Control from INTERFACE Mode	575
Configure Storm Control from CONFIGURATION Mode	575
33 Spanning Tree Protocol (STP)	577
Overview	577
Configuring Spanning Tree	578
Related Configuration Tasks	578
Important Points to Remember	578
Configuring Interfaces for Layer 2 Mode	579
Enabling Spanning Tree Protocol Globally	580
Adding an Interface to the Spanning Tree Group	582
Removing an Interface from the Spanning Tree Group	582
Modifying Global Parameters	583
Modifying Interface STP Parameters	584
Enabling PortFast	584
Preventing Network Disruptions with BPDU Guard	585
BPDU Filtering	587
STP Root Selection	589
STP Root Guard	589
Root Guard Scenario	589

Root Guard Configuration	592
SNMP Traps for Root Elections and Topology Changes	592
Displaying STP Guard Configuration	593
34 System Time and Date	595
Network Time Protocol	595
Overview	596
Implementation Information	597
Configuring Network Time Protocol	597
Related Configuration Tasks	597
Enable NTP	597
Set the Hardware Clock with the Time Derived from NTP	598
Configure NTP Broadcasts	599
Disable NTP on an Interface	599
Configure a Source IP Address for NTP Packets	599
Configure NTP Authentication	600
FTOS Time and Date	603
Configuring Time and Date Settings	603
Set the Time and Date for the Switch Hardware Clock	603
Set the Time and Date for the Switch Software Clock.	604
Set the Timezone.	604
Set Daylight Savings Time	605
Set Daylight Saving Time Once.	605
Set Recurring Daylight Saving Time	606
35 Uplink Failure Detection (UFD)	609
Feature Description	609
How Uplink Failure Detection Works	610
UFD and NIC Teaming	611
Important Points to Remember	611
Configuring Uplink Failure Detection	613
Clearing a UFD-Disabled Interface	614
Displaying Uplink Failure Detection	616
Sample Configuration: Uplink Failure Detection	619
36 Upgrade Procedures	621
Find the Upgrade Procedures	621
Get Help with Upgrades	621
37 Virtual LANs (VLAN)	623
Default VLAN	624
Port-Based VLANs	625

VLANs and Port Tagging	625
Configuration Task List for VLANs	626
Create a Port-Based VLAN	626
Assign Interfaces to a VLAN	627
Assign an IP Address to a VLAN	630
Native VLANs	630
Enable Null VLAN as the Default VLAN	631
38 Virtual Router Redundancy Protocol (VRRP)	633
Overview	633
VRRP Benefits	635
VRRP Implementation	635
VRRP Configuration	636
Configuration Task List for VRRP	636
Create a Virtual Router	636
Assign Virtual IP addresses	637
Set the VRRP Group (Virtual Router) Priority	639
Configure VRRP Authentication	640
Disable Preempt	641
Change the Advertisement Interval	642
Track an Interface or Object	643
VRRP Initialization Delay	645
Sample Configurations	646
VRRP for IPv4 Configuration	646
39 Debugging and Diagnostics.	651
Offline Diagnostics	651
Important Points to Remember	652
Running Offline Diagnostics	652
Trace Logs	653
Auto Save on Crash or Rollover	653
Show Hardware Commands	653
Environmental Monitoring	654
Recognize an Over-Temperature Condition	656
Troubleshoot an Over-Temperature Condition	656
Recognize an Under-Voltage Condition	657
Troubleshoot an Under-Voltage Condition	657
Buffer Tuning	658
Deciding to Tune Buffers	660
Buffer Tuning Commands	660
Using a Pre-Defined Buffer Profile.	662
Sample Buffer Profile Configuration	663
Troubleshooting Packet Loss	664

Displaying Drop Counters	664
Dataplane Statistics	666
Displaying Stack Port Statistics	668
Displaying Stack Member Counters	668
Application Core Dumps	669
Mini Core Dumps	669
TCP Dumps	671
40 Standards Compliance	673
IEEE Compliance	673
RFC and I-D Compliance	674
General Internet Protocols	674
General IPv4 Protocols	675
Border Gateway Protocol (BGP)	676
Routing Information Protocol (RIP)	677
Open Shortest Path First (OSPF)	677
Network Management	678
MIB Location	682
41 Index	683

About this Guide

Objectives

This guide describes the supported protocols and software features, and provides configuration instructions and examples, for the Dell Force10 MXL 10/40GbE Switch IO Module running FTOS version 8.3.16.1.

The MXL 10/40GbE Switch IO Module is installed in a Dell PowerEdge M1000e Enclosure. For information about how to install and perform the initial switch configuration, refer to the *Getting Started Guides* on the Dell Support website at <http://support.dell.com/manuals>.

Though this guide contains information about protocols, it is not intended to be a complete reference. This guide is a reference for configuring protocols on Dell Force10 systems. For complete information about protocols, refer to other documentation, including IETF requests for comment (RFCs). The instructions in this guide cite relevant RFCs, and [Standards Compliance](#) 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 must be entered in the CLI as listed.
<i>parameter</i>	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 1-1 describes symbols contained in this guide.

Table 1-1. Information Symbols

Symbol	Warning	Description
	Note	This symbol informs you of important operational information.
	FTOS Behavior	This symbol informs you of an FTOS behavior. These behaviors are inherent to the Dell Force10 system or FTOS feature and are non-configurable.
	Exception	This symbol is a note associated with some other text on the page that is marked with an asterisk.

Related Documents

For more information about the Dell Force10 MXL 10/40GbE Switch IO Module, refer to the following documents:

- *FTOS Command Reference*
- *Users's Guide for the MXL 10/40GbE Switch IO Module*
- *FTOS Release Notes*

Configuration Fundamentals

The Dell Force10 operating software (FTOS) command line interface (CLI) is a text-based interface through which you can configure interfaces and protocols. 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 you enable a command, 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. For more information, refer to [Save the Running-Configuration](#).



Note: You can use the chassis management controller (CMC) out-of-band management interface to access and manage an MXL Switch using the FTOS command-line interface. For information about how to access the CMC to configure an MXL Switch, refer to the *Dell Chassis Management Controller (CMC) User's Guide* on the Dell Support website at <http://support.dell.com/support/edocs/systems/pem/en/index.htm>.

Accessing the Command Line

Access the command line through a serial console port or a Telnet session ([Figure 2-1](#)). When the system successfully boots, enter the command line in EXEC mode.

Figure 2-1. 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:
FTOS> ← 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 exception of EXEC mode commands preceded by the command `do`; for more information, refer to [The do Command](#) and EXEC Privilege Mode commands). You can set user access rights to commands and command modes using privilege levels; for more information about privilege levels and security options, refer to [Security](#).

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 the show commands, which allow you to view system information.
- **EXEC Privilege mode** 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. For more information, refer to [Configure the Enable Password](#).
- **CONFIGURATION mode** allows you to configure security features, time settings, set logging and simple network management protocol (SNMP) functions, and static address resolution protocol (ARP) and MAC addresses on the system.

Beneath CONFIGURATION mode are sub-modes that apply to interfaces, protocols, and features. [Figure 2-2](#) shows this sub-mode command structure. When configuring the chassis for the first time, the following two sub-CONFIGURATION modes are important:

- **INTERFACE sub-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, 10-Gigabit Ethernet, or 40-Gigabit Ethernet) or logical (Loopback, Null, port channel, or VLAN).
- **LINE sub-mode** is the mode in which you configure the console and virtual terminal lines.



Note: At any time, entering a question mark (?) displays the available command options. For example, when you are in CONFIGURATION mode, entering the question mark first lists all the available commands, including the possible sub-modes.

Figure 2-2. CLI Modes in FTOS

```

EXEC
EXEC Privilege
CONFIGURATION
  INTERFACE
    TEN GIGABIT ETHERNET
    FORTY GIGABIT ETHERNET
    INTERFACE RANGE
    LOOPBACK
    MANAGEMENT ETHERNET
    MONITOR SESSION
    NULL
    PORT-CANNEL
    VLAN
  IP
  IP ACCESS-LIST
    STANDARD ACCESS-LIST
    EXTENDED ACCESS-LIST
  LINE
    CONSOLE
    VIRTUAL TERMINAL
  MAC ACCESS-LIST
  MONITOR SESSION
  MULTIPLE SPANNING TREE
  PROTOCOL GVRP
  PROTOCOL LLDP
  Per-VLAN SPANNING TREE
  RAPID SPANNING TREE
  ROUTE-MAP
  ROUTER OSPF
  ROUTER RIP
  SPANNING TREE

```

Navigating CLI Modes

The FTOS prompt changes to indicate the CLI mode. [Table 2-1](#) lists the CLI mode, its prompt, and information about 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 and the exit command moves you up one command mode level.



Note: Sub-CONFIGURATION modes all have the letters “conf” in the prompt with additional modifiers to identify the mode and slot/port information. These are shown in [Table 2-1](#).

Table 2-1. FTOS Command Modes

CLI Command Mode	Prompt	Access Command
EXEC	FTOS>	Access the router through the console or Telnet.
EXEC Privilege	FTOS#	<ul style="list-style-type: none"> From EXEC mode, enter the command enable. From any other mode, enter the command end.

Table 2-1. FTOS Command Modes


CLI Command Mode	Prompt	Access Command	
CONFIGURATION	FTOS(conf)#	<ul style="list-style-type: none"> From EXEC privilege mode, enter the command configure. From every mode except EXEC and EXEC Privilege, enter the command exit. 	
 Note: Access the following modes from CONFIGURATION mode:			
INTERFACE modes	10 Gigabit Ethernet Interface	FTOS(conf-if-te-0/1)#	
	40 Gigabit Ethernet Interface	FTOS(conf-if-fo-0/33)#	
	Interface Range	FTOS(conf-if-range)#	
	Loopback Interface	FTOS(conf-if-lo-0)#	
	Management Ethernet Interface	FTOS(conf-if-ma-0/0)#	interface
	Null Interface	FTOS(conf-if-nu-0)#	
	Port-channel Interface	FTOS(conf-if-po-1)#	
	VLAN Interface	FTOS(conf-if-vl-1)#	
Monitor Session	FTOS(conf-mon-sess)		
IP ACCESS-LIST	STANDARD ACCESS-LIST	FTOS(conf-std-nacl)#	ip access-list standard
	EXTENDED ACCESS-LIST	FTOS(conf-ext-nacl)#	ip access-list extended
	IP COMMUNITY-LIST	FTOS(conf-community-list)#	ip community-list
LINE	CONSOLE	FTOS(conf-line-console)#	line
	VIRTUAL TERMINAL	FTOS(conf-line-vty)#	
MAC ACCESS-LIST	STANDARD ACCESS-LIST	FTOS(conf-std-macl)#	mac access-list standard
	EXTENDED ACCESS-LIST	FTOS(conf-ext-macl)#	mac access-list extended
	MULTIPLE SPANNING TREE	FTOS(conf-mstp)#	protocol spanning-tree mstp

Table 2-1. FTOS Command Modes

CLI Command Mode	Prompt	Access Command
PROTOCOL GVRP	FTOS(conf-gvrp)	protocol gvrp
PROTOCOL LLDP	FTOS(conf-lldp)	protocol lldp
Per-VLAN SPANNING TREE Plus	FTOS(conf-pvst)#	protocol spanning-tree pvst
RAPID SPANNING TREE	FTOS(conf-rstp)#	protocol spanning-tree rstp
ROUTE-MAP	FTOS(conf-route-map)#	route-map
ROUTER OSPF	FTOS(conf-router_ospf)#	router ospf
ROUTER RIP	FTOS(conf-router_rip)#	router rip
SPANNING TREE	FTOS(conf-stp)#	protocol spanning-tree 0

Figure 2-3 shows how to change the command mode from CONFIGURATION mode to PROTOCOL SPANNING TREE.

Figure 2-3. Changing CLI Modes

```
FTOS(conf)#protocol spanning-tree 0
FTOS(conf-stp)# ← New command prompt
```

The do Command

Enter an EXEC mode or EXEC privilege 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 2-4 illustrates the do command.



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

Figure 2-4. Using the do Command

```

FTOS(conf)#do show system brief
Stack MAC : 00:1e:c9:f1:04:22

Reload Type : normal-reload [Next boot : normal-reload]

-- Stack Info --
Unit  UnitType   Status      ReqTyp      CurTyp      Version     Ports
-----
0     Management  online     MXL-10/40GbE  MXL-10/40GbE  8-3-16-47   56
1     Member      not present
2     Member      not present
3     Member      not present
4     Member      not present
5     Member      not present

```

“do” form of show command

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 2-5](#).



Note: Use the help or ? command as described in [Obtaining Help](#) to help you construct the *no* form of a command.

Figure 2-5. Undoing a command with the no Command

```

FTOS(conf)#interface tengigabitethernet 5/1
FTOS(conf-if-te-5/1)#ip address 192.168.10.1/24
FTOS(conf-if-te-5/1)#show config
!
interface TenGigabitEthernet 5/1
!
ip address 192.168.10.1/24 ← IP address assigned

shutdown
FTOS(conf-if-te-5/1)#no ip address ← “no” form of IP address command
FTOS(conf-if-te-5/1)#show config
!
interface TenGigabitEthernet 5/1
no ip address ← IP address removed

shutdown
FTOS(conf-if-te-5/1)#

```


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 commands:

- 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 2-6. ? Command Example

```
FTOS#? ← "?" at prompt for list of commands
start      Start Shell
capture    Capture Packet
cd          Change current directory
clear      Reset functions
clock      Manage the system clock
configure  Configuring from terminal
copy       Copy from one file to another
--More--
```

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

Figure 2-7. Keyword? Command Example

```
FTOS(conf)#cl? ← partial keyword plus "[space]?" for matching keywords
class-map
clock
FTOS(conf)#cl
```

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

Figure 2-8. Keyword ? Command Example

```
FTOS(conf)#clock ? ← keyword plus "[space]?" for compatible keywords
summer-time      Configure summer (daylight savings) time
timezone         Configure time zone
FTOS(conf)#clock
```

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`.” You can, however, enter `clo` as a partial keyword because only one command begins with those three letters.
 - The TAB key auto-completes keywords in commands.
 - The UP and DOWN arrow keys display previously entered commands (refer to [Command History](#)).
 - The BACKSPACE and DELETE keys erase the previous letter.
 - Key combinations are available to move quickly across the command line, refer to [Table 2-2](#).

Table 2-2. 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 the 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.
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 the previously-entered EXEC mode commands.
- When you are in CONFIGURATION mode, the UP or DOWN arrows keys recall 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 | save] specified_text` after the command. The variable *specified_text* is the text for which you are filtering and it IS case sensitive unless you use the ignore-case sub-option.

The `grep` command accepts an ignore-case sub-option that forces the search to be case-*insensitive*. For example, the commands:

- `show run | grep Ethernet` returns a search result with instances containing a capitalized “Ethernet,” such as `interface TenGigabitEthernet 0/0`.
- `show run | grep ethernet` would not return that search result because it only searches for instances containing a non-capitalized “ethernet.”

Executing the `show run | grep Ethernet ignore-case` command would return instances containing both “Ethernet” and “ethernet.”

- `grep` displays only the lines containing specified text. [Figure 2-9](#) shows this command used in combination with the `do show stack-unit all stack-ports pfc details | grep 0` command.

Figure 2-9. Filtering Command Outputs with the grep Command

```
FTOS(conf)#do show stack-unit all stack-ports all pfc details | grep 0
stack unit 0 stack-port all
  0 Pause Tx pkts, 0 Pause Rx pkts
  0 Pause Tx pkts, 0 Pause Rx pkts
  0 Pause Tx pkts, 0 Pause Rx pkts
  0 Pause Tx pkts, 0 Pause Rx pkts
  0 Pause Tx pkts, 0 Pause Rx pkts
  0 Pause Tx pkts, 0 Pause Rx pkts
```



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 2-10](#) shows this command used in combination with the do show stack-unit all stack-ports all pfc details | except 0 command.

Figure 2-10. Filtering Command Outputs with the except Command

```
FTOS(conf)#do show stack-unit all stack-ports all pfc details | except 0

Admin mode is On
Admin is enabled
Local is enabled
Link Delay 45556 pause quantum

stack unit 1 stack-port all
Admin mode is On
Admin is enabled
```

- find displays the output of the show command beginning from the first occurrence of specified text [Figure 2-11](#) shows this command.

Figure 2-11. Filtering Command Outputs with the find Command

```
FTOS(conf)#do show stack-unit all stack-ports all pfc details | find 0
stack unit 0 stack-port all
Admin mode is On
Admin is enabled
Local is enabled
Link Delay 45556 pause quantum
0 Pause Tx pkts, 0 Pause Rx pkts

stack unit 1 stack-port all
```

- no-more displays the output all at once rather than one screen at a time. This is similar to the terminal length command except that the no-more option affects the output of the specified command only.
- save copies the output to a file for future reference.



Note: You can filter a single command output multiple times. The save option should be the last option entered. For example:

```
FTOS# command | grep regular-expression | except regular-expression | grep
other-regular-expression | find regular-expression | save
```

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 appear, Dell Force10 recommends coordinating with the users listed in the message so that you do not unintentionally overwrite each other's configuration changes.

Getting Started

This chapter contains the following major sections:

- [Console access](#)
- [Boot Process](#)
- [Default Configuration](#)
- [Configure a Host Name](#)
- [Access the System Remotely](#)
- [Configure the Enable Password](#)
- [Configuration File Management](#)
- [File System Management](#)
- [View the Command History](#)
- [Upgrading and Downgrading FTOS](#)

When the boot process is complete, the console monitor displays the Dell Force10 operating software (FTOS) banner and EXEC mode prompt ([Figure 3-2](#)).

For details about using the command line interface (CLI), refer to the [Accessing the Command Line](#) section in the [Configuration Fundamentals](#) chapter.


Console access

The MXL 10/40GbE Switch IO Module has two management ports available for system access: a serial console port and an out-of-bounds (OOB) port.

Serial Console

A universal serial bus (USB) (A-Type) connector is located at the front panel. The USB can be defined as an External Serial Console (RS-232) port, and is labeled on the MXL 10/40GbE Switch IO Module chassis. The USB is present on the lower side, as you face the I/O side of the chassis ([Figure 3-1](#)).

Figure 3-1. Serial Console



For the console port pinout, refer to [Table 3-1](#).

To access the console port, follow these steps.

Step	Task
1	Connect the USB connector to the front panel. Use the RS-232 Serial Line cable to connect the MXL 10/40GbE Switch IO Module console port to a terminal server.
2	Connect the other end of the cable to the DTE terminal server.
3	Terminal settings on the console port cannot be changed in the software and are set as follows: <ul style="list-style-type: none">• 9600 baud rate• No parity• 8 data bits• 1 stop bit• No flow control

External Serial Port with a USB Connector

[Table 3-1](#) lists the pin assignments.

Table 3-1. Pin Assignments

USB Pin Number	Signal Name
Pin 1	RTS
Pin 2	RX
Pin 3	TX
Pin 4	CTS
Pin 5, 6	GND
RxD	Chassis GND

Boot Process

After you follow the *Installation Procedure* in the *Getting Started Guide*, the MXL Switch boots up. The MXL Switch with FTOS version 8.3.16.1 requires boot flash version 4.0.1.0 and boot selector version 4.0.0.0. [Figure 3-2](#) and [Figure 3-3](#) show the completed boot process.

Figure 3-2. Completed Boot Process

```

syncing disks... done
unmounting file systems...
unmounting /f10/flash (/dev/ld0e)...
unmounting /usr (mfs:31)...
unmounting /lib (mfs:23)...
unmounting /f10 (mfs:20)...
unmounting /tmp (mfs:15)...
unmounting /kern (kernfs)...
unmounting / (/dev/md0a)... done
rebooting...

NetLogic XLP Stage 1 Loader
Built by build at tools-sjc-01 on Thu May 31 23:53:38 2012
IOM Boot Selector Label 4.0.0.0

Nodes online: 1
  GPIO 22 init'ed as an output
  GPIO 23 init'ed as an output
I2C0 speed = 30 KHz, prescaler = 0x0377.
Initialized I2C0 Controller.
I2C1 speed = 100 KHz, prescaler = 0x0109.
Initialized I2C1 Controller.
DDR SPD: Node 0 Channel 0 Mem size = 2048 MB
DDR SPD: Node 0 DRAM frequency 666 MHz
DDR SPD: Node 0 CPU frequency 1200 MHz
RTT Norm:44
NBU0 DRAM BAR0 base: 00000000 limit: 0013f000 xlate: 00000001 node: 00000000 ( 0 MB -> 320 MB
, size: 320 MB)
NBU0 DRAM BAR1 base: 001d0000 limit: 0088f000 xlate: 00090001 node: 00000000 ( 464 MB -> 2192 MB
, size: 1728 MB)
Modifying Default Flash Address map..Done
Initialized eMMC Host Controller
Detected SD Card
BLC is 1 (preset 10)
Hit any key to stop autoboot: 0
Boot Image selection
Reading the Boot Block Info...Passed !!
Images are OK A:0x0 B:0x0
Boot Selector set to Bootflash Partition A image...
Verifying Copyright Information..success for Image - 0
Boot Selector: Booting Bootflash Partition A image...
Copying stage-2 loader from 0xb6120000 to 0x8c100000(size = 0x100000)
Boot Image selection DONE.
## Starting application at 0x8C100000 ...

U-Boot 2010.03-rc1(Dell Force10)
Built by build at tools-sjc-01 on Thu May 31 23:53:38 2012
IOM Boot Label 4.0.1.0

```

Figure 3-3. Completed Boot Process (Contd.)

```
DRAM: 2 GB
Initialized CPLD on CS3
Detected [XLP308 (Lite+) Rev A0]
Initializing I2C0: speed = 30 KHz, prescaler = 0x0377 -- done.
Initializing I2C1: speed = 100 KHz, prescaler = 0x0109 -- done.
Initialized eMMC Host Controller
Detected SD Card
Now running in RAM - U-Boot [N64 ABI, Big-Endian] at: ffffffff8c100000
Flash: 256 MB
PCIE (B0:D01:F0) : Link up.
PCIE (B0:D01:F1) : No Link.
In: serial
Out: serial
Err: serial
Net: nae-0: PHY is Broadcom BCM54616S

--More--

SOFTWARE IMAGE HEADER DATA :
-----

--More--

Starting Dell Force10 application

Welcome to Dell Easy Setup Wizard

The setup wizard guides you through the initial switch configuration, and gets
you up and running as quickly as possible. You can skip the setup wizard, and
enter CLI mode to manually configure the switch. You must respond to the next
question to run the setup wizard within 60 seconds, otherwise the system will
continue with normal operation using the default system configuration.
Note: You can exit the setup wizard at any point by entering [ctrl+c].

Would you like to run the setup wizard (you must answer this question within
60 seconds)? [Y/N]: N
00:00:40: %STKUNIT0-M:CP %IFMGR-5-ASTATE_UP: Changed interface Admin state to up: V1 1
00:00:42: %STKUNIT0-M:CP %IFMGR-5-IFM_ISCSI_ENABLE: iSCSI has been enabled causing flow control
to be enabled on all interfaces.
EQL detection and enabling iscsi profile-compellent on an interface may cause some automatic
configurations to occur like jumbo frames on all ports and no storm control
and spanning tree port-fast on the port of detection
00:00:42: %STKUNIT0-M:CP %SEC-5-LOGIN_SUCCESS: Login successful for user on line console
FTOS>en
Password:
```

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 FTOS). You must configure the system using the CLI.

Configure a Host Name

The host name appears in the prompt. The default host name is FTOS.

- 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, follow this steps:

Step	Task	Command Syntax	Command Mode
1	Create a new host name.	hostname <i>name</i>	CONFIGURATION

Figure 3-4 shows the hostname command.

Figure 3-4. Configuring a Hostname



Access the System Remotely

You can configure the system to access it remotely by Telnet. The MXL 10/40GbE Switch IO Module has a dedicated management port and a management routing table that is separate from the IP routing table.

Access the MXL Switch Remotely

Configuring the system for Telnet is a three-step process:

1. Configure an IP address for the management port. Refer to [Configure the Management Port IP Address](#).
2. Configure a management route with a default gateway. Refer to [Configure a Management Route](#).

3. Configure a username and password. Refer to [Configure a Username and Password](#).

Configure the Management Port IP Address

Assign IP addresses to the management ports in order to access the system remotely.

To configure the management port IP address, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Enter INTERFACE mode for the Management port.	interface ManagementEthernet <i>slot/port</i> <ul style="list-style-type: none">• <i>slot</i>: 0• <i>port</i>: 0	CONFIGURATION
2	Assign an IP address to the interface.	ip address <i>ip-address/mask</i> <ul style="list-style-type: none">• <i>ip-address</i>: an address in dotted-decimal format (A.B.C.D).• <i>mask</i>: a subnet mask in /prefix-length format (/xx).	INTERFACE
3	Enable the interface.	no shutdown	INTERFACE

Configure 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 through the management port.

To configure a management route, follow this step:

Step	Task	Command Syntax	Command Mode
1	Configure a management route to the network from which you are accessing the system.	management route <i>ip-address/mask gateway</i> <ul style="list-style-type: none">• <i>ip-address</i>: the network address in dotted-decimal format (A.B.C.D).• <i>mask</i>: a subnet mask in /prefix-length format (/xx).• <i>gateway</i>: the next hop for network traffic originating from the management port.	CONFIGURATION

Configure a Username and Password

Configure a system username and password to access the system remotely.

To configure a username and password, follow this step:

Step	Task	Command Syntax	Command Mode
1	Configure a username and password to access the system remotely.	<pre>username <i>username</i> password [<i>encryption-type</i>] <i>password</i> <i>encryption-type</i> specifies how you are inputting the password, is 0 by default, and is not required.</pre> <ul style="list-style-type: none"> 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 Dell Force10 system. 	CONFIGURATION

Configure the Enable Password

Access EXEC Privilege mode using the enable command. EXEC Privilege mode is unrestricted by default. Configure 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 DES encryption method.
- enable secret is stored in the running/startup configuration in using a stronger, MD5 encryption method.


Dell Force10 recommends using the enable secret password.

To configure an enable password:

Task	Command Syntax	Command Mode
Create a password to access EXEC Privilege mode.	<pre>enable [password secret] [level <i>level</i>] [<i>encryption-type</i>] <i>password</i></pre> <p><i>level</i> is the privilege level, is 15 by default, and is not required.</p> <p><i>encryption-type</i> specifies how you are inputting the password, is 0 by default, and is not required.</p> <ul style="list-style-type: none"> 0 is for inputting the password in clear text. 7 is for inputting a password that is already encrypted using a DES hash. Obtain the encrypted password from the configuration file of another Dell Force10 system. Can be used only for enable password. 5 is for inputting a password that is already encrypted using an MD5 hash. Obtain the encrypted password from the configuration file of another Dell Force10 system. Can be used only for enable secret password. 	CONFIGURATION

Configuration File Management

You can store on and access files from various storage media. Rename, delete, and copy files on the system from EXEC Privilege mode.

 **Note:** Using flash memory cards in the system that have not been approved by Dell Force10 can cause unexpected system behavior, including a reboot.

Copy Files to and from the System

The command syntax for copying files is similar to UNIX. The copy command uses the format `copy source-file-url destination-file-url`.

 **Note:** For a detailed description of the copy command, refer to the *FTOS Command Reference Guide*.

- To copy a local file to a remote system, combine the *file-origin* syntax for a local file location with the *file-destination* syntax for a remote file location (Table 3-2).
- To copy a remote file to a Dell Force10 system, combine the *file-origin* syntax for a remote file location with the *file-destination* syntax for a local file location (Table 3-2).

Table 3-2. Forming a copy Command

	<i>source-file-url</i> Syntax	<i>destination-file-url</i> Syntax
Local File Location		
Internal flash:		
flash	copy flash://filename	flash://filename
USB flash:		
usbflash	usbflash://filename	usbflash://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://username:password@{hostip hostname}/filepath/filename	scp://username:password@{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.
- When copying to a server, you can only use a hostname if a DNS server is configured.

Figure 3-5 shows an example of using the copy command to save a file to an FTP server.

Figure 3-5. Copying a file to a Remote System


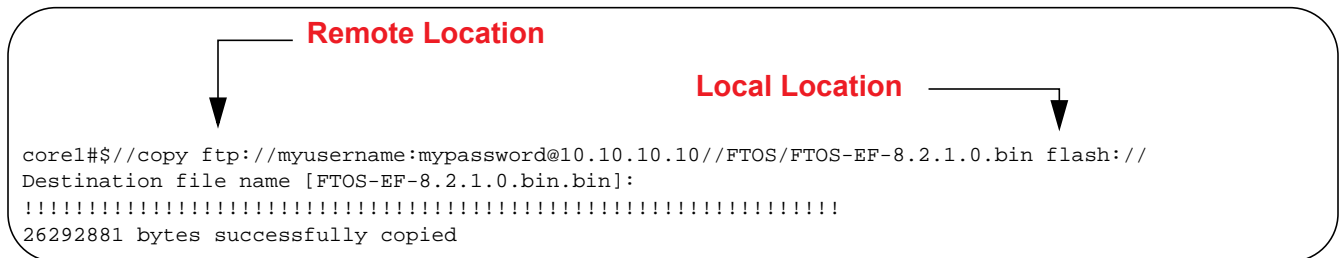


Figure 3-6 shows an example of using the copy command to import a file to the Dell Force10 system from an FTP server.

Figure 3-6. Copying a file from a Remote System



Save the Running-Configuration

The running-configuration contains the current system configuration. Dell Force10 recommends copying 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 IOM by default, but you can save the startup-configuration to a USB flash device or on a remote server.

To save the running-configuration:



Note: The commands in this section follow the same format as those in [Copy Files to and from the System on page 45](#) but use the filenames *startup-config* and *running-config*. 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:		
the startup-configuration on the internal flash	copy running-config startup-config	
the usb flash on the IOM	copy running-config usbflash://filename	
an FTP server	copy running-config ftp:// username:password@{hostip hostname}/filepath/ filename	EXEC Privilege
a TFTP server	copy running-config tftp://{hostip hostname}/filepath/ filename	
an SCP server	copy running-config scp:// username:password@{hostip hostname}/filepath/ filename	



Note: When copying to a server, you can only use a hostname if a DNS server is configured.

View Files

You can only view file information and content on local file systems.

To view a list of files on the internal or external Flash, follow this step:

Step	Task	Command Syntax	Command Mode
1	View a list of files on:		
	the internal flash	dir flash:	EXEC Privilege
	the usbflash	dir usbflash:	

The output of the command `dir` also shows the read/write privileges, size (in bytes), and date of modification for each file (Figure 3-7).

Figure 3-7. Viewing a List of Files in the Internal Flash

```

FTOS#dir
Directory of flash:

 1  drwx      4096   Jan 01 1980 00:00:00 +00:00 .
 2  drwx      2048   May 10 2011 14:45:15 +00:00 ..
 3  drwx      4096   Feb 17 2011 00:28:00 +00:00 TRACE_LOG_DIR
 4  drwx      4096   Feb 17 2011 00:28:02 +00:00 CORE_DUMP_DIR
 5  d---      4096   Feb 17 2011 00:28:02 +00:00 ADMIN_DIR
 6  -rwx      1272   Apr 29 2011 16:15:14 +00:00 startup-config
 7  -rwx     10093   Feb 17 2011 20:48:02 +00:00 abhi-jan26.cfg
 8  -rwx     217155 Feb 22 2011 23:14:34 +00:00 show-tech-cfg.txt
 9  -rwx      5162   Mar 02 2011 04:02:58 +00:00 runn-feb6
10  -rwx     10507   Mar 03 2011 01:17:16 +00:00 abhi-feb7.cfg
11  -rwx         4   May 06 2011 22:05:06 +00:00 dhcpBindConflict
12  -rwx      6900   Feb 17 2011 04:43:12 +00:00 startup-config.bak
13  -rwx    1244038 Feb 13 2011 04:27:16 +00:00 f10cp_sysd_110213042625.acore.gz

flash: 2143281152 bytes total (2123755520 bytes free)
--More--

```

To view the contents of a file, follow this step:

Step	Task	Command Syntax	Command Mode
1	View the:		
	contents of a file in the internal flash	<code>show file flash://filename</code>	
	contents of a file in the usb flash	<code>show file usbflash://filename</code>	EXEC Privilege
	running-configuration	<code>show running-config</code>	
	startup-configuration	<code>show startup-config</code>	

View Configuration Files

Configuration files have three commented lines at the beginning of the file (Figure 3-8), to help you track the last time any user made a change to the file, which user made the change(s), and when the file was last saved to the startup-configuration.

In the running-configuration file, if there is a difference between the timestamp on the “Last configuration change,” and “Startup-config last updated,” you have made changes that have not been saved and will not be preserved upon a system reboot.

Figure 3-8. Tracking Changes with Configuration Comments

```

FTOS#show running-config
Current Configuration ...
Current Configuration ...
! Version E8-3-16-0
! Last configuration change at Tue Mar 6 11:51:50 2012 by default
! Startup-config last updated at Tue Mar 6 07:41:23 2012 by default
!
boot system stack-unit 5 primary tftp://10.11.200.241/dt-m1000e-3-a2
boot system stack-unit 5 secondary system: B:
boot system stack-unit 5 default tftp://10.11.200.241/dt-m1000e-3-b2
boot system gateway 10.11.209.254
--More--

```

File System Management

The Dell Force10 system can use the internal Flash, USB Flash, or remote devices to store files. The system stores files on the internal Flash by default, but you can configure it to store files elsewhere.

To view file system information:

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

The output of the show file-systems command (Figure 3-9) shows the total capacity, amount of free memory, file structure, media type, and read/write privileges for each storage device in use.

Figure 3-9. show file-systems Command Example

```

FTOS#show file-systems

Size(b)      Free(b)      Feature      Type  Flags  Prefixes
2143281152   2000785408   FAT32        USERFLASH  rw  flash:
15848660992   831594496   FAT32        USBFLASH   rw  usbflash:
-            -            -            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:

Task	Command Syntax	Command Mode
Change the default directory.	cd <i>directory</i>	EXEC Privilege

You can change the default storage location to the USB Flash (Figure 3-10). File management commands then apply to the USB Flash rather than the internal Flash.

Figure 3-10. Alternative Storage Location

```

FTOS#cd usbflash:
FTOS#copy running-config test ← No File System Specified
!
3998 bytes successfully copied

FTOS#dir
Directory of usbflash:

 1 drwx      4096   Jan 01 1980 00:00:00 +00:00 .
 2 drwx      2048   May 02 2012 07:05:06 +00:00 ..
 3 -rwx       1272  Apr 29 2011 16:15:14 +00:00 startup-config
 4 -rwx       3998  May 11 2011 23:36:12 +00:00 test ← File Saved to USB Flash

```

View the Command History

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 (Figure 3-11).

Figure 3-11. show command-history Command Example

```

FTOS#show command-history
[5/18 21:58:32]: CMD-(TEL0):[enable]by admin from vty0 (10.11.68.5)
[5/18 21:58:48]: CMD-(TEL0):[configure]by admin from vty0 (10.11.68.5)
  - Repeated 1 time.
[5/18 21:58:57]: CMD-(TEL0):[interface port-channel 1]by admin from vty0 (10.11.68.5)
[5/18 21:59:9]: CMD-(TEL0):[show config]by admin from vty0 (10.11.68.5)
[5/18 22:4:32]: CMD-(TEL0):[exit]by admin from vty0 (10.11.68.5)
[5/18 22:4:41]: CMD-(TEL0):[show interfaces port-channel brief]by admin from vty0
(10.11.68.5)

```

Upgrading and Downgrading FTOS



Note: To upgrade or downgrade FTOS, refer to the Release Notes for the version you want to load on the system.

Management

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

- [Configure Privilege Levels](#)
- [Configure Logging](#)
- [File Transfer Services](#)
- [Terminal Lines](#)
- [Lock CONFIGURATION Mode](#)
- [Recovering from a Forgotten Password](#)
- [Recovering from a Failed Start](#)

Configure Privilege Levels

Privilege levels restrict access to commands based on user or terminal line. There are 15 privilege levels, of which two are pre-defined. The default privilege level is 1.

- **Level 1**—Access to the system begins at EXEC mode, and EXEC mode commands are limited to very basic commands, some of which are enable, disable, and exit.
- **Level 15**—To access all commands, you must enter EXEC Privilege mode. Normally, you must enter a password to enter this mode.

Create a Custom Privilege Level

Custom privilege levels start with the default EXEC mode command set. You can then customize privilege levels 2-14 by:

- removing commands from the EXEC mode commands
- moving commands from EXEC Privilege mode to EXEC mode
- allowing access to CONFIGURATION mode commands
- allowing access to INTERFACE, LINE, ROUTE-MAP, and ROUTER mode commands

A user can access all commands at his privilege level and below.

Removing a Command from EXEC Mode

Remove a command from the list of available commands in EXEC mode for a specific privilege level using the `privilege exec command` from CONFIGURATION mode. In the command, specify a level *greater* than the level given to a user or terminal line, followed by the first keyword of each command to be restricted.

Move a Command from EXEC Privilege Mode to EXEC Mode

Move a command from EXEC Privilege to EXEC mode for a privilege level using the `privilege exec command` from CONFIGURATION mode. In the command, specify the privilege level of the user or terminal line, and specify *all* keywords in the command to which you want to allow access.

Allow Access to CONFIGURATION Mode Commands

Allow access to CONFIGURATION mode using the `privilege exec configure level level` command from CONFIGURATION mode. A user that enters CONFIGURATION mode remains at the same privilege level and has access to only two commands, `end` and `exit`. You must individually specify each CONFIGURATION mode command to which you want to allow access using the `privilege configure level level` command. In this command, specify the privilege level of the user or terminal line, and specify *all* keywords in the command to which you want to allow access.

Allow Access to INTERFACE, LINE, ROUTE-MAP, and ROUTER Mode

To allow access to INTERFACE, LINE, ROUTE-MAP, and ROUTER mode, follow these steps:

1. Similar to allowing access to CONFIGURATION mode, to allow access to INTERFACE, LINE, ROUTE-MAP, and ROUTER modes, you must first allow access to the command that enters you into the mode. For example, allow a user to enter INTERFACE mode using the `privilege configure level level interface tengigabitethernet` command.
2. Then, individually identify the INTERFACE, LINE, ROUTE-MAP, or ROUTER commands to which you want to allow access using the `privilege {interface | line | route-map | router} level level` command. In this command, specify the privilege level of the user or terminal line, and specify *all* keywords in the command to which you want to allow access.

The following table lists the configuration tasks you can use to customize a privilege level:

Task	Command Syntax	Command Mode
Remove a command from the list of available commands in EXEC mode.	<code>privilege exec level level {command ... command}</code>	CONFIGURATION
Move a command from EXEC Privilege to EXEC mode.	<code>privilege exec level level {command ... command}</code>	CONFIGURATION
Allow access to CONFIGURATION mode.	<code>privilege exec configure level level</code>	CONFIGURATION

Task	Command Syntax	Command Mode
Allow access to INTERFACE, LINE, ROUTE-MAP, and/or ROUTER mode. Specify <i>all</i> keywords in the command.	privilege configure level <i>level</i> {interface line route-map router} { <i>command-keyword</i> ... <i>command-keyword</i> }	CONFIGURATION
Allow access to a CONFIGURATION, INTERFACE, LINE, ROUTE-MAP, and/or ROUTER mode command.	privilege {configure interface line route-map router} level <i>level</i> { <i>command</i> ... <i>command</i> }	CONFIGURATION

The configuration in [Figure 4-1](#) creates privilege level 3. This level:

- removes the resequence command from EXEC mode by requiring a minimum of privilege level 4
- moves the capture bgp-pdu max-buffer-size command from EXEC Privilege to EXEC mode by requiring a minimum privilege level 3, which is the configured level for VTY 0
- allows access to CONFIGURATION mode with the banner command
- allows access to INTERFACE and LINE modes with the no command

Figure 4-1. Create a Custom Privilege Level Apply a Privilege Level to a Username

```

FTOS(conf)#do show run privilege
!
FTOS(conf)#privilege exec level 3 capture
FTOS(conf)#privilege exec level 3 configure
FTOS(conf)#privilege exec level 4 resequence
FTOS(conf)#privilege exec level 3 clear arp-cache
FTOS(conf)#privilege exec level 3 clear arp-cache max-buffer-size
FTOS(conf)#privilege configure level 3 line
FTOS(conf)#privilege configure level 3 interface
FTOS(conf)#do telnet 10.11.80.201
[telnet output omitted]
FTOS#show priv
Current privilege level is 3.
FTOS#?
capture          Capture packet
configure        Configuring from terminal
disable          Turn off privileged commands
enable           Turn on privileged commands
exit             Exit from the EXEC
ip               Global IP subcommands
monitor          Monitoring feature
ntrace           Trace reverse multicast path from destination to source
ping             Send echo messages
quit            Exit from the EXEC
show             Show running system information
[output omitted]
FTOS#config
[output omitted]
FTOS(conf)#do show priv
Current privilege level is 3.
FTOS(conf)#?
end              Exit from configuration mode
exit             Exit from configuration mode
interface        Select an interface to configure
FTOS(conf)#interface ?
loopback         Loopback interface
managementethernet Management Ethernet interface
null             Null interface
port-channel     Port-channel interface
range            Configure interface range
tengigabitethernet TenGigabit Ethernet interface
vlan             VLAN interface
FTOS(conf)#interface tengigabitethernet 1/1
FTOS(conf-if-te-1/1)#?
end              Exit from configuration mode
exit             Exit from interface configuration mode
FTOS(conf-if-te-1/1)#exit
FTOS(conf)#line ?
console          Primary terminal line
vty              Virtual terminal
FTOS(conf)#line vty 0
FTOS(conf-line-vty)#?
exit            Exit from line configuration mode
FTOS(conf-line-vty)#

```



To set a privilege level for a user:

Task	Command Syntax	Command Mode
Configure a privilege level for a user.	<code>username <i>username</i> privilege <i>level</i></code>	CONFIGURATION

Apply a Privilege Level to a Terminal Line

To set a privilege level for a terminal line:

Task	Command Syntax	Command Mode
Configure a privilege level for a terminal line.	<code>privilege level <i>level</i></code>	LINE

 **Note:** When you assign a privilege level between 2 and 15, access to the system begins at EXEC mode, but the prompt is `hostname#`, rather than `hostname>`.

Configure Logging

FTOS tracks changes in the system using event and error messages. By default, FTOS logs these messages on:

- the internal buffer
- console and terminal lines
- any configured syslog servers

Log Messages in the Internal Buffer

All error messages, except those beginning with `%BOOTUP (Message)`, are logged in the internal buffer.

Message 1 BootUp Events

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

Configuration Task List for System Log Management

The following sections include the configuration tasks for system log management:

- [Disable System Logging](#)
- [Send System Messages to a Syslog Server](#)

Disable System Logging

By default, logging is enabled and log messages are sent to the logging buffer, all terminal lines, console, and syslog servers.

To enable and disable system logging:

Task	Command Syntax	Command Mode
Disable all logging except on the console.	no logging on	CONFIGURATION
Disable logging to the logging buffer.	no logging buffer	CONFIGURATION
Disable logging to terminal lines.	no logging monitor	CONFIGURATION
Disable console logging.	no logging console	CONFIGURATION

Send System Messages to a Syslog Server

To send system messages to a syslog server:

Task	Command Syntax	Command Mode
Specify the server to which you want to send system messages. You can configure up to eight syslog servers.	logging {ip-address hostname}	CONFIGURATION

Configure a Unix System as a Syslog Server

Configure a UNIX system as a syslog server by adding the following lines to */etc/syslog.conf* on the Unix system and assigning write permissions to the file.

- on a 4.1 BSD UNIX system, add the line: `local7.debugging /var/log/log7.log`
- on a 5.7 SunOS UNIX system, add the line: `local7.debugging /var/adm/ftos.log`

In the lines above, `local7` is the logging facility level and `debugging` is the severity level.

Change System 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 the severity level of messages logged to a syslog server, use any or all of the following commands in CONFIGURATION mode:

Task	Command Syntax	Command Mode
Specify the minimum severity level for logging to the logging buffer.	logging buffered <i>level</i>	CONFIGURATION
Specify the minimum severity level for logging to the console.	logging console <i>level</i>	CONFIGURATION
Specify the minimum severity level for logging to terminal lines.	logging monitor <i>level</i>	CONFIGURATION
Specifying the minimum severity level for logging to a syslog server.	logging trap <i>level</i>	CONFIGURATION
Specify the minimum severity level for logging to the syslog history table.	logging history <i>level</i>	CONFIGURATION

Task	Command Syntax	Command Mode
Specify the size of the logging buffer. Note: When you decrease the buffer size, FTOS deletes all messages stored in the buffer. Increasing the buffer size does not affect messages in the buffer.	logging buffered <i>size</i>	CONFIGURATION
Specify the number of messages that FTOS saves to its logging history table.	logging history size <i>size</i>	CONFIGURATION

To view the logging buffer and configuration, enter the show logging command in EXEC privilege mode (Figure 4-2).

To view the logging configuration, enter the show running-config logging command in EXEC privilege mode (Figure 4-3).

Display the Logging Buffer and the Logging Configuration

To display the current contents of the logging buffer and the logging settings for the system, enter the show logging command in EXEC privilege mode (Figure 4-2).

Figure 4-2. show logging Command Example

```

FTOS#show logging
Syslog logging: enabled
  Console logging: level debugging
  Monitor logging: level debugging
  Buffer logging: level debugging, 58 Messages Logged, Size (40960 bytes)
  Trap logging: level informational
    Logging to 172.31.1.4
    Logging to 172.16.1.162
    Logging to 133.33.33.4
    Logging to 10.10.10.4
    Logging to 10.1.2.4
May 20 20:00:10: %STKUNIT0-M:CP %SYS-5-CONFIG_I: Configured from vty0 ( 10.11.68
.5 )by admin
May 20 19:57:45: %STKUNIT0-M:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable pas
sword authentication success on vty0 ( 10.11.68.5 )
May 20 19:57:40: %STKUNIT0-M:CP %SEC-5-LOGIN_SUCCESS: Login successful for user
admin on vty0 (10.11.68.5)
May 20 19:37:08: %STKUNIT0-M:CP %SEC-5-LOGOUT: Exec session is terminated for us
er admin on line vty0 (10.11.68.5)
May 20 18:59:36: %STKUNIT0-M:CP %SYS-5-CONFIG_I: Configured from vty0 ( 10.11.68
.5 )by admin
May 20 18:45:44: %STKUNIT0-M:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable pas
sword authentication success on vty0 ( 10.11.68.5 )
May 20 18:45:39: %STKUNIT0-M:CP %SEC-5-LOGIN_SUCCESS: Login successful for user
admin on vty0 (10.11.68.5)
May 20 17:18:08: %STKUNIT0-M:CP %SEC-5-LOGOUT: Exec session is terminated for us
er admin on line vty0 (10.11.68.5)
May 20 16:42:40: %STKUNIT0-M:CP %SYS-5-CONFIG_I: Configured from vty0 ( 10.11.68
.5 )by admin
- repeated 2 times
May 20 16:37:41: %STKUNIT0-M:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable pas
sword authentication success on vty0 ( 10.11.68.5 )
May 20 16:37:28: %STKUNIT0-M:CP %SEC-5-LOGIN_SUCCESS: Login successful for user
admin on vty0 (10.11.68.5)
May 20 16:37:17: %STKUNIT0-M:CP %SEC-5-LOGOUT: Exec session is terminated for us
er admin on line vty0 (10.11.68.5)
May 20 16:37:08: %STKUNIT0-M:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable pas
sword authentication success on vty0 ( 10.11.68.5 )
--More--

```

To view any changes made, use the show running-config logging command (Figure 4-3) in the EXEC privilege mode.

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 CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
logging facility [<i>facility-type</i>]	CONFIGURATION	<p>Specify one of the following parameters.</p> <ul style="list-style-type: none">• 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) <p>The default is local7.</p>

To view non-default settings, use the show running-config logging command (Figure 4-3) in EXEC mode.

Figure 4-3. show running-config logging Command Example

```
FTOS#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
FTOS#
```

Synchronize log messages

You can configure FTOS to filter and consolidate 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 CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	line {console 0 vty <i>number</i> [<i>end-number</i>]}	CONFIGURATION	Enter the LINE mode. Configure the following parameters for the virtual terminal lines: <ul style="list-style-type: none"> <i>number</i> range: zero (0) to 9. <i>end-number</i> range: 1 to 8. You can configure multiple virtual terminals at one time by entering a <i>number</i> and an <i>end-number</i> .
2	logging synchronous [level <i>severity-level</i> all] [<i>limit</i>]	LINE	Configure a level and set the maximum number of messages to be printed. Configure the following optional parameters: <ul style="list-style-type: none"> <i>level severity-level</i> range: 0 to 7. Default is 2. Use the all keyword to include all messages. <i>limit</i> range: 20 to 300. Default is 20.

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

Enable timestamp on Syslog Messages

By default, syslog messages 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 CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
<code>service timestamps [log debug] [datetime [localtime] [msec] [show-timezone] uptime]</code>	CONFIGURATION	Add timestamp to syslog messages. Specify the following optional parameters: <ul style="list-style-type: none">datetime: You can add the keyword <code>localtime</code> to include the <code>localtime</code>, <code>msec</code>, and <code>show-timezone</code>. If you do not add the keyword <code>localtime</code>, the time is UTC.uptime. To view time since last boot. If neither parameter is specified, FTOS configures <code>uptime</code>.

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

To disable time stamping on syslog messages, enter the `no service timestamps [log | debug]` command.

File Transfer Services

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

For more information about 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 the FTP Server](#) (mandatory)
- [Configure the FTP Server Parameters](#) (optional)
- [Configure FTP Client Parameters](#) (optional)

For a complete listing of FTP related commands, refer to [RFC 959, File Transfer Protocol](#).

Enable the FTP Server

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

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

To view the FTP configuration, enter the show running-config ftp command in EXEC privilege mode (Figure 4-4).

Figure 4-4. show running-config ftp Command Example

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

Configure the FTP Server Parameters

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

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

Command Syntax	Command Mode	Purpose
ftp-server topdir <i>dir</i>	CONFIGURATION	Specify the directory for users using FTP to reach the system. The default is the internal flash directory.
ftp-server username <i>username</i> password [<i>encryption-type</i>] <i>password</i>	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: <ul style="list-style-type: none"> <i>username</i>: Enter a text string <i>encryption-type</i>: Enter 0 for plain text or 7 for encrypted text. <i>password</i>: Enter a text string.



Note: You cannot use the change directory (cd) command until you configure ftp-server topdir.

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

Configure FTP Client Parameters

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

Command Syntax	Command Mode	Purpose
ip ftp source-interface <i>interface</i>	CONFIGURATION	Enter the following keywords and slot/port or number information: <ul style="list-style-type: none">• For a loopback interface, enter the keyword <code>loopback</code> followed by a number between 0 and 16383.• For a port channel interface, enter the keyword <code>port-channel</code> followed by a number from 1 to 128. For a 10-Gigabit Ethernet interface, enter the keyword <code>TenGigabitEthernet</code> followed by the slot/port information.• For a VLAN interface, enter the keyword <code>vlan</code> followed by a number from 1 to 4094.• For a 40-Gigabit Ethernet interface, enter the keyword <code>fortyGigE</code> followed by the slot/port information.
ip ftp password <i>password</i>	CONFIGURATION	Configure a password.
ip ftp username <i>name</i>	CONFIGURATION	Enter username to use on FTP client.

To view FTP configuration, use the `show running-config ftp` command in EXEC privilege mode (Figure 4-4).

Terminal Lines

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 virtual terminal lines (VTY) connect you through Telnet to the system.

Deny and Permit Access to a Terminal Line

Dell Force10 recommends applying only standard access control lists (ACLs) to deny and permit access to VTY lines.

- Layer 3 ACL deny all traffic that is not explicitly permitted, but in the case of VTY lines, an ACL with no rules does not deny any traffic.
- You cannot use the `show ip accounting access-list` command to display the contents of an ACL that is applied only to a VTY line.

To apply an IP ACL to a line:

Task	Command Syntax	Command Mode
Apply an ACL to a VTY line.	<code>ip access-class <i>access-list</i></code>	LINE

To view the configuration, enter the `show config` command in LINE mode (Figure 4-5).

Figure 4-5. Applying an Access List to a VTY Line

```

FTOS(conf-std-nacl)#show config
!
ip access-list standard myvtyacl
 seq 5 permit host 10.11.0.1
FTOS(conf-std-nacl)#line vty 0
FTOS(conf-line-vty)#show config
line vty 0
 access-class myvtyacl

```



FTOS Behavior: Prior to FTOS version 7.4.2.0, in order to deny access on a VTY line, you must apply an ACL and AAA authentication to the line. Then users are denied access only *after* they enter a username and password. Beginning in FTOS version 7.4.2.0, only an ACL is required, and users are denied access *before* they are prompted for a username and password.

Configure Login Authentication for Terminal Lines

You can use any combination of up to six authentication methods to authenticate a user on a terminal line. A combination of authentication methods is called a “method list”. If the user fails the first authentication method, FTOS prompts the next method until all methods are exhausted, at which point the connection is terminated. The available authentication methods are:

- **enable**—Prompt for the enable password.
- **line**—Prompt for the password you assigned to the terminal line. You must configure a password for the terminal line to which you assign a method list that contains the line authentication method. Configure a password using the password command from LINE mode.
- **local**—Prompt for the system username and password.
- **none**—Do not authenticate the user.
- **radius**—Prompt for a username and password and use a RADIUS server to authenticate.
- **tacacs+**—Prompt for a username and password and use a TACACS+ server to authenticate.

To configure authentication for a terminal line, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Create an authentication method list. You may use a mnemonic name or use the keyword default. The default authentication method for terminal lines is local, and the default method list is empty.	aaa authentication login { <i>method-list-name</i> default} [<i>method-1</i>] [<i>method-2</i>] [<i>method-3</i>] [<i>method-4</i>] [<i>method-5</i>] [<i>method-6</i>]	CONFIGURATION
2	Apply the method list from Step 1 to a terminal line.	login authentication { <i>method-list-name</i> default}	CONFIGURATION

Step	Task	Command Syntax	Command Mode
3	If you used the line authentication method in the method list you applied to the terminal line, configure a password for the terminal line.	password	LINE

VTY lines 0-2 use a single authentication method, line (Figure 4-6).

Figure 4-6. Configuring Login Authentication on a Terminal Line

```

FTOS(conf)#aaa authentication login myvtymethodlist line
FTOS(conf)#line vty 0 2
FTOS(conf-line-vty)#login authentication myvtymethodlist
FTOS(conf-line-vty)#password myvtypassword
FTOS(conf-line-vty)#show config
line vty 0
 password myvtypassword
 login authentication myvtymethodlist
line vty 1
 password myvtypassword
 login authentication myvtymethodlist
line vty 2
 password myvtypassword
 login authentication myvtymethodlist
FTOS(conf-line-vty)#

```

Time Out of EXEC Privilege Mode

EXEC timeout is a basic security feature that returns FTOS to EXEC mode after a period of inactivity on terminal lines.

To change the timeout period or disable EXEC timeout:

Task	Command Syntax	Command Mode
Set the number of minutes and seconds. Default: 10 minutes on console, 30 minutes on VTY. Disable EXEC timeout by setting the timeout period to 0.	exec-timeout <i>minutes</i> [<i>seconds</i>]	LINE
Return to the default timeout values.	no exec-timeout	LINE

To view the configuration, enter the show config command from LINE mode (Figure 4-7).

Figure 4-7. Configuring EXEC Timeout

```

FTOS(conf)#line con 0
FTOS(conf-line-console)#exec-timeout 0
FTOS(conf-line-console)#show config
line console 0
  exec-timeout 0 0
FTOS(conf-line-console)#

```

Telnet to Another Network Device

To telnet to another device ([Figure 4-8](#)):

Task	Command Syntax	Command Mode
Telnet to the stack-unit. You do not need to configure the management port on the stack-unit to be able to telnet to it.	telnet-peer-stack-unit	EXEC Privilege
Telnet to a device with an IPv4 address. If you do not enter an IP address, FTOS enters a Telnet dialog that prompts you for one. <ul style="list-style-type: none"> Enter an IPv4 address in dotted decimal format (A.B.C.D) 	telnet [<i>ip-address</i>]	EXEC Privilege

Figure 4-8. Telnet to Another Network Device

```

FTOS# telnet 10.11.80.203
Trying 10.11.80.203...
Connected to 10.11.80.203.
Exit character is '^]'.
Login:
Login: admin
Password:
FTOS>exit
FTOS#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.forcel0networks.com) (ttypl)
login: admin
FTOS#

```

Lock CONFIGURATION Mode

FTOS allows multiple users to make configurations at the same time. You can lock CONFIGURATION mode so that only one user can be in CONFIGURATION mode at any time ([Figure 4-9](#)).

You can set two types of locks: auto and manual.

- Set an auto-lock using the configuration mode `exclusive auto` command from CONFIGURATION mode. When you set an auto-lock, every time a user is in CONFIGURATION mode, all other users are denied access. This means that you can exit to EXEC Privilege mode, and re-enter CONFIGURATION mode, without having to set the lock again.
- Set a manual lock using the `configure terminal lock` command from CONFIGURATION mode. When you configure a manual lock, which is the default, you must enter this command each time you want to enter CONFIGURATION mode and deny access to others.

Figure 4-9. Locking CONFIGURATION mode

```
FTOS(conf)#configuration mode exclusive auto
BATMAN(conf)#exit
3d23h35m: %RPM0-P:CP %SYS-5-CONFIG_I: Configured from console by console

FTOS#config
! Locks configuration mode exclusively.
FTOS(conf)#
```

If another user attempts to enter CONFIGURATION mode while a lock is in place, [Message 1](#) appears on their terminal.

Message 1 CONFIGURATION Mode Locked Error

```
% Error: User "" on line console0 is in exclusive configuration mode
```

If *any* user is already in CONFIGURATION mode when a lock is placed, [Message 2](#) appears on their terminal.

Message 2 Cannot Lock CONFIGURATION Mode Error

```
% Error: Can't lock configuration mode exclusively since the following users are currently
configuring the system:
User "admin" on line vty1 ( 10.1.1.1 )
```



Note: CONFIGURATION mode lock corresponds to a VTY session, not a user. Therefore, if you configure a lock and then exit CONFIGURATION mode, and another user enters CONFIGURATION mode, when you attempt to re-enter CONFIGURATION mode, you are denied access even though *you* are the one that configured the lock.



Note: If your session times out and you return to EXEC mode, the CONFIGURATION mode lock is unconfigured.

Viewing the Configuration Lock Status

If you attempt to enter CONFIGURATION mode when another user has locked it, you may view which user has control of CONFIGURATION mode using the `show configuration lock` command from EXEC Privilege mode.

You can then send any user a message using the `send` command from EXEC Privilege mode. Alternatively you can clear any line using the `clear` command from EXEC Privilege mode. If you clear a console session, the user is returned to EXEC mode.

Recovering from a Forgotten Password

If you configure authentication for the console and you exit out of EXEC mode or your console session times out, you are prompted to re-enter the password.

If you forget your password, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Log onto the system using the console.		
2	Power-cycle the chassis by switching off all of the power modules and then switching them back on.		
3	Hit any key to abort the boot process. You enter uBoot i mme id at ely, as indicated by the => prompt.	hit any key	(during bootup)
4	Set the system parameters to ignore the startup configuration file when the system reloads.	setenv stconfigignore true	uBoot
5	To save the changes use the saveenv command.	saveenv	uBoot
6	Reload the system.	reset	uBoot
7	Copy startup-config.bak to the running config.	copy flash://startup-config.bak running-config	EXEC Privilege
8	Remove all authentication statements you might have for the console.	no authentication login no password	LINE
9	Save the running-config.	copy running-config startup-config	EXEC Privilege
10	Set the system parameters to use the startup configuration file when the system reloads.	setenv stconfigignore false	uBoot
11	Save the running-config.	copy running-config startup-config	EXEC Privilege

Recovering from a Forgotten Enable Password

If you forget the enable password, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Log onto the system via console.		

Step	Task	Command Syntax	Command Mode
2	Power-cycle the chassis by switching off all of the power modules and then switching them back on.		
3	Hit any key to abort the boot process. You enter uBoot immediately, as indicated by the => prompt.	hit any key	(during bootup)
4	Set the system parameters to ignore the enable password when the system reloads.	setenv enablepwdignore true	uBoot
5	Reload the system.	reset	uBoot
6	Configure a new enable password.	enable {secret password}	CONFIGURATION
7	Save the running-config to the startup-config.	copy running-config startup-config	EXEC Privilege

Recovering from a Failed Start

A system that does not start correctly might be attempting to boot from a corrupted FTOS image or from a mis-specified location. In that case, you can restart the system and interrupt the boot process to point the system to another boot location.

For more information about the `setenv` command, its supporting commands, and other commands that can help recover from a failed start, refer to the Boot User chapter in the *FTOS Command Line Reference for the MXL 10/40GbE Switch IO Module*.

To recover from a failed start using the `setenv` command, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Power-cycle the chassis (pull the power cord and reinsert it).		
2	Hit any key to abort the boot process. You enter uBoot immediately, as indicated by the => prompt.	hit any key	(during bootup)
3	Assign the new location to the FTOS image to be used when the system reloads.	setenv [primary_image f10boot location secondary_image f10boot location default_image f10boot location]	uBoot
4	Assign an IP address to the Management Ethernet interface.	setenv ipaddre address	uBoot
5	Assign an IP address as the default gateway for the system.	setenv gatewayip address	uBoot
6	Reload the system.	reset	uBoot

Access Control Lists (ACLs)

This chapter describes the access control lists (ACLs), prefix lists, and route-maps.

This chapter contains the following sections:

- [IP Access Control Lists \(ACLs\)](#)
- [IP Fragment Handling](#)
- [Configure a Standard IP ACL](#)
- [Configure an Extended IP ACL](#)
- [Configuring Layer 2 and Layer 3 ACLs on an Interface](#)
- [Assign an IP ACL to an Interface](#)
- [Configuring Ingress ACLs](#)
- [Configuring Egress ACLs](#)
- [IP Prefix Lists](#)
- [ACL Resequencing](#)
- [Route Maps](#)

Overview

At their simplest, ACLs, prefix lists, and route-maps permit or deny traffic based on MAC and/or IP addresses. This chapter describes implementing IP ACLs, IP prefix lists, and route-maps. For MAC ACLs, refer to [“Layer 2” on page 305](#).

An ACL is a filter containing some criteria to match (examine IP, transmission control protocol [TCP], or user datagram protocol [UDP] packets) and an action to take (permit or deny). ACLs 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 a packet matches a filter, the switch drops or forwards the packet based on the filter’s specified action. If the packet does not match any of the filters in the ACL, the packet is dropped (implicit deny).

The number of ACLs supported on a system depends on your CAM size. For more information, refer to [Content Addressable Memory \(CAM\)](#).

IP Access Control Lists (ACLs)

In the Dell Force10 switch/routers, 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 packets based on the following criteria:

- 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 more information about ACL supported options, refer to the *FTOS Command Reference Guide*.

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 assigns 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 append or delete new rules into an existing ACL (already written into the content addressable memory [CAM]) without disrupting traffic flow. Existing entries in CAM are shuffled to accommodate the new entries. Hot Lock ACLs are enabled by default and support both standard and extended ACLs.



Note: Hot Lock ACLs are supported for Ingress ACLs only.

Implementing ACLs on FTOS

You can assign one IP ACL per interface with FTOS. If an IP ACL is not assigned to an interface, it is not used by the software in any other capacity.

The number of entries allowed per ACL is hardware-dependent.

If counters are enabled on IP ACL rules that are already configured, those counters are reset when a new rule is inserted or prepended. If a rule is appended, the existing counters are not affected. This is applicable to the following features:

- L2 Ingress Access list
- L2 Egress Access list

- L3 Ingress Access list
- L3 Egress Access list



Note: IP ACLs are supported over VLANs in Version 6.2.1.1 and higher.

ACLs and VLANs

There are some differences when assigning ACLs to a VLAN rather than a physical port. For example, when using a single port-pipe, if you apply an ACL to a VLAN, one copy of the ACL entries gets installed in the ACL CAM on the port-pipe. The entry would look for the incoming VLAN in the packet. Whereas, if you apply an ACL on individual ports of a VLAN, separate copies of the ACL entries are installed for each port belonging to a port-pipe.

ACL Optimization

If an access list contains duplicate entries, FTOS deletes one of the entries to conserve CAM space.

Standard and extended ACLs take up the same amount of CAM space. A single ACL rule uses two CAM entries whether it is identified as a standard or extended ACL.

Determine the Order in Which ACLs are Used to Classify Traffic

When you link class-maps to queues using the `service-queue` command, FTOS matches the class-maps according to queue priority (queue numbers closer to 0 have lower priorities). For example, in [Figure 5-1](#), class-map *cmap2* is matched against ingress packets before *cmap1*.

ACLs *acl1* and *acl2* have overlapping rules because the address range 20.1.1.0/24 is within 20.0.0.0/8. Therefore, (without the keyword `order`) packets within the range 20.1.1.0/24 match positive against *cmap1* and are buffered in queue 7, though you intended for these packets to match positive against *cmap2* and be buffered in queue 4.

In cases such as these, where class-maps with overlapping ACL rules are applied to different queues, use the `order` keyword to specify the order in which you want to apply ACL rules ([Figure 5-1](#)). The order can range from 0 to 254. FTOS writes to the CAM ACL rules with lower order numbers (order numbers closer to 0) before rules with higher order numbers so that packets are matched as you intended. By default, all ACL rules have an order of 254.

Figure 5-1. Using the Order Keyword in ACLs

```

FTOS(conf)#ip access-list standard acl1
FTOS(conf-std-nacl)#permit 20.0.0.0/8
FTOS(conf-std-nacl)#exit
FTOS(conf)#ip access-list standard acl2
FTOS(conf-std-nacl)#permit 20.1.1.0/24 order 0
FTOS(conf-std-nacl)#exit
FTOS(conf)#class-map match-all cmap1
FTOS(conf-class-map)#match ip access-group acl1
FTOS(conf-class-map)#exit
FTOS(conf)#class-map match-all cmap2
FTOS(conf-class-map)#match ip access-group acl2
FTOS(conf-class-map)#exit
FTOS(conf)#policy-map-input pmap
FTOS(conf-policy-map-in)#service-queue 3 class-map cmap1
FTOS(conf-policy-map-in)#service-queue 1 class-map cmap2
FTOS(conf-policy-map-in)#exit
FTOS(conf)#interface tengig 1/0
FTOS(conf-if-ti-1/0)#service-policy input pmap

```

IP Fragment Handling

FTOS supports a configurable option to explicitly deny IP fragmented packets, particularly second and subsequent packets. It extends the existing ACL command syntax with the `fragments` keyword for all Layer 3 rules applicable to all Layer protocols (`permit/deny ip/tcp/udp/icmp`).

- Both standard and extended ACLs support IP fragments.
- Second and subsequent fragments are allowed because a Layer 4 rule cannot be applied to these fragments. If the packet is to be denied eventually, the first fragment would be denied and hence the packet as a whole cannot be reassembled.
- Implementing the required rules uses a significant number of CAM entries per TCP/UDP entry.
- For IP ACL, FTOS always applies implicit deny. You do not have to configure it.
- For IP ACL, FTOS applies implicit permit for second and subsequent fragments just prior to the implicit deny.
- If an *explicit* deny is configured, the second and subsequent fragments do not hit the implicit permit rule for fragments.

IP Fragments ACL Examples

The following configuration permits all packets (both fragmented & non-fragmented) with destination IP 10.1.1.1. The second rule does not get hit at all ([Figure 5-2](#)).

Figure 5-2. Permit All Packets

```

FTOS(conf)#ip access-list extended ABC
FTOS(conf-ext-nacl)#permit ip any 10.1.1.1/32
FTOS(conf-ext-nacl)#deny ip any 10.1.1.1./32 fragments
FTOS(conf-ext-nacl)

```

To deny second/subsequent fragments, use the same rules in a different order. These ACLs deny all second & subsequent fragments with destination IP 10.1.1.1 but permit the first fragment & non fragmented packets with destination IP 10.1.1.1 (Figure 5-3).

Figure 5-3. Deny Second Packets

```
FTOS(conf)#ip access-list extended ABC
FTOS(conf-ext-nacl)#deny ip any 10.1.1.1/32 fragments
FTOS(conf-ext-nacl)#permit ip any 10.1.1.1/32
FTOS(conf-ext-nacl)
```

Layer 4 ACL Rules Examples

In Figure 5-4, first fragments or non-fragmented TCP packets from 10.1.1.1 with TCP destination port equal to 24 are permitted. All other fragments are denied.

Figure 5-4. Layer 4 ACL Rules

```
FTOS(conf)#ip access-list extended ABC
FTOS(conf-ext-nacl)#permit tcp host 10.1.1.1 any eq 24
FTOS(conf-ext-nacl)#deny ip any any fragment
FTOS(conf-ext-nacl)
```

In (Figure 5-5), TCP packets that are first fragments or non-fragmented from host 10.1.1.1 with TCP destination port equal to 24 are permitted. Additionally, all TCP non-first fragments from host 10.1.1.1 are permitted. All other IP packets that are non-first fragments are denied.

Figure 5-5. TCP Packets

```
FTOS(conf)#ip access-list extended ABC
FTOS(conf-ext-nacl)#permit tcp host 10.1.1.1 any eq 24
FTOS(conf-ext-nacl)#permit tcp host 10.1.1.1 any fragment
FTOS(conf-ext-nacl)#deny ip any any fragment
FTOS(conf-ext-nacl)
```



Note the following when configuring ACLs with the fragments keyword.

When an ACL filters packets, it looks at the fragment offset (FO) to determine whether or not it is a fragment.

FO = 0 means it is either the first fragment or the packet is a non-fragment.

FO > 0 means it is dealing with the fragments of the original packet.

Permit ACL line with L3 information only and the fragments keyword is present:

If a packet's L3 information matches the L3 information in the ACL line, the packet's fragment offset (FO) is checked.

- If a packet's FO > 0, the packet is permitted.
- If a packet's FO = 0, the next ACL entry is processed.

Deny ACL line with L3 information only and the fragments keyword is present:

If a packet's L3 information matches the L3 information in the ACL line, the packet's fragment offset (FO) is checked.

- If a packet's FO > 0, the packet is denied.
- If a packet's FO = 0, the next ACL line is processed.

Configure a Standard IP ACL

To configure an ACL, use commands in IP ACCESS LIST mode and INTERFACE mode. For a complete listing of all commands related to IP ACLs, refer to the *FTOS Command Line Interface Reference Guide*.

To set up extended ACLs, refer to [Configure an Extended IP ACL](#).

A standard IP ACL uses the source IP address as its match criterion.

To configure a standard IP ACL, follow these steps:

Step	Command Syntax	Command Mode	Purpose
1	ip access-list standard <i>access-listname</i>	CONFIGURATION	Enter IP ACCESS LIST mode by naming a standard IP access list.
2	seq <i>sequence-number</i> {deny permit} { <i>source</i> [<i>mask</i>] any host <i>ip-address</i> } [count [byte]] [order] [fragments]	CONFIG-STD-NACL	Configure a drop or forward filter.



Note: 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.

To view the rules of a particular ACL configured on a particular interface, use the show ip accounting access-list *ACL-name* interface *interface* command in EXEC Privilege mode ([Figure 5-6](#)).

Figure 5-6. Command Example: show ip accounting access-list

```
FTOS#show ip accounting access ToOspf interface tengig 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
FTOS#
```

Figure 5-7 shows 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.

Figure 5-7. Command example: seq

```
FTOS(conf-std-nacl)#seq 25 deny ip host 10.5.0.0 any
FTOS(conf-std-nacl)#seq 15 permit tcp 10.3.0.0 /16 any
FTOS(conf-std-nacl)#show config
!
ip access-list standard dilling
  seq 15 permit tcp 10.3.0.0/16 any
FTOS(conf-std-nacl)#
```

To delete a filter, use the no seq *sequence-number* command in 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, follow these steps starting in CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	ip access-list standard <i>access-list-name</i>	CONFIGURATION	Create a standard IP ACL and assign it a unique name.
2	{deny permit} {source [<i>mask</i>] any host <i>ip-address</i> } [count [byte]] [order] [fragments]	CONFIG-STD-NACL	Configure a drop or forward IP ACL filter.

Figure 5-8 shows a standard IP 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 IP ACCESS LIST mode displays the two filters with the sequence numbers 5 and 10.

Figure 5-8. Standard IP ACL

```
FTOS(conf-route-map)#ip access standard kigali
FTOS(conf-std-nacl)#permit 10.1.0.0/16
FTOS(conf-std-nacl)#show config
!
ip access-list standard kigali
  seq 5 permit 10.1.0.0/16
FTOS(conf-std-nacl)#
```

To view all configured IP ACLs, use the show ip accounting access-list command in EXEC Privilege mode (Figure 5-5).

Figure 5-9. Command Example: show ip accounting access-list

```
FTOS#show ip accounting access example interface tengig 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
```

To delete a filter, enter the show config command in 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 IP ACCESS LIST mode.

Configure an Extended IP ACL

Extended IP ACLs filter based on source and destination IP addresses, IP host addresses, TCP addresses, TCP host addresses, UDP addresses, and UDP host addresses.

Because traffic passes through the filter in the order of the filter's sequence, you can configure the extended IP ACL by first entering IP ACCESS LIST mode and then assigning a sequence number to the filter.

Configure Filters with a Sequence Number

To create a filter for packets with a specified sequence number, follow these steps, starting in CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	<code>ip access-list extended access-list-name</code>	CONFIGURATION	Enter the IP ACCESS LIST mode by creating an extended IP ACL.
2	<code>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]] [order] [fragments]</code>	CONFIG-EXT-NAACL	Configure a drop or forward filter.

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, you might need to insert a new filter. To prevent reconfiguring multiple filters, assign sequence numbers in multiples of five or another number.

Figure 5-10 shows 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.

Figure 5-10. Command Example: seq

```
FTOS(conf-ext-nacl)#seq 15 deny ip host 112.45.0.0 any
FTOS(conf-ext-nacl)#seq 5 permit tcp 12.1.3.45 0.0.255.255 any
FTOS(conf-ext-nacl)#show config
!
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
FTOS(conf-ext-nacl)#
```

Configure Filters Without a Sequence Number

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 IP ACCESS LIST mode:

Command Syntax	Command Mode	Purpose
{deny permit} {source mask any host ip-address} [count [byte]] [order] [fragments]	CONFIG-EXT-NACL	Configure a deny or permit filter to examine IP packets.
{deny permit} tcp {source mask} any host ip-address}} [count [byte]] [order] [fragments]	CONFIG-EXT-NACL	Configure a deny or permit filter to examine TCP packets.
{deny permit} udp {source mask any host ip-address}} [count [byte]] [order] [fragments]	CONFIG-EXT-NACL	Configure a deny or permit filter to examine UDP packets.

Figure 5-11 shows an extended IP ACL 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 IP ACCESS LIST mode displays the two filters with the sequence numbers 5 and 10.

Figure 5-11. Extended IP ACL

```
FTOS(conf-ext-nacl)#deny tcp host 123.55.34.0 any
FTOS(conf-ext-nacl)#permit udp 154.44.123.34 0.0.255.255 host 34.6.0.0
FTOS(conf-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
FTOS(conf-ext-nacl)#
```

To view all configured IP ACLs and the number of packets processed through the ACL, use the show ip accounting access-list command in EXEC Privilege mode (Figure 5-6).

Established Flag

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

You can configure both Layer 2 and Layer 3 ACLs 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 FTOS are governed by the L3 ACL only because they are not filtered against an L2 ACL.
- The packets switched by FTOS are first filtered by the L3 ACL, then by the L2 ACL.
- When packets are switched by FTOS, the egress L3 ACL does not filter the packet.


For the following features, if you enable counters on rules that have already been configured and 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
- L3 Ingress Access list

If a rule is simply appended, existing counters are not affected.

Table 5-1. 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 L3 ACL
Permit	Permit	Permitted by L3 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 MAC ACLs, refer to [Layer 2 on page 305](#).

Assign an IP ACL to an Interface

To pass traffic through a configured IP ACL, you must assign that ACL to a physical interface, a port channel interface, or a VLAN. 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.

You can apply the same ACL 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.

For more information about Layer-3 interfaces, refer to [Interfaces](#).

To apply an IP ACL (standard or extended) to a physical or port channel interface, follow these steps, in INTERFACE mode:

Step	Command Syntax	Command Mode	Purpose
1	interface interface slot/port	CONFIGURATION	Enter the interface number.
2	ip address <i>ip-address</i>	INTERFACE	Configure an IP address for the interface, placing it in Layer-3 mode.
3	ip access-group <i>access-list-name</i> {in out} [implicit-permit] [vlan <i>vlan-range</i>]	INTERFACE	Apply an IP ACL to traffic entering or exiting an interface. out: configure the ACL to filter outgoing traffic. <ul style="list-style-type: none"> • Note: The number of entries allowed per ACL is hardware-dependent.
4	ip access-list [standard extended] <i>name</i>	INTERFACE	Apply rules to the new ACL.

To view which IP ACL is applied to an interface, use the show config command ([Figure 5-12](#)) in INTERFACE mode or the show running-config command in EXEC mode.

Figure 5-12. Command example: show config command in the INTERFACE Mode

```
FTOS(conf-if)#show conf
!
interface TenGigabitEthernet 0/0
 ip address 10.2.1.100 255.255.255.0
 ip access-group nimule in
 no shutdown
FTOS(conf-if)#
```

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. In the MXL Switch, either count (packets) or count (bytes) can be configured. However, for an ACL with multiple rules, you can configure some ACLs with count (packets) and others as count (bytes) at any given time.

To view the number of packets matching an ACL that is applied to an interface, follow these steps:

Step	Task
1	Create an ACL that uses rules with the count option. Refer to Configure a Standard IP ACL
2	Apply the ACL as an inbound or outbound ACL on an interface. Refer to Assign an IP ACL to an Interface
3	View the number of packets matching the ACL using the show ip accounting access-list command 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 ACL, use the ip access-group command in EXEC Privilege mode (Figure 5-13). This example also shows applying the ACL, applying rules to the newly created access group, and viewing the access list.

Figure 5-13. Creating an Ingress ACL

```
FTOS(conf)#interface tengig 0/0
FTOS(conf-if-tengig0/0)#ip access-group abcd in
FTOS(conf-if-tengig0/0)#show config
!
tengigetherenet 0/0
no ip address
ip access-group abcd in
no shutdown
FTOS(conf-if-tengig0/0)#end
FTOS#configure terminal
FTOS(conf)#ip access-list extended abcd
FTOS(conf-ext-nacl)#permit tcp any any
FTOS(conf-ext-nacl)#deny icmp any any
FTOS(conf-ext-nacl)#permit 1.1.1.2
FTOS(conf-ext-nacl)#end
FTOS#show ip accounting access-list
!
Extended Ingress IP access list abcd on tengigetherenet 0/0
seq 5 permit tcp any any
seq 10 deny icmp any any
seq 15 permit 1.1.1.2
```

Use the "in" keyword to specify ingress.

Begin applying rules to the ACL named "abcd."

View the access-list.

Configuring Egress ACLs

Configuring egress ACLs onto physical interfaces protects the system infrastructure from attack—malicious and incidental—by explicitly 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.

Use an egress ACL when you would like to restrict egress traffic. For example, when a DOS attack traffic is isolated to one particular interface, you 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 in EXEC Privilege mode (Figure 5-14). This example also shows viewing the configuration, applying rules to the newly created access group, and viewing the access list.

Figure 5-14. Creating an Egress ACL

```

FTOS(conf)#interface tengig 0/0
FTOS(conf-if-tengig0/0)#ip access-group abcd out
FTOS(conf-if-tengig0/0)#show config
!
tengigetherenet 0/0
no ip address
ip access-group abcd out
no shutdown
FTOS(conf-if-tengig0/0)#end
FTOS#configure terminal
FTOS(conf)#ip access-list extended abcd
FTOS(conf-ext-nacl)#permit tcp any any
FTOS(conf-ext-nacl)#deny icmp any any
FTOS(conf-ext-nacl)#permit 1.1.1.2
FTOS(conf-ext-nacl)#end
FTOS#show ip accounting access-list
!
Extended Ingress IP access list abcd on tengigetherenet 0/0
seq 5 permit tcp any any
seq 10 deny icmp any any
seq 15 permit 1.1.1.2

```

Use the “out” keyword to specify egress.

Begin applying rules to the ACL named “abcd.”

View the access-list.

Egress Layer 3 ACL Lookup for Control-Plane IP Traffic

By default, packets originated from the system are not filtered by egress ACLs. If you initiate a ping session from the system, for example, and apply an egress ACL to block this type of traffic on the interface, the ACL does not affect that ping traffic.

The Control Plane Egress Layer 3 ACL feature enhances IP reachability debugging by implementing control-plane ACLs for CPU-generated and CPU-forwarded traffic. Using permit rules with the count option, you can track on a per-flow basis whether CPU-generated and CPU-forwarded packets were transmitted successfully..

Task	Command Syntax	Command Mode
Apply Egress ACLs to IPv4 system traffic.	<code>ip control-plane [egress filter]</code>	CONFIGURATION
Create a Layer 3 ACL using permit rules with the count option to describe the desired CPU traffic	<code>permit ip {source mask any host ip-address} {destination mask any host ip-address} count</code>	CONFIG-NACL



FTOS Behavior: VRRP hellos and IGMP packets are not affected when you enable egress ACL filtering for CPU traffic. Packets sent by the CPU with the source address as the VRRP virtual IP address have the interface MAC address instead of VRRP virtual MAC address.

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.
- After 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, router information protocol [RIP], open shortest path first [OSPF], and border gateway protocol [BGP]).



Note: The MXL Switch platform does not support all protocols. It is important to know which protocol you are supporting prior to implementing prefix lists.

Configuration Task List for Prefix Lists

To configure a prefix list, you must use commands in PREFIX LIST, ROUTER RIP, ROUTER OSPF, and ROUTER BGP modes. You create the prefix list in PREFIX LIST mode, and assign that list to commands in ROUTER RIP, ROUTER OSPF, and ROUTER BGP modes.

The following list includes the configuration tasks for prefix lists:

- [Configure a Prefix List](#)
- [Use a Prefix List for Route Redistribution](#)

For a complete listing of all commands related to prefix lists, refer to the *FTOS Command Line Interface Reference Guide*.

Configure a Prefix List

To configure a prefix list, follow these steps, starting in CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	<code>ip prefix-list <i>prefix-name</i></code>	CONFIGURATION	Create a prefix list and assign it a unique name. You are in PREFIX LIST mode.
2	<code>seq <i>sequence-number</i> {deny permit} <i>ip-prefix</i> [<i>ge min-prefix-length</i>] [<i>le max-prefix-length</i>]</code>	CONFIG-NPREFIXL	Create a prefix list with a sequence number and a deny or permit action. The optional parameters are: <ul style="list-style-type: none"> • <i>ge min-prefix-length</i>: is the minimum prefix length to be matched (0 to 32). • <i>le max-prefix-length</i>: 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 must be the last filter in your prefix list. To permit the default route only, enter permit 0.0.0.0/0.

Figure 5-15 shows 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.

Figure 5-15. Command Example: seq

```

FTOS(conf-nprefixl)#seq 20 permit 0.0.0.0/0 le 32
FTOS(conf-nprefixl)#seq 12 deny 134.23.0.0 /16
FTOS(conf-nprefixl)#seq 15 deny 120.23.14.0 /8 le 16
FTOS(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
FTOS(conf-nprefixl)#

```

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 PREFIX LIST mode.

If you are creating a standard prefix 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 five.

To configure a filter without a specified sequence number, follow these steps, starting in CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	ip prefix-list <i>prefix-name</i>	CONFIGURATION	Create a prefix list and assign it a unique name.
2	{deny permit} <i>ip-prefix</i> [ge <i>min-prefix-length</i>] le <i>max-prefix-length</i>]	CONFIG-NPREFIXL	Create a prefix list filter with a deny or permit action. The optional parameters are: <ul style="list-style-type: none"> ge <i>min-prefix-length</i>: is the minimum prefix length to be matched (0 to 32). le <i>max-prefix-length</i>: is the maximum prefix length to be matched (0 to 32).

Figure 5-16 shows a prefix 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 PREFIX LIST mode displays the two filters with the sequence numbers 5 and 10.

Figure 5-16. Prefix List

```
FTOS(conf-nprefix1)#permit 123.23.0.0 /16
FTOS(conf-nprefix1)#deny 133.24.56.0 /8
FTOS(conf-nprefix1)#show conf
!
ip prefix-list awe
  seq 5 permit 123.23.0.0/16
  seq 10 deny 133.0.0.0/8
FTOS(conf-nprefix1)#
```

To delete a filter, enter the show config command in PREFIX LIST mode and locate the sequence number of the filter you want to delete; then use the no seq *sequence-number* command in PREFIX LIST mode.

To view all configured prefix lists, use either of the following commands in EXEC mode (Figure 5-17) and (Figure 5-18):

Command Syntax	Command Mode	Purpose
show ip prefix-list detail [<i>prefix-name</i>]	EXEC Privilege	Show detailed information about configured Prefix lists.
show ip prefix-list summary [<i>prefix-name</i>]	EXEC Privilege	Show a table of summarized information about configured Prefix lists.

Figure 5-17. Command example: show ip prefix-list detail

```
FTOS>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)
FTOS>
```

Figure 5-18. Command Example: show ip prefix-list summary

```
FTOS>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
FTOS>
```

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, use either of the following commands in ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
<code>router rip</code>	CONFIGURATION	Enter RIP mode
<code>distribute-list <i>prefix-list-name</i> in [<i>interface</i>]</code>	CONFIG-ROUTER-RIP	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.
<code>distribute-list <i>prefix-list-name</i> out [<i>interface</i> connected static ospf]</code>	CONFIG-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 5-19) or the `show running-config rip` command in EXEC mode.

Figure 5-19. Command Example: show config in ROUTER RIP Mode

```
FTOS(conf-router_rip)#show config
!
router rip
  distribute-list prefix juba out
  network 10.0.0.0
FTOS(conf-router_rip)#router ospf 34
```

To apply a filter to routes in OSPF, use either of the following commands in ROUTER OSPF mode:

Command Syntax	Command Mode	Purpose
<code>router ospf</code>	CONFIGURATION	Enter OSPF mode
<code>distribute-list <i>prefix-list-name</i> in [<i>interface</i>]</code>	CONFIG-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.
<code>distribute-list <i>prefix-list-name</i> out [connected rip static]</code>	CONFIG-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 5-20) or the `show running-config ospf` command in EXEC mode.

Figure 5-20. Command Example: show config in ROUTER OSPF Mode

```
FTOS(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
FTOS(conf-router_ospf)#
```

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 resequencing whenever there is no longer an opportunity to order new rules as desired using current numbering scheme.

For example, Table 5-2 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 5-3. In the same example, apply resequencing if more than two rules must be placed between rules 7 and 10.

You can resequence IPv4 ACLs, prefix lists, and MAC ACLs. If resequencing, no CAM writes happen as a result, so there is no packet loss (the behavior is similar to Hot-lock ACLs).



Note: ACL resequencing does not affect the rules or remarks or the order in which they are applied. It merely rennumbers them so that new rules can be placed within the list as desired.

Table 5-2. 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 5-3. 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 ACLs, prefix lists, and MAC ACLs. To resequence an ACL or prefix list, use the appropriate command in [Table 5-4](#). When using these commands, you must specify the list name, starting number, and increment.

Table 5-4. Resequencing ACLs and Prefix Lists

List	Command	Command Mode
IPv4 or MAC ACL	<code>resequence access-list {ipv4 mac} {access-list-name StartingSeqNum Step-to-Increment}</code>	Exec
IPv4 prefix-list	<code>resequence prefix-list {ipv4} {prefix-list-name StartingSeqNum Step-to-Increment}</code>	Exec

[Figure 5-21](#) shows the resequencing of an IPv4 access-list beginning with the number 2 and incrementing by 2.

Figure 5-21. Resequencing ACLs

```
FTOS(conf-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
FTOS# end
FTOS# resequence access-list ipv4 test 2 2
FTOS# 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 continue to have the same sequence number after the resequence command is applied. Remarks that do not have a corresponding rule are incremented as a rule. These two mechanisms allow remarks to retain their original position in the list.

For example, in [Figure 5-22](#), 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 5-22. Resequencing Remarks

```

FTOS(conf-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
FTOS# end
FTOS# resequence access-list ipv4 test 2 2
FTOS# 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

Similar to 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

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 ROUTE-MAP mode and apply them in various commands in ROUTER RIP and ROUTER OSPF modes.

The following list includes the configuration tasks for route maps:

- [Create a Route Map](#) (mandatory)
- [Configure Route Map Filters](#) (optional)
- [Configure a Route Map for Route Redistribution](#) (optional)
- [Configure a Route Map for Route Tagging](#) (optional)

Create a Route Map

Route maps, ACLs, and prefix lists are similar in composition because all three contain filters, but route map filters 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 ROUTE-MAP mode, follow these steps, starting in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
<code>route-map <i>map-name</i> [permit deny] [<i>sequence-number</i>]</code>	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 you use the keyword deny in configuring a route map, routes that meet the match filters are not redistributed.

To view the configuration, use the show config command in ROUTE-MAP mode ([Figure 5-23](#)).

Figure 5-23. Command Example: show config in the ROUTE-MAP Mode

```
FTOS(conf-route-map)#show config
!
route-map dilling permit 10
FTOS(conf-route-map)#
```

You can create multiple instances of this route map 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, such as redistribute, traffic passes through all instances of that route map until a match is found. [Figure 5-24](#) shows an example with two instances of a route map.

Figure 5-24. Command Example: show route-map with Multiple Instances of a Route Map

```
FTOS#show route-map
route-map zakho, permit, sequence 10
Match clauses:
Set clauses:
route-map zakho, permit, sequence 20
Match clauses:
interface TenGigabitEthernet 0/1
Set clauses:
tag 35
level stub-area
FTOS#
```

Route map zakho has two instances

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 5-25](#)).

Figure 5-25. Deleting One Instance of a Route Map

```
FTOS(conf)#no route-map zakho 10
FTOS(conf)#end
FTOS#show route-map
route-map zakho, permit, sequence 20
Match clauses:
interface TenGigabitEthernet 0/1
Set clauses:
tag 35
level stub-area
FTOS#
```

[Figure 5-26](#) 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.

Figure 5-26. Command Example: show route-map

```
FTOS#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
FTOS#
```

To delete a route map, use the `no route-map map-name` command in CONFIGURATION mode.

Configure Route Map Filters

Within ROUTE-MAP mode, there are match and set commands. match commands search for a certain criterion in the routes and set commands change the characteristics of those routes, either by adding something or by specifying a level.

When there are multiple match commands of the same parameter under one instance of a route-map, FTOS does a match between either of those match commands. If there are multiple match commands of different parameter, FTOS does a match **ONLY** if there is a match among **ALL** match commands. Refer to the following examples:

Example 1

```
FTOS(conf)#route-map force permit 10
FTOS(conf-route-map)#match tag 1000
FTOS(conf-route-map)#match tag 2000
FTOS(conf-route-map)#match tag 3000
```

In the above route-map, if a route has any of the tag value specified in the match commands, there is a match.

Example 2

```
FTOS(conf)#route-map force permit 10
FTOS(conf-route-map)#match tag 1000
FTOS(conf-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 there is a match.

Also, if there are different instances of the same route-map, it is sufficient if a permit match happens in *any* instance of that route-map, for example:

```
FTOS(conf)#route-map force permit 10
FTOS(conf-route-map)#match tag 1000

FTOS(conf)#route-map force deny 20
FTOS(conf-route-map)#match tag 1000

FTOS(conf)#route-map force deny 30
FTOS(conf-route-map)#match tag 1000
```

In the above route-map, instance 10 permits the route having a tag value of 1000 and instances 20 and 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 ROUTE-MAP mode:

Command Syntax	Command Mode	Purpose
match interface <i>interface</i>	CONFIG-ROUTE-MAP	Match routes whose next hop is a specific interface. The parameters are: <ul style="list-style-type: none"> For a loopback interface, enter the keyword loopback followed by a number between zero (0) and 16383. 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. For a 40-Gigabit Ethernet interface, enter the keyword fortyGigE followed by the slot/port information.
match ip address <i>prefix-list-name</i>	CONFIG-ROUTE-MAP	Match destination routes specified in a prefix list (IPv4).
match ip next-hop { <i>access-list-name</i> <i>prefix-list prefix-list-name</i> }	CONFIG-ROUTE-MAP	Match next-hop routes specified in a prefix list (IPv4).
match ip route-source { <i>access-list-name</i> <i>prefix-list prefix-list-name</i> }	CONFIG-ROUTE-MAP	Match source routes specified in a prefix list (IPv4).
match metric <i>metric-value</i>	CONFIG-ROUTE-MAP	Match routes with a specific value.
match route-type {external [type-1 type-2] internal level-1 level-2 local }	CONFIG-ROUTE-MAP	Match routes specified as internal or external to OSPF, or locally generated.
match tag <i>tag-value</i>	CONFIG-ROUTE-MAP	Match routes with a specific tag.

To configure a set condition, use any or all of the following commands in ROUTE-MAP mode:

Command Syntax	Command Mode	Purpose
set automatic-tag	CONFIG-ROUTE-MAP	Generate a tag to be added to redistributed routes.
set level {backbone level-1 level-1-2 level-2 stub-area }	CONFIG-ROUTE-MAP	Specify an OSPF area for redistributed routes.
set metric {+ - <i>metric-value</i> }	CONFIG-ROUTE-MAP	Specify a value for redistributed routes.
set metric-type {external internal type-1 type-2}	CONFIG-ROUTE-MAP	Specify an OSPF type for redistributed routes.
set next-hop <i>ip-address</i>	CONFIG-ROUTE-MAP	Assign an IP address as the route's next hop.
set tag <i>tag-value</i>	CONFIG-ROUTE-MAP	Specify a tag for the redistributed routes.

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, you must call or include that route map in a command such as the redistribute or default-information originate commands in OSPF 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 the route tag. Use the redistribute command in OSPF, RIP, 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 5-27](#), the redistribute command calls the route map `static ospf` to redistribute only certain static routes into OSPF. According to the route map `static ospf`, only routes that have a next hop of TenGigabitEthernet interface 0/0 and that have a metric of 255 are redistributed into the OSPF backbone area.



Note: When re-distributing routes using route-maps, you 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.

Figure 5-27. Route Redistribution into OSPF

```
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 TenGigabitEthernet 0/0
  match metric 255
  set level backbone
```

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. Use this tag when the route leaves a routing domain to redistribute those routes again.

In [Figure 5-28](#), the redistribute ospf command with a route map is used in ROUTER RIP mode to apply a tag of 34 to all internal OSPF routes that are redistributed into RIP.

Figure 5-28. Tagging OSPF Routes Entering a RIP Routing Domain

```
!
router rip
  redistribute ospf 34 metric 1 route-map torip
!
route-map torip permit 10
  match route-type internal
  set tag 34
!
```

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 you configure the continue command at the end of a module, the next module (or a specified module) is processed even after a match is found. [Figure 5-29](#) 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 are processed.



Note: If you configure the continue clause without specifying a module, the next sequential module is processed.

Figure 5-29. Command Example: continue

```
!
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!
```


Bare Metal Provisioning (BMP)

Bare metal provisioning (BMP) improves accessibility to the MXL 10/40GbE Switch IO Module system. BMP performs auto configuration using a configuration file and an approved version of FTOS from a network source. BMP not only allows you to configure a stack with a minimum of effort, but it is also useful for quick configuration of a standalone system.

BMP eases configuration in the following key areas:

- Obtaining an IP address, running configuration and boot image information from a dynamic host configuration protocol (DHCP) server.
- Allowing access to the system through an Ethernet management port and data ports with or without DHCP-based dynamic IP address configuration of the user device. This does not stop the BMP process.
- Booting up in Layer 3 mode with interfaces already in no shutdown mode. Only Management mode is in no shutdown mode and the ip address dhcp is enabled, the front end ports are in shut mode. You can configure the username root password if the configuration file is not received.

Overview

The system is configured for BMP mode when it leaves the factory. The system retrieves a boot file and configuration information from a DHCP server and downloads that information from the trivial file transfer protocol (TFTP) server. Each system has its own configuration file and is automatically connected to the management network.



Note: PRIOR to booting the system, you must update the **dhcp.conf** file on the DHCP server you use. This enables the system to connect to the DHCP server and download the correct configuration and boot files.

There are two reload modes available in BMP:

- **BMP mode:** As the system boots up, it connects to a DHCP server containing the required FTOS and start-up configuration files. These files are downloaded to the system and the system is reloaded with those images.
- **Normal mode:** The system loads the FTOS image on the flash and boots up with the start-up configuration files on the flash. New configurations require that the Management IP and Management Interface IP addresses be configured manually. This mode is implemented with the standard reload command (with no additional parameters entered).

Use reload mode to boot up, the system remains in the system memory. If the system undergoes an automatic reload, it reloads using the previously used mode. To use a different mode when the system reloads automatically, reboot the system in a new mode. The new mode is then retained in system memory. To view the current reload mode, use the show reload type or show bootvar command (Figure 6-1) and (Figure 6-2).

Figure 6-1. Command Example: show reload type (normal)

```
FTOS#show reload-type
Reload-Type      :   normal-reload [Next boot : normal-reload]

FTOS#show bootvar
PRIMARY IMAGE FILE =  tftp://10.11.9.3/WJ_m1000e-2-c2
SECONDARY IMAGE FILE =  variable does not exist
DEFAULT IMAGE FILE =  variable does not exist
LOCAL CONFIG FILE =  variable does not exist
PRIMARY HOST CONFIG FILE =  variable does not exist
SECONDARY HOST CONFIG FILE =  variable does not exist
PRIMARY NETWORK CONFIG FILE =  variable does not exist
SECONDARY NETWORK CONFIG FILE =  variable does not exist
CURRENT IMAGE FILE =  tftp://10.11.9.3/WJ_m1000e-2-c2
CURRENT CONFIG FILE 1 =  flash://startup-config
CURRENT CONFIG FILE 2 =  variable does not exist
CONFIG LOAD PREFERENCE =  local first
BOOT INTERFACE GATEWAY IP ADDRESS =  10.11.9.254
Reload Mode =  normal-reload
```

Figure 6-2. Command Example: show reload type (jumpstart)

```
FTOS#show reload-type
Reload-Type      :   jump-start [Next boot :jump-start]
auto-save       :   disable
config-download :   enable
dhcp-timeout    :   50
retry-count     :   1
FTOS#

FTOS#show boot jumpstart
Config download enabled via DHCP/BOOTP
Autoconfig State : Autoconfig process has started
Autoconfig State : Waiting for boot options
FTOS#
```


Auto-Configuration

The system boot status is output to the console as the reload progresses. The messages include connections to the servers, assigned IP addresses and gateways, and the success or failure of those connections.

- [BMP Mode](#)
 - [MAC-Based IP Assignment](#)
 - [DHCP Configuration](#)
 - [IP Server](#)
 - [Domain Name Server](#)
 - [Boot Commands](#)
 - [System Boot and Set-Up Behavior](#)

BMP Mode

BMP mode is the boot mode configured for a new system arriving from Dell Force10. This mode obtains the FTOS image and configuration file from a network source (a DHCP server).

Before implementing this mode, you must set up a DHCP server and an IP server. The necessary FTOS image and start-up configuration files must be located on the server for the system to retrieve.

The DNS Server is not required, but it is recommended.

MAC-Based IP Assignment

One way to use BMP mode most efficiently is to configure the DHCP server to assign a fixed IP address and configuration file based on the management/front end MAC address. When this is done, the same IP address is assigned to the system even on repetitive reloads and the same configuration file is retrieved when using the DNS server or the **network-config** file to determine the hostname.

A dynamic IP address assignment may create a situation where the intended configuration is not applied to the system because the IP address can vary every time the system is reloaded.

The following is the configuration to be included in the **dhcp.conf** file so that the MAC-based IP and configuration file assignment are fixed:

Parameter Example	Description
<pre>host HOST1 { ##### Mac to IP mapping hardware ethernet 00:01:e8:8c:4d:0e; fixed-address 30.0.0.20; ##### Config file name could be given in the following way option configfile "ftp://admin:admin@30.0.0.1/ pt-MXLSwitchIO-12"; option configfile "http://Guest-1/pt-MXLSwitchIO-12";</pre>	<p>FTP URL with IP address</p> <p>HTTP URL with DNS</p>

```
option configfile "pt-MXLSwitchIO-12";
##### bootfile-name could be given in the following way
option bootfile-name "ftp://admin:admin@Guest-1/jumpstart";
option bootfile-name "http://30.0.0.1/jumpstart";
option bootfile-name "tftp://30.0.0.1/jumpstart";
)
```

TFTP

FTP URL with DNS

HTTP URL with IP address

TFTP URL with IP address

DHCP Configuration

Prior to implementing BMP mode, you must update the **dhcp.conf** file on the appropriate DHCP server.

- Set up a DHCP server. For more information, refer to the *FTOS Configuration Guide Dynamic Host Configuration Protocol* chapter. The DHCP server is configured to assign an IP to the system and other parameters.
 - **Boot File Name:** The image to be loaded on the system. The boot file name is expected to use option 67.
 - **Configuration File Name:** The configurations to be applied to the system. The configuration file name is expected to use option 209.
 - **IP Server Address:** The server where the Image and Configuration files are placed. The address is assumed to be a TFTP address unless it is given as a URL. The system supports TFTP, HTTP, FTP, Flash, and USB file names.
 - **Domain Name Server:** The DNS server to be contacted to resolve for the hostname.
 - **IP Address:** Dynamic IP address for the system.
- The DHCP option codes used are:
 - 6 Domain Name Server IP
 - 209 Configuration File



Note: The boot file name and configuration file name must be in the correct format. If it is not, the chassis loaded in jump-start mode is unable to download the file from the DHCP server, and it behaves as if the server could not be reached. The discovery process continues, despite configured time-out, until you enter the stop jump-start command.

A sample **dhcp.conf** file:

Parameter Example	Description
lease-file-name "/var/lib/dhcpd/dhcpd.leases";	
##### configuration file name can be given as	
option configfile code 209 = text;	
option tftp-server-address code 150 = ip-address;	TFTP server IP address
option tftp-server-address code 150 = text;	TFTP server name
##### bootfile name could be given in the following way	
option bootfile-name code 67 = text;	Boot file name
option routers code 3 = ip-address;	

```

option routers code 3 = ip-address;
subnet 30.0.0.0 netmask 255.255.0.0 {
    range 30.0.1.17 30.0.1.100;
    option tftp-server-address 30.0.0.1; (IP address)

    option tftp-server-address "Guest-1" (DNS)
    option domain-name-servers 30.0.0.1;
    option routers 30.0.0.14;
}

```

Boot file location IP address

DNS server hostname

IP Server

- Set up an IP server and ensure connectivity.

The server that holds the boot and configuration files must be configured as the network source for the system. By default, the necessary files are stored in the `/tftpboot` directory. However, the system also recognizes HTTP, FTP, Flash, and USB URLs. For example:

- `tftp://server ip or name/filename`
- `ftp://user:passwd@serverip or name//mypath/FTOS-A.B.C.D.bin`
- `flash://filename`
- `usbflash://FTOS-1.2.3.4.bin`
- `http://server ip`
- `name/filename`



Note: The boot file name and configuration file name must be in the correct format. If it is not, the jump-start reload is unable to download the file from the DHCP server, and it behaves as if the server could not be reached. The discovery process continues, despite configured time-out, until you enter the stop jump-start command.

When loading the FTOS image, if the FTOS image on the server is different from the image on the local flash, the system downloads the image from the server onto the local flash and reloads using that image. If it is the same image, the system reloads from the flash without downloading a new image.

Domain Name Server

- Set up a DNS server. For more information, refer to the *FTOS Configuration Guide IPv4 Addressing* chapter.

You must configure the Domain Name Server to determine the hostname for applying the configuration to the system when the DHCP offer does not have a configuration file specified. The DNS server is contacted only when there is no configuration file specified in the DHCP offer and hostname is not resolved from network-config file.

Boot Commands

Command Syntax	Command Mode	Purpose
reload-type jump-start auto-save dhcp-timeout <i>minutes</i> config-download [enable disable] retry-count	EXEC Privilege	<p>Reload the system in BMP mode. To reload in non-BMP mode, enter reload-type normal command.</p> <p>Enter config-download enable to download the configuration file from the DHCP server. Enter config-download disable so that the system uses the start-up configuration file on the flash.</p> <p>Enter auto-save enable to configure the auto-save option for the downloaded configuration file.</p> <p>The retry-count is to configure the number of entries for configuration download.</p> <p>The dhcp-timeout (1 to 50 minutes) is the amount of time the system waits for a DHCP server response before reverting to Normal mode and loading the start-up configuration files from the flash. The default time is infinity; if no time is set, the system continues to wait unless the stop jump-start command is given.</p>
stop jump-start	EXEC Privilege	<p>This command stops the jump-start reload process while it is in progress. If the command is initiated while the system is downloading an image or configuration file, the command takes effect when the DHCP release is sent.</p>

System Boot and Set-Up Behavior



Note: Configure the system to boot in Normal mode (NOT using BMP) by entering the reload-type normal command.

When the system boots up in BMP mode all ports, including the management port on the master unit are placed in **L3** mode in a **no shut** state. The system acts as a DHCP client on these ports for a period of time (dhcp-timeout). This allows the system time to send out a DHCP DISCOVER on all the **interface up** ports to the DHCP Server in order to obtain its IP address, boot image filename, and configuration file from the DHCP server.

1. System begins boot up process in BMP mode (default mode).

2. The system sends DHCP Discover on all the interface up ports.

```
00:01:31: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_DISCOVER: DHCP DISCOVER sent on Ma 0/0.
00:01:31: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_DISCOVER: DHCP DISCOVER sent on Te 0/0.
00:01:31: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_DISCOVER: DHCP DISCOVER sent on Te 0/5.
00:01:31: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_DISCOVER: DHCP DISCOVER sent on Te 0/6.
00:01:47: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_DISCOVER: DHCP DISCOVER sent on Te 0/8.
00:01:47: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_DISCOVER: DHCP DISCOVER sent on Te 0/35.
00:01:47: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_DISCOVER: DHCP DISCOVER sent on Fo 0/56.
00:01:47: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_DISCOVER: DHCP DISCOVER sent on Fo 0/60.
00:01:47: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_DISCOVER: DHCP DISCOVER sent on Ma 0/0.
```

3. The IP address, boot image filename, and configuration filename are reserved for the system and provided in the DHCP reply (one-file read method). The system receives its IP address, subnet mask, DHCP server IP, TFTP server address, DNS server IP, boot file name, and the configuration filename from the DHCP server.

If a DHCP offer has no image path or configuration file path, it is considered to be an invalid BMP DHCP offer and the offer is ignored. The first DHCP offer with IP address, FTOS image and configuration file, *or* the IP address and FTOS image, *or* the IP address and configuration file is chosen.

4. The DHCP OFFER is selected.

```
***** DHCP OFFER DETAILS *****
DHCP acquired IP           = 30.0.0.20
subnet-mask                = 255.255.0.0
DHCP provided Image file  = jumpstart
DHCP provided Config file = pt-dt-m1000e-3-a2-12
DHCP Server IP            = 30.0.0.1
TFTP IP                    = 30.0.0.1
DNS IP                     = 30.0.0.1
Routers                    = 30.0.0.14
*****
```

The details are displayed in the syslog messages:

```
00:01:21: %STKUNIT0-M:CP %JUM00:01:21: %STKUNIT0-M:CP %JUMPSTART-5-BOOT_OFFER: DHCP acquired IP
10.16.134.57 mask 255.255.0.0 server IP 10.16.134.207.
PSTART-5-BOOT_OFFER: DHCP tftp IP NIL dns IP NIL router IP 30.1.1.1.
00:01:21: %STKUNIT0-M:CP %JUMPSTART-5-BOOT_OFFER: DHCP image file tftp://10.16.134.207/jumpstart.
00:01:21: %STKUNIT0-M:CP %JUMPSTART-5-BOOT_OFFER: DHCP config file tftp://10.16.134.207/
pt-dt-m1000e-3-a2-12.
```

5. The system sends a unicast message to the server to retrieve the named configuration file and/or boot file from the base directory of the server.

- a The FTOS image is expected to be a Boot filename in the DHCP offer (128 bytes). The name can be a fully qualified URL or it can be a file name only.
- b When an FTOS image is found, the system compares that image to the version the chassis currently has loaded.

- If there is a mismatch, the system upgrades to the downloaded version and reloads.

```
*****VALID IMAGE*****
```

```
DOWNLOADED RELEASE HEADER :
Release Image Major Version : 8
Release Image Minor Version : 3
Release Image Main Version : 8
Release Image Patch Version : 33
```

```
FLASH RELEASE HEADER B :
Release Image Major Version : 8
Release Image Minor Version : 3
Release Image Main Version : 8
Release Image Patch Version : 28
```

```
00:04:05: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_DOWNLOAD: The FTOS image download is
successful.
```

```
Erasing MXL 10/40GbE Switch IO Module Primary Image, please wait
```

```
.....
.....
.....
.....
.....00:09:50: %STKUNsyncing disks... IT0-M:CP %CHMGR-1
5-RELOAD: User done
request to reload the chassis
rebooting
```

- If there is no version mismatch, the system downloads the configuration file.

```
00:03:27: %STKUNIT0-M:CP %JUMPSTART-5-CFG_APPLY: The downloaded config from dhcp
server is being applied
```

```
00:03:27: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_RELEASE: DHCP RELEASE sent on Fo 0/56.
```

```
00:03:27: %STKUNIT0-M:CP %SYS-5-CONFIG_LOAD: Loading configuration file
```

- c If the configuration file is downloaded from the server, any saved start-up configuration file on the flash is ignored. If no configuration file is downloaded from the server, the start-up configuration file on the flash is loaded as in normal reload.

6. When the FTOS image and the configuration file has been downloaded, the IP address is released.

7. When the boot image has been downloaded, the DHCP RELEASE is sent

```
00:04:06: %STKUNIT0-M:CP %JUMPSTART-5-JUMPSTART_RELEASE: DHCP RELEASE sent on Fo 0/56.
```

```
00:04:06: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN: Changed interface Admin state to down: Fo 0/56
```

The system is now up and running, and it can be managed as usual.

Content Addressable Memory (CAM)

Content addressable memory (CAM) is a type of memory that stores information in the form of a look-up table (LUT). On Dell Force10 systems, the CAM stores Layer 2 and Layer 3 forwarding information, access-lists (ACL), flows, and routing policies.

This chapter contains the following sections:

- [CAM Allocation](#)
- [Test CAM Usage](#)
- [View CAM-ACL Settings](#)
- [CAM Optimization](#)

CAM Allocation

Allocate space for IPV4 ACLs and quality of service (QoS) regions by using the `cam-acl` command in CONFIGURATION mode.

The CAM space is allotted in field processor (FP) blocks. The total space allocated must equal 13 FP blocks. Note that there are 16 FP blocks, but the system flow requires three blocks that cannot be reallocated. The default CAM allocation settings are:

- L3 ACL (`ipv4acl`): 2
- L2 ACL (`l2acl`): 2
- IPv6 L3 ACL (`ipv6acl`): 0
- L3 QoS (`ipv4qos`): 2
- L2 QoS (`l2qos`): 1
- L2PT (`l2pt`): 0
- MAC ACLs (`ipmacacl`): 0
- ECFMAACL (`ecfmacl`): 0
- FCOEACL (`fcoeacl`): 4
- ISCSIOPTACL (`iscsioptacl`): 2
- VMAN QoS (`vman-qos`): 0
- VMAN Dual QoS (`vman-dual-qos`): 0

The ipv6acl and vman-dual-qos allocations must be entered as a factor of 2 (2, 4, 6, 8, 10). All other profile allocations can use either even or odd numbered ranges.



Note: On the MXL 10/40GbE Switch IO Module, there can be *only one* odd number of blocks in the command line interface (CLI) configuration; the other blocks must be in factors of two. For example, a CLI configuration of 5+4+2+1+1 blocks is not supported; a configuration of 6+4+2+1 blocks is supported.

You must save the new CAM settings to the startup-config (write-mem or copy run start) then reload the system for the new settings to take effect.

To configure the IPv4 ACLs and QoS regions on the entire system, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Select a cam-acl action	cam-acl [default l2acl]	CONFIGURATION
	Note: Selecting default resets the CAM entries to the default settings. Select l2acl to allocate space for the ACLs and QoS regions.		
2	Enter the number of FP blocks for each region.	l2acl <i>number</i> ipv4acl <i>number</i> ipv6acl <i>number</i> , ipv4qos <i>number</i> l2qos <i>number</i> , l2pt <i>number</i> ipmacacl <i>number</i> ecfmacacl <i>number</i> [vman-qos vman-dual-qos <i>number</i>]	EXEC Privilege
3	Reload the system.	reload	EXEC Privilege
4	Verify that the new settings will be written to the CAM on the next boot.	show cam-acl	EXEC Privilege

Test CAM Usage

This command applies to both IPv4 CAM profiles, but is best used when verifying QoS optimization for IPv6 ACLs.

Use this command to determine whether sufficient ACL CAM space is available to enable a service-policy. Create a class map with all required ACL rules, then execute the test cam-usage command in Privilege mode to verify the actual CAM space required. [Figure 7-1](#) gives a sample of the output shown when executing the command. The status column indicates whether or not you can enable the policy.

Figure 7-1. Command Example: test cam-usage

```
FTOS#test cam-usage service-policy input pmap stack-unit all
```

Stack-Unit	Portpipe	CAM Partition	Available CAM	Estimated CAM per Port	Status
2	0	L2ACL	28	1	Allowed (28)

View CAM-ACL Settings

View the current cam-acl settings for the system chassis and each component using the show cam-acl command (Figure 7-2).

Figure 7-2. View CAM-ACL settings

```
FTOS#show cam-acl
```

```
-- Chassis Cam ACL --
Current Settings(in block sizes)
L2Acl      :      6
Ipv4Acl    :      2
Ipv6Acl    :      0
Ipv4Qos    :      2
L2Qos      :      1
L2PT       :      0
IpMacAcl   :      0
VmanQos    :      0
VmanDualQos :      0
EcfmAcl    :      0
FcoeAcl    :      0
iscsiOptAcl :      2
```

```
-- Stack unit 5 --
Current Settings(in block sizes)
L2Acl      :      6
Ipv4Acl    :      2
Ipv6Acl    :      0
Ipv4Qos    :      2
L2Qos      :      1
L2PT       :      0
IpMacAcl   :      0
VmanQos    :      0
VmanDualQos :      0
EcfmAcl    :      0
FcoeAcl    :      0
iscsiOptAcl :      2
```

```
FTOS#
```

CAM Optimization

When you enable the CAM optimization command, if a policy map containing classification rules (ACL and/or dscp/ip-precedence rules) is applied to more than one physical interface on the same port-pipe, only a single copy of the policy is written (only one FP entry is used).

When you disable this command, the system behaves as described in this chapter. However, enabling CAM optimization would apply a single rate policy FP entry. If the input service policy maps applied to several ports are the same, rate policing is applied to all the ports as a group and not individually.

Data Center Bridging (DCB)

The data center bridging (DCB) features are supported on the MXL 10/40GbE Switch.

This chapter describes the following data center bridging topics:

- [Ethernet Enhancements in Data Center Bridging](#)
- [Enabling Data Center Bridging](#)
- [Configuring Priority-Based Flow Control](#)
- [Configuring Enhanced Transmission Selection](#)
- [Applying DCB Policies in a Switch Stack](#)
- [Configuring DCBX Operation](#)
- [Verifying DCB Configuration](#)
- [PFC and ETS Configuration Examples](#)

Ethernet Enhancements in Data Center Bridging

Data center bridging (DCB) refers to a set of IEEE Ethernet enhancements that provide data centers with a single, robust, converged network to support multiple traffic types, including local area network (LAN), server, and storage traffic. Through network consolidation, DCB results in reduced operational cost, simplified management, and easy scalability by avoiding the need to deploy separate application-specific networks.

For example, instead of deploying an Ethernet network for LAN traffic, additional storage area networks (SANs) to ensure lossless fibre-channel traffic, and a separate InfiniBand network for high-performance inter-processor computing within server clusters, only one DCB-enabled network is required in a data center. The Dell Force10 switches that support a unified fabric and consolidate multiple network infrastructures use a single input/output (I/O) device called a converged network adapter (CNA).

A CNA is a computer input/output device that combines the functionality of a host bus adapter (HBA) with a network interface controller (NIC). Multiple adapters on different devices for several traffic types are no longer required.

Data center bridging satisfies the needs of the following types of data center traffic in a unified fabric:

- LAN traffic consists of a large number of flows that are generally insensitive to latency requirements, while certain applications, such as streaming video, are more sensitive to latency. Ethernet functions as a best-effort network that may drop packets in case of network congestion. IP networks rely on transport protocols (for example, TCP) for reliable data transmission with the associated cost of greater processing overhead and performance impact.
- Storage traffic based on Fibre Channel media uses the SCSI protocol for data transfer. This traffic typically consists of large data packets with a payload of 2K bytes that cannot recover from frame loss. To successfully transport storage traffic, data center Ethernet must provide no-drop service with lossless links.
- Servers use InterProcess Communication (IPC) traffic within high-performance computing clusters to share information. Server traffic is extremely sensitive to latency requirements.

To ensure lossless delivery and latency-sensitive scheduling of storage and service traffic and I/O convergence of LAN, storage, and server traffic over a unified fabric, IEEE data center bridging adds the following extensions to a classical Ethernet network:

- 802.1Qbb - Priority-based Flow Control (PFC)
- 802.1Qaz - Enhanced Transmission Selection (ETS)
- 802.1Qau - Congestion Notification
- Data Center Bridging Exchange (DCBX) protocol



Note: In FTOS version 8.3.16.1, only the PFC, ETS, and DCBX features are supported in data center bridging.


Priority-Based Flow Control

In a data center network, priority-based flow control (PFC) manages large bursts of one traffic type in multiprotocol links so that it does not affect other traffic types and no frames are lost due to congestion. When PFC detects congestion on a queue for a specified priority, it sends a pause frame for the 802.1p priority traffic to the transmitting device. In this way, PFC ensures that large amounts of queued LAN traffic do not cause storage traffic to be dropped, and that storage traffic does not result in high latency for high-performance computing (HPC) traffic between servers.

PFC enhances the existing 802.3x pause and 802.1p priority capabilities to enable flow control based on 802.1p priorities (classes of service). Instead of stopping all traffic on a link (as performed by the traditional Ethernet pause mechanism), PFC pauses traffic on a link according to the 802.1p priority set on a traffic type. You can create lossless flows for storage and server traffic while allowing for loss in case of LAN traffic congestion on the same physical interface.

Figure 8-1 shows how PFC handles traffic congestion by pausing the transmission of incoming traffic with dot1p priority 3.

Figure 8-1. Priority-Based Flow Control



PFC is implemented as follows in the Dell Force10 operating software (FTOS):

- PFC is supported on specified 802.1p priority traffic (dot1p 0 to 7) and is configured per interface. However, only two lossless queues are supported on an interface: one for FCoE converged traffic and one for SCSI storage traffic. You must configure the same lossless queues on all ports.
- PFC delay constraints place an upper limit on the transmit time of a queue after receiving a message to pause a specified priority.
- By default, PFC is enabled on an interface with no dot1p priorities configured. You can configure the PFC priorities if the switch negotiates with a remote peer using DCBX.
- During DCBX negotiation with a remote peer:
 - If the negotiation succeeds and the port is in DCBX Willing mode to receive a peer configuration, PFC parameters from the peer are used to configure PFC priorities on the port. If you enable the link-level flow control mechanism on the interface, DCBX negotiation with a peer is not performed.
 - If the negotiation fails and PFC is enabled on the port, any user-configured PFC input policies are applied. If no PFC input policy has been previously applied, the PFC default setting is used (no priorities configured). If you do not enable PFC on an interface, you can enable the 802.3x link-level pause function. By default, the link-level pause is disabled.
- PFC supports buffering to receive data that continues to arrive on an interface while the remote system reacts to the PFC operation.
- PFC uses the DCB MIB IEEE802.1azd2.5 and the PFC MIB IEEE802.1bb-d2.2.


Enhanced Transmission Selection

Enhanced transmission selection (ETS) supports optimized bandwidth allocation between traffic types in multiprotocol (Ethernet, FCoE, SCSI) links. ETS allows you to divide traffic according to its 802.1p priority into different priority groups (traffic classes) and configure bandwidth allocation and queue scheduling for each group to ensure that each traffic type is correctly scheduled and receives its required bandwidth. For example, you can prioritize low-latency storage or server cluster traffic in a traffic class to receive more bandwidth and restrict best-effort LAN traffic assigned to a different traffic class.

Although you can configure strict-priority queue scheduling for a priority group, ETS introduces flexibility that allows the bandwidth allocated to each priority group to be dynamically managed according to the amount of LAN, storage, and server traffic in a flow. Unused bandwidth in a priority-group is dynamically allocated to other priority groups for which traffic is available to be scheduled. Traffic is queued according to its 802.1p priority assignment, while flexible bandwidth allocation and the configured queue-scheduling for a priority group is supported.

Figure 8-2 shows how ETS allows you to allocate bandwidth when different traffic types are classed according to 802.1p priority and mapped to priority groups.

Figure 8-2. Enhanced Transmission Selection



ETS uses the following traffic groupings to select multiprotocol traffic for transmission:

- Priority group: A group of 802.1p priorities used for bandwidth allocation and queue scheduling. All 802.1p priority traffic in a group should have the same traffic handling requirements for latency and frame loss.
- Group ID: A 4-bit identifier assigned to each priority group. Valid values are from 0 to 7.
- Group bandwidth: Percentage of available bandwidth allocated to a priority group.
- Group transmission selection algorithm (TSA): Type of queue scheduling used by a priority group.

ETS is implemented as follows in FTOS:

- ETS supports groups of 802.1p priorities that have:
 - PFC enabled or disabled
 - No bandwidth limit or no ETS processing
- Bandwidth allocated by the ETS algorithm is made available after strict-priority groups are serviced. If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups.
- For ETS traffic selection, an algorithm is applied to priority groups using:
 - Strict-priority shaping
 - ETS shaping

Credit-based shaping is not supported.
- ETS uses the DCB MIB IEEE802.1azd2.5.

Data Center Bridging Exchange Protocol (DCBX)

The data center bridging exchange (DCBX) protocol is enabled by default on any switch on which PFC or ETS are enabled. DCBX allows a switch to automatically discover DCB-enabled peers and exchange configuration information. PFC and ETS use DCBX to exchange and negotiate parameters with peer devices. DCBX capabilities include:

- Discovery of DCB capabilities on peer-device connections
- Determination of possible mismatch in DCB configuration on a peer link
- Configuration of a peer device over a DCB link

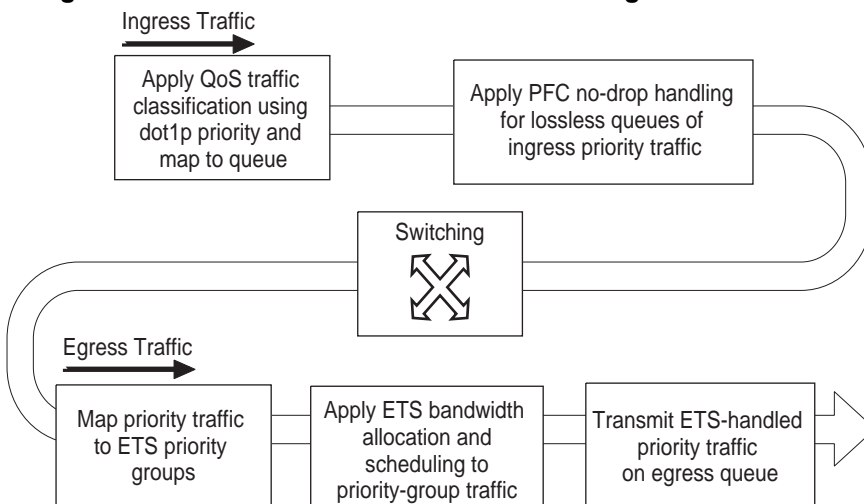
DCBX requires the link layer discovery protocol (LLDP) to provide the path to exchange DCB parameters with peer devices. Exchanged parameters are sent in organizationally specific type, length, values (TLVs) in LLDP data units. For more information, refer to the [Link Layer Discovery Protocol \(LLDP\)](#) chapter. The following LLDP TLVs are supported for DCB parameter exchange:

- PFC parameters: PFC Configuration TLV and Application Priority Configuration TLV.
- ETS parameters: ETS Configuration TLV and ETS Recommendation TLV.

Data Center Bridging in a Traffic Flow

Figure 8-3 shows how DCB handles a traffic flow on an interface.

Figure 8-3. DCB PFC and ETS Traffic Handling



Enabling Data Center Bridging

Data center bridging is enabled by default on an MXL 10/40GbE Switch to support converged enhanced Ethernet (CEE) in a data center network, and is a prerequisite for configuring:

- Priority-based flow control
- Enhanced transmission selection
- Data center bridging exchange protocol
- FCoE initialization protocol (FIP) snooping

DCB processes virtual local area network (VLAN)-tagged packets and dot1p priority values. Untagged packets are treated with a dot1p priority of 0.

For DCB to operate effectively, you can classify ingress traffic according to its dot1p priority so that it maps to different data queues. The dot1p-queue assignments used on an MXL Switch is shown in [Table 8-1 in QoS dot1p Traffic Classification and Queue Assignment](#).

On the MXL Switch, by default, DCB is enabled and MMU buffers are reserved to achieve no-drop traffic handling for PFC. Disabling DCB does not release the buffers reserved by default. To utilize reserved buffers for non-DCB applications, you have to explicitly release the buffers (Refer to [Configuring the PFC Buffer in a Switch Stack](#)).

To disable or re-enable DCB on a switch, enter the following commands:

Task	Command	Command Mode
Disable DCB.	no dcb enable	CONFIGURATION
Re-enable DCB.	dcb enable	CONFIGURATION



FTOS Behavior:

DCB is not supported if you enable link-level flow control on one or more interfaces (refer to [Layer 2 Flow Control Using Ethernet Pause Frames on page 247](#)).

After you disable DCB, if link-level flow control is not automatically enabled on an interface, manually shut down the interface (**shutdown** command) and re-enable it (**no shutdown** command) to enable flow control.

QoS dot1p Traffic Classification and Queue Assignment

DCB supports PFC, ETS, and DCBX to handle converged Ethernet traffic that is assigned to an egress queue according to the following quality of service (QoS) methods:

- Honor dot1p: Using the `service-class dynamic dot1p` command in INTERFACE Configuration mode, you can honor dot1p priorities in ingress traffic at the port or global switch level (refer to [Honoring dot1p Values on Ingress Packets](#)).
- Layer 2 class maps: You can use dot1p priorities to classify traffic in a class map and apply a service policy to an ingress port to map traffic to egress queues (refer to [Policy-Based QoS Configurations](#)).



Note: Dell Force10 does not recommend mapping all ingress traffic to a single queue when using PFC and ETS. Ingress traffic classification using the `service-class dynamic dot1p` command (honor dot1p) is recommended on all DCB-enabled interfaces. If you use L2 class maps to map dot1p priority traffic to egress queues, take into account the default dot1p-queue assignments in [Table 8-1](#) and the maximum number of two lossless queues supported on a port (refer to [Configuring Lossless Queues](#)).

Although FTOS allows you to change the default dot1p priority-queue assignments (refer to [service-class dot1p-mapping on page 462](#)), DCB policies applied to an interface may become invalid if dot1p-queue mapping is reconfigured. If the configured DCB policy remains valid, the change in the dot1p-queue assignment is allowed. For DCB ETS enabled interfaces, traffic destined to queue that is not mapped to any dot1p priority will be dropped.

Table 8-1. dot1p Priority-Queue Assignment

dot1p Value in Incoming Frame	Egress Queue Assignment
0	0
1	0
2	0
3	1
4	2
5	3
6	3
7	3



Note: If you reconfigure the global dot1p-queue mapping, an automatic re-election of the DCBX configuration source port is performed (see [Configuration Source Election](#)).

Configuring Priority-Based Flow Control

Priority-based flow control provides a flow control mechanism based on the 802.1p priorities in converged Ethernet traffic received on an interface and is enabled by default. As an enhancement to the existing Ethernet pause mechanism, PFC stops traffic transmission for specified priorities (CoS values) without impacting other priority classes. Different traffic types are assigned to different priority classes.

When traffic congestion occurs, PFC sends a pause frame to a peer device with the CoS priority values of the traffic that needs to be stopped. DCBX provides the link-level exchange of PFC parameters between peer devices. PFC allows network administrators to create zero-loss links for SAN traffic that requires no-drop service, while at the same time retaining packet-drop congestion management for LAN traffic.

To ensure complete no-drop service, you must apply the same DCB input policy with the same pause time and dot1p priorities on all PFC-enabled peer interfaces.

To configure PFC and apply a PFC input policy to an interface, follow these steps:

Step	Task	Command	Command Mode
1	Create a DCB input policy to apply pause or flow control for specified priorities using a configured delay time. Maximum: 32 alphanumeric characters.	<code>dcb-input <i>policy-name</i></code>	CONFIGURATION
2	Configure the link delay used to pause specified priority traffic. One quantum is equal to a 512-bit transmission. Valid values (in quanta): 712-65535. Default: 45556 quantum in link delay.	<code>pfc link-delay <i>value</i></code>	DCB INPUT POLICY
3	Configure the CoS traffic to be stopped for the specified delay. Enter the 802.1p values of the frames to be paused. Valid values: 0-7. Default: None. Maximum number of loss less queues supported on the switch: 2. Separate priority values with a comma; specify a priority range with a dash; for example: <code>pfc priority 1,3,5-7</code> .	<code>pfc priority <i>priority-range</i></code>	DCB INPUT POLICY
4	Enable the PFC configuration on the port so that the priorities are included in DCBX negotiation with peer PFC devices. Default: PFC mode is on.	<code>pfc mode on</code>	DCB INPUT POLICY
5	(Optional) Enter a text description of the input policy. Maximum: 32 characters.	<code>description <i>text</i></code>	DCB INPUT POLICY
6	Exit DCB input policy configuration mode.	<code>exit</code>	DCB INPUT POLICY
7	Enter interface configuration mode.	<code>interface type <i>slot/port</i></code>	CONFIGURATION
8	Apply the input policy with the PFC configuration to an ingress interface.	<code>dcb-policy input <i>policy-name</i></code>	INTERFACE
9	Repeat Steps 1 to 8 on all PFC-enabled peer interfaces to ensure lossless traffic service.		



FTOS Behavior:

As soon as you apply a DCB policy with PFC enabled on an interface, DCBX starts exchanging information with PFC-enabled peers. The IEEE802.1Qbb, CEE and CIN versions of PFC TLV are supported. DCBX also validates PFC configurations received in TLVs from peer devices.

By applying a DCB input policy with PFC enabled, you enable PFC operation on ingress port traffic. To achieve complete lossless handling of traffic, you must also enable PFC on all DCB egress ports or configure the dot1p priority-queue assignment of PFC priorities to lossless queues (refer to [Configuring Lossless Queues](#)).

To remove a DCB input policy, including the PFC configuration it contains, use the **no dcb-input policy-name** command in INTERFACE Configuration mode. To disable PFC operation on an interface, use the **no pfc mode on** command in DCB Input Policy Configuration mode. PFC is enabled or disabled as global DCB operation is enabled (**dcb enable**) or disabled (**no dcb enable**).

You can enable any number of 802.1p priorities for PFC. Queues to which PFC priority traffic is mapped are lossless by default. Traffic may be interrupted due to an interface flap (going down and coming up) when you reconfigure the lossless queues for no-drop priorities in a PFC input policy and re-apply the policy to an interface.

For PFC to be applied, the configured priority traffic must be supported by a PFC peer (as detected by DCBX).

To honor a PFC pause frame multiplied by the number of PFC-enabled ingress ports, the minimum link delay should be greater than the round trip transmission time needs to honor PFC pause frame by peer multiplied by the number of PFC-enabled ingress ports.

If you apply an input policy with PFC disabled (**no pfc mode on**):

- Link-level flow control can be enabled on the interface (refer to [Layer 2 Flow Control Using Ethernet Pause Frames on page 247](#)). To delete the input policy, you must first disable link-level flow control. PFC is then automatically enabled on the interface because an interface is by default PFC-enabled.
- PFC still allows you to configure lossless queues on a port to ensure no-drop handling of lossless traffic (refer to [Configuring Lossless Queues](#)).

PFC and link-level flow control cannot be enabled at the same time on an interface.

When you apply an input policy to an interface, an error message is displayed if:

- The PFC dot1p priorities result in more than two lossless port queues globally on the switch.
- Link-level flow control is already enabled. PFC and link-level flow control cannot be enabled at the same time on an interface.

In a switch stack, you must configure all stacked ports with the same PFC configuration.

A DCB input policy for PFC applied to an interface may become invalid if dot1p-queue mapping is reconfigured (refer to [Create Input Policy Maps](#)). This situation occurs when the new dot1p-queue assignment exceeds the maximum number (2) of lossless queues supported globally on the switch. In this case, all PFC configurations received from PFC-enabled peers are removed and re-synchronized with the peer devices.

Traffic may be interrupted when you reconfigure PFC no-drop priorities in an input policy or re-apply the policy to an interface.

FTOS does not support MACsec Bypass Capability (MBC).

Configuring Lossless Queues

DCB also supports the manual configuration of lossless queues on an interface when PFC mode is turned off and priority classes are disabled in a DCB input policy applied to the interface. The configuration of no-drop queues provides flexibility for ports on which PFC is not needed but lossless traffic should egress from the interface.

Lossless traffic egresses out the no-drop queues. Ingress dot1p traffic from PFC-enabled interfaces is automatically mapped to the no-drop egress queues.

Prerequisite: A DCB input policy with PFC configuration is applied to the interface with the following conditions:

- PFC mode is off (no pfc mode on).
- No PFC priority classes are configured (no pfc priority *priority-range*).

To configure lossless queues on a port interface, follow these steps:

Step	Task	Command	Command Mode
1	Enter INTERFACE Configuration mode.	<code>interface type slot/port</code>	CONFIGURATION
2	Configure the port queues that will still function as no-drop queues for lossless traffic. For the dot1p-queue assignments, refer to Table 8-1 . Maximum number of lossless queues globally supported on the switch: 2. Valid values: 0-3. Separate queue values with a comma; specify a priority range with a dash; for example: <code>pfc no-drop queues 1,3</code> or <code>pfc no-drop queues 2-3</code> Default: No lossless queues are configured.	<code>pfc no-drop queues queue-range</code>	INTERFACE



FTOS Behavior:

By default, no lossless queues are configured on a port.

A limit of two lossless queues are supported on a port. If the amount of priority traffic that you configure to be paused exceeds the two lossless queues, an error message is displayed. You must reconfigure the input policy using a smaller number of PFC priorities.

If you configure lossless queues on an interface that already has a DCB input policy with PFC enabled (**pfc mode on**), an error message is displayed.

Traffic may be interrupted due to an interface flap (going down and coming up) when you reconfigure lossless queues on no-drop priorities in an input policy and re-apply the policy to an interface.

Configuring the PFC Buffer in a Switch Stack

In a switch stack, you must configure all stacked ports with the same PFC configuration. In addition, you must configure a separate buffer of memory allocated exclusively to a service pool accessed by queues on which priority-based control flows are mapped.

These PFC-enabled queues ensure the lossless transmission of storage and server traffic. The buffer required for the PFC service pool is calculated based on the number of ports and port queues used by PFC traffic.

You can configure the size of the PFC buffer for all switches in a stack or all port pipes on a specified stack unit by entering the following commands on the master switch:

Task	Command	Command Mode
Configure the PFC buffer for all switches in the stack. Default: The PFC buffer is enabled on all ports on the stack unit.	[no] dcb stack-unit all pfc-buffering pfc-port-count { 1-56 } pfc-queues { 1-2 }	CONFIGURATION
Configure the PFC buffer for all port pipes in a specified stack unit by specifying the port-pipe number, number of PFC-enabled ports, and number of configured lossless queues. Valid stack-unit IDs are 0 to 5. The only valid port-set ID (port-pipe number) on an MXL Switch is 0.	[no] dcb stack-unit <i>stack-unit-id</i> [port-set <i>port-set-id</i>] pfc-buffering pfc-ports { 1-56 } pfc-queues { 1-2 }	CONFIGURATION



FTOS Behavior:

If you configure PFC on a 40GbE port, count the 40GbE port as four PFC-enabled ports in the **pfc-port** number you enter in the command syntax.

To achieve lossless PFC operation, the PFC port count and queue number used for the reserved buffer size that is created must be greater than or equal to the buffer size required for PFC-enabled ports and lossless queues on the switch.

For the PFC buffer configuration to take effect, you must reload the stack or a specified stack unit (**reload** command at the EXEC Privilege level).

If you configure the PFC buffer on all stack units, delete the startup configuration on both the master and standby, and reload the stack, the new master (previously standby) generates the following syslog message for each stack unit when it boots up:

PFC_BUFFER_CONFIG_CHANGED is generated for all stack units.

Configuring Enhanced Transmission Selection

Enhanced transmission selection (ETS) provides a way to optimize bandwidth allocation to outbound 802.1p classes of converged Ethernet traffic. Different traffic types have different service needs. Using ETS, you can create groups within an 802.1p priority class to configure different treatment for traffic with different bandwidth, latency, and best-effort needs.

For example, storage traffic is sensitive to frame loss; interprocess communication (IPC) traffic is latency-sensitive. ETS allows different traffic types to coexist without interruption in the same converged link by:

- Allocating a guaranteed share of bandwidth to each priority group
- Allowing each group to exceed its minimum guaranteed bandwidth if another group is not fully using its allotted bandwidth

To configure ETS and apply an ETS output policy to an interface, you must:

1. Create a QoS output policy with ETS scheduling and bandwidth allocation settings.
2. Create a priority group of 802.1p traffic classes.
3. Configure a DCB output policy in which you associate a priority group with a QoS ETS output policy.
4. Apply the DCB output policy to an interface.

ETS Prerequisites and Restrictions

The following prerequisites and restrictions apply when you configure ETS bandwidth allocation or queue scheduling and apply a QoS ETS output policy on an interface:

- Configuring ETS bandwidth allocation or a queue scheduler for dot1p priorities in a priority group is applicable only if the DCBX version used on a port is CIN (refer to [Configuring DCBX on an Interface](#)).
- When allocating bandwidth or configuring a queue scheduler for dot1p priorities in a priority group on a DCBX CIN interface, take into account the CIN bandwidth allocation ([Configuring Bandwidth Allocation for DCBX CIN](#)) and dot1p-queue mapping ([Table 8-1](#)).
- Although an ETS output policy does not support WRED, ECN, rate shaping, and rate limiting because these parameters are not negotiated by DCBX with peer devices, you can apply a QoS output policy with WRED and/or rate shaping on a DCBX CIN-enabled interface (refer to [Configure Port-based Rate Shaping on page 420](#) and [Weighted Random Early Detection](#)). In this case, the WRED or rate shaping configuration in the QoS output policy should take into account the bandwidth allocation or queue scheduler configured in the ETS output policy.
- You can only use a QoS ETS output policy in association with a priority group in a DCB output policy and cannot be applied to an interface as a normal QoS output policy (refer to [Applying an ETS Output Policy for a Priority Group to an Interface](#) and [Create an Output QoS Policy](#)).



Note: The IEEE 802.1Qaz, CEE, and CIN versions of ETS are supported.

Creating a QoS ETS Output Policy

A QoS output policy that you create to optimize bandwidth on an output interface for specified priority traffic consists of the ETS settings used in DCBX negotiations with peer devices:

- Bandwidth percentage
- Queue scheduling

To create a QoS output policy with ETS settings, follow these steps:

Step	Task	Command	Command Mode
1	Create a QoS output policy to configure the ETS bandwidth allocation and scheduling for priority traffic. Maximum: 32 characters.	<code>qos-policy-output <i>policy-name ets</i></code>	CONFIGURATION
2	(Optional) Configure the method used to schedule priority traffic in port queues. Valid values: <ul style="list-style-type: none">• <code>strict</code> - Strict priority traffic is serviced before any other queued traffic (refer to Strict-Priority Queueing).• <code>werr</code> - Weighted elastic round robin provides low-latency scheduling for priority traffic on port queues. Default: WERR scheduling is used to queue priority traffic. Note: If you configure a scheduling method, you cannot configure bandwidth allocation in Step 3.	<code>scheduler <i>value</i></code>	POLICY-MAP-OUT-ETS
3	(Optional) Configure the bandwidth percentage allocated to priority traffic in port queues. Percentage range: 1 to 100% in units of 1%. The sum of bandwidth percentage assigned to dot1p priorities/queues in a priority group should be 100%. Default: None. Note: If you configure bandwidth allocation, you cannot configure a scheduling method in Step 2.	<code>bandwidth-percentage <i>percentage</i></code>	POLICY-MAP-OUT-ETS
4	Exit ETS Output Policy Configuration mode.	<code>exit</code>	POLICY-MAP-OUT-ETS

**FTOS Behavior:**

Traffic in priority groups is assigned to strict-queue or WERR scheduling in an ETS output policy and is managed using the ETS bandwidth-assignment algorithm. FTOS dequeues all frames of strict-priority traffic before servicing any other queues. A queue with strict-priority traffic can starve other queues in the same port.

ETS-assigned bandwidth allocation and scheduling apply only to data queues, not to control queues.

FTOS supports hierarchical scheduling on an interface. FTOS control traffic is redirected to control queues as higher priority traffic with strict priority scheduling. After control queues drain out, the remaining data traffic is scheduled to queues according to the bandwidth and scheduler configuration in the ETS output policy. The available bandwidth calculated by the ETS algorithm is equal to the link bandwidth after scheduling non-ETS higher-priority traffic.

The configuration of bandwidth allocation and strict-queue scheduling is not supported at the same time for a priority group. If both are configured, the configured bandwidth allocation is ignored for priority-group traffic when you apply the output policy on an interface (refer to [Applying an ETS Output Policy for a Priority Group to an Interface](#)).

Bandwidth assignment in a dot.1p priority-queue: By default, equal bandwidth is assigned to each port queue and each dot1p priority in a priority group. Use the **bandwidth-percentage** command to configure bandwidth amounts in associated dot1p queues. When specified bandwidth is assigned to some port queues and not to others, the remaining bandwidth (100% minus assigned bandwidth amount) is equally distributed to unassigned non-strict priority queues in the priority group. The sum of the allocated bandwidth to all queues in a priority group should be 100% of the bandwidth on the link.

Bandwidth assignment in a priority group: By default, equal bandwidth is assigned to each priority group in the ETS output policy applied to an egress port if you did not configure bandwidth allocation. The sum of configured bandwidth allocation to dot1p priority traffic in all ETS priority groups must be 100%. You must allocate at least 1% of the total bandwidth to each priority group and queue. If you assign bandwidth to some priority groups but not to others, the remaining bandwidth (100% minus assigned bandwidth amount) is equally distributed to non-strict-priority groups which have no configured scheduler.

Scheduling of priority traffic: dot1p priority traffic on the switch is scheduled according to the current dot1p-queue mapping. dot1p priorities within the same queue should have the same traffic properties and scheduling method.

ETS output-policy error: If an error occurs in an ETS output-policy configuration, the configuration is ignored and the scheduler and bandwidth allocation settings are reset to the ETS default values (all priorities are in the same ETS priority group and bandwidth is allocated equally to each priority). If an error occurs when a port receives a peer's ETS configuration, the port's configuration is reset to the previously configured ETS output policy. If no ETS output policy was previously applied, the port is reset to the default ETS parameters.

Creating an ETS Priority Group

An ETS priority group specifies the range of 802.1p priority traffic to which a QoS output policy with ETS settings is applied on an egress interface. You can associate a priority group to more than one ETS output policy on different interfaces.

To create a priority group for ETS, follow these steps:

Step	Task	Command	Command Mode
1	Create an ETS priority group to use with an ETS output policy. Maximum: 32 characters.	<code>priority-group <i>group-name</i></code>	CONFIGURATION
2	Configure the priority-group identifier. Valid values: 0 to 7. Default: None.	<code>set-pgid <i>value</i></code>	PRIORITY-GROUP
3	Configure the 802.1p priorities for the traffic on which you want to apply an ETS output policy. Valid values: 0 to 7. Default: None. Separate priority values with a comma; specify a priority range with a dash; for example: <code>priority-list 3,5-7.</code>	<code>priority-list <i>value</i></code>	PRIORITY-GROUP
4	Exit priority-group configuration mode.	<code>exit</code>	PRIORITY-GROUP
5	Repeat Steps 1 to 4 to configure all remaining dot1p priorities in an ETS priority group.		



FTOS Behavior:

A priority group consists of 802.1p priority values that are grouped together for similar bandwidth allocation and scheduling, and that share the same latency and loss requirements. All 802.1p priorities mapped to the same queue should be in the same priority group.

All 802.1p priorities should be configured in priority groups associated with an ETS output policy (refer to [Applying an ETS Output Policy for a Priority Group to an Interface](#)). You can assign each dot1p priority to only one priority group.

By default:

- All 802.1p priorities are grouped in priority group 0.
- 100% of the port bandwidth is assigned to priority group 0. The complete bandwidth is equally assigned to each priority class so that each class has 12 to 13%.

The maximum number of priority groups supported in ETS output policies on an interface is equal to the number of data queues (4) on the port. The 802.1p priorities in a priority group can map to multiple queues.

If you configure more than one priority queue as strict priority or more than one priority group as strict priority, the higher numbered priority queue is given preference when scheduling data traffic.

Applying an ETS Output Policy for a Priority Group to an Interface

To apply ETS on egress port traffic, you must associate a priority group with an ETS output policy which has scheduling and bandwidth configuration in a DCB output policy, and then apply the output policy to an interface. To apply ETS on egress port traffic, follow these steps:

Step	Task	Command	Command Mode
1	Create a DCB output policy to associate an ETS configuration with priority traffic. Maximum: 32 alphanumeric characters.	<code>dcb-output <i>policy-name</i></code>	CONFIGURATION
2	Enable the ETS configuration so that scheduling and bandwidth allocation configured in an ETS output policy or received in a DCBX TLV from a peer can take effect on an interface. Default: ETS mode is on.	<code>ets mode on</code>	DCB OUTPUT POLICY
3	Associate the 802.1p priority traffic in a priority group with the ETS configuration in a QoS output policy.	<code>priority-group <i>group-name</i> qos-policy <i>ets-policy-name</i></code>	DCB OUTPUT POLICY
4	(Optional) Enter a text description of the output policy. Maximum: 32 characters.	<code>description <i>text</i></code>	DCB OUTPUT POLICY
5	Repeat Steps 1 to 4 to configure all remaining ETS priority groups with an ETS output policy.		
6	Exit DCB Output Policy Configuration mode.	<code>exit</code>	DCB OUTPUT POLICY
7	Enter INTERFACE Configuration mode.	<code>interface type <i>slot/port</i></code>	CONFIGURATION
8	Apply the output policy with the ETS configuration to an egress interface.	<code>dcb-policy output <i>policy-name</i></code>	INTERFACE



FTOS Behavior:

Create a DCB output policy to associate a priority group with an ETS output policy with scheduling and bandwidth configuration. You can apply a DCB output policy on multiple egress ports.

The ETS configuration associated with 802.1p priority traffic in a DCB output policy is used in DCBX negotiation with ETS peers.

When you apply an ETS output policy to an interface, ETS-configured scheduling and bandwidth allocation take precedence over any configured settings in the QoS output policies.

To remove an ETS output policy from an interface, use the **no dcb-policy output *policy-name*** command. ETS is enabled by default with the default ETS configuration applied (all dot1p priorities in the same group with equal bandwidth allocation).

If you disable ETS in an output policy applied to an interface (the **no ets mode on** command), any previously configured QoS settings at the interface or global level take effect. If QoS settings are configured at the interface or global level and in an output policy map (the **service-policy output** command), the QoS configuration in the output policy take precedence.

When you apply a DCB output policy with ETS bandwidth allocation to an egress interface which uses default ETS settings, the configured bandwidth allocation may not be applied to dot1p priority traffic in the specified priority group.

ETS Operation with DCBX

In DCBX negotiation with peer ETS devices, ETS configuration is handled as follows:

- ETS TLVs are supported in DCBX versions CIN, CEE, and IEEE2.5.
- ETS operational parameters are determined by the DCBX port-role configurations ([Configuring DCBX Operation](#)).
- ETS configurations received from TLVs from a peer are validated.
- In case of a hardware limitation or TLV error:
 - DCBX operation on an ETS port goes down.
 - New ETS configurations are ignored and existing ETS configurations are reset to the previously configured ETS output policy on the port or to the default ETS settings if no ETS output policy was previously applied.
- ETS operates with legacy DCBX versions as follows:
 - In the CEE version, the priority group/traffic class group (TCG) ID 15 represents a non-ETS priority group. Any priority group configured with a scheduler type is treated as a strict-priority group and is given the priority-group (TCG) ID 15.
 - The CIN version supports two types of strict-priority scheduling:
 - Group strict priority: Allows a single priority flow in a priority group to increase its bandwidth usage to the bandwidth total of the priority group. A single flow in a group can use all the bandwidth allocated to the group.
 - Link strict priority: Allows a flow in any priority group to increase to the maximum link bandwidth.

CIN supports only the dot1p priority-queue assignment in a priority group. To configure a dot1p priority flow in a priority group to operate with link strict priority, you must configure:

- The dot1p priority for strict-priority scheduling (**strict-priority** command; [Strict-Priority Queuing](#))
- The priority group for strict-priority scheduling (**scheduler strict** command; [Creating a QoS ETS Output Policy](#))

If you configure only the priority group in an ETS output policy or only the dot1p priority for strict-priority scheduling, the flow is handled with group strict priority.

Configuring Bandwidth Allocation for DCBX CIN

After you apply an ETS output policy to an interface, if the DCBX version used in your data center network is CIN, you may need to configure a QoS output policy to overwrite the default CIN bandwidth allocation. This default setting divides the bandwidth allocated to each port queue equally between the dot1p priority traffic assigned to the queue.

For more information, refer to [Allocate Bandwidth to the Queue](#).

To create a QoS output policy that allocates different amounts of bandwidth to the different traffic types/dot1p priorities assigned to a queue and apply the output policy to the interface, follow these steps.

Step	Task	Command	Command Mode
1	Create a QoS output policy. Maximum: 32 alphanumeric characters.	qos-policy-output <i>output-policy-name</i>	CONFIGURATION
2	Configure the percentage of bandwidth to be allocated to the dot1p priority/queue traffic in the associated L2 class map. Default: None.	bandwidth-percentage <i>percentage</i>	QoS OUTPUT POLICY
3	Repeat Step 2 to configure bandwidth percentages for other priority queues on the port.	bandwidth-percentage <i>percentage</i>	QoS OUTPUT POLICY
4	Create a priority group for strict-priority scheduling.	scheduler strict	QoS OUTPUT POLICY
5	Exit QoS Output Policy Configuration mode.	exit	QoS OUTPUT POLICY
6	Enter INTERFACE Configuration mode.	interface type <i>slot/port</i>	CONFIGURATION
7	Apply the QoS output policy with the bandwidth percentage for specified priority queues to an egress interface.	service-policy output <i>output-policy-name</i>	INTERFACE

Applying DCB Policies in a Switch Stack

You can apply a DCB input policy with PFC configuration to all stacked ports in a switch stack or on a stacked switch. You can apply different DCB input policies to different stacked switches.

Task	Command	Command Mode
Apply the specified DCB input policy on all ports of the switch stack or a single stacked switch.	<code>dcb-policy input stack-unit {all <i>stack-unit-id</i>} stack-ports all <i>dcb-input-policy-name</i></code>	CONFIGURATION



FTOS Behavior:

A **dcb-policy input stack-unit all** command overwrites any previous **dcb-policy input stack-unit *stack-unit-id*** configurations. Similarly, a **dcb-policy input stack-unit *stack-unit-id*** command overwrites any previous **dcb-policy input stack-unit all** configuration.

Entering the **no dcb-policy input stack-unit all** command removes all DCB input policies applied to stacked ports and resets PFC to its default settings. The **no dcb-policy input stack-unit *stack-unit-id*** command removes only the DCB input policy applied to the specified switch.

You can apply a DCB output policy with ETS configuration to all stacked ports in a switch stack or an individual stacked switch. In addition, you can apply different DCB output policies to different stack units.

Task	Command	Command Mode
Apply the specified DCB output policy on all ports of the switch stack or a stacked switch.	<code>dcb-policy output stack-unit {all <i>stack-unit-id</i>} stack-ports all <i>dcb-output-policy-name</i></code>	CONFIGURATION



FTOS Behavior:

A **dcb-policy output stack-unit all** command overwrites any previous **dcb-policy output stack-unit *stack-unit-id*** configurations. Similarly, a **dcb-policy output stack-unit *stack-unit-id*** command overwrites any previous **dcb-policy output stack-unit all** configuration.

Entering the **no dcb-policy output stack-unit all** command removes all DCB output policies applied to stacked ports. The **no dcb-policy output stack-unit *stack-unit-id*** command removes only the DCB output policy applied to the specified switch.

Configuring DCBX Operation

The data center bridging exchange protocol (DCBX) is used by DCB devices to exchange configuration information with directly connected peers using the link layer discovery protocol (LLDP) protocol. DCBX can detect the misconfiguration of a peer DCB device, and optionally, configure peer DCB devices with DCB feature settings to ensure consistent operation in a data center network.

DCBX is a prerequisite for using DCB features, such as priority-based flow control (PFC) and enhanced traffic selection (ETS), to exchange link-level configurations in a converged Ethernet environment. DCBX is also deployed in topologies that support lossless operation for FCoE or iSCSI traffic. In these scenarios, all network devices are DCBX-enabled (DCBX is enabled end-to-end). For more information about how these features are implemented and used on an MXL Switch, refer to:

- [Configuring Priority-Based Flow Control](#)
- [Configuring Enhanced Transmission Selection](#)
- [FIP Snooping](#)
- [iSCSI Optimization](#)

The following versions of DCBX are supported on an MXL Switch: CIN, CEE, and IEEE2.5.

Prerequisite: DCBX requires the LLDP to be enabled on all DCB devices.

DCBX Operation

DCBX performs the following operations:

- Discovers DCB configuration (such as PFC and ETS) in a peer device.
- Detects DCB misconfiguration in a peer device; that is, when DCB features are not compatibly configured on a peer device and the local switch. Misconfiguration detection is feature-specific because some DCB features support asymmetric configuration.
- Reconfigures a peer device with the DCB configuration from its configuration source if the peer device is willing to accept configuration.
- Accepts the DCB configuration from a peer if a DCBX port is in “willing” mode to accept a peer’s DCB settings and then internally propagates the received DCB configuration to its peer ports.

DCBX Port Roles

Use the following DCBX port roles to enable the auto-configuration of DCBX-enabled ports and propagate DCB configurations learned from peer DCBX devices internally to other switch ports:

- **Auto-upstream:** The port advertises its own configuration to DCBX peers and receives its configuration from DCBX peers (ToR or FCF device). The port also propagates its configuration to other ports on the switch.

The first auto-upstream that is capable of receiving a peer configuration is elected as the *configuration source*. The elected configuration source then internally propagates the configuration to auto-downstream ports. A port that receives an internally propagated configuration overwrites its local configuration with the new parameter values.

When an auto-upstream port (besides the configuration source) receives and overwrites its configuration with internally propagated information, one of the following actions is taken:

- If the peer configuration received is compatible with the internally propagated port configuration, the link with the DCBX peer is enabled.
- If the received peer configuration is not compatible with the currently configured port configuration, the link with the DCBX peer port is disabled and a syslog message for an incompatible configuration is generated. The network administrator must then reconfigure the peer device so that it advertises a compatible DCB configuration.

The configuration received from a DCBX peer or from an internally propagated configuration is not stored in the switch's running configuration.

On a DCBX port in an auto-upstream role, the PFC and application priority TLVs are enabled. ETS recommend TLVs are disabled and ETS configuration TLVs are enabled.

- **Auto-downstream -** The port advertises its own configuration to DCBX peers but is *not willing* to receive remote peer configuration. The port always accepts internally propagated configurations from a configuration source. An auto-downstream port that receives an internally propagated configuration overwrites its local configuration with the new parameter values.

When an auto-downstream port receives and overwrites its configuration with internally propagated information, one of the following actions is taken:

- If the peer configuration received is compatible with the internally propagated port configuration, the link with the DCBX peer is enabled.
- If the received peer configuration is not compatible with the currently configured port configuration, the link with the DCBX peer port is disabled and a syslog message for an incompatible configuration is generated. The network administrator must then reconfigure the peer device so that it advertises a compatible DCB configuration.

The internally propagated configuration is not stored in the switch's running configuration.

On a DCBX port in an auto-downstream role, all PFC, application priority, ETS recommend, and ETS configuration TLVs are enabled.

- **Configuration source -** The port is configured to serve as a source of configuration information on the switch. Peer DCB configurations received on the port are propagated to other DCBX auto-configured ports. If the peer configuration is compatible with a port configuration, DCBX is enabled on the port.

On a configuration-source port, the link with a DCBX peer is enabled when the port receives a DCB configuration that can be internally propagated to other auto-configured ports.

The configuration received from a DCBX peer is not stored in the switch's running configuration.

On a DCBX port that is the configuration source, all PFC and application priority TLVs are enabled. ETS recommend TLVs are disabled and ETS configuration TLVs are enabled.

- **Manual** - The port is configured to operate only with administrator-configured settings and does not auto-configure with DCB settings received from a DCBX peer or from an internally propagated configuration from the configuration source. If you enable DCBX, ports in Manual mode advertise their configurations to peer devices but do not accept or propagate internal or external configurations. Unlike other user-configured ports, the configuration of DCBX ports in Manual mode is saved in the running configuration.

On a DCBX port in a manual role, all PFC, application priority, ETS recommend, and ETS configuration TLVs are enabled.

Default DCBX port role: Manual.



Note: On a DCBX port, application priority TLV advertisements are handled as follows:

- The application priority TLV is transmitted only if the priorities in the advertisement match the configured PFC priorities on the port.
- On auto-upstream and auto-downstream ports:
 - If a configuration source is elected, the ports send an application priority TLV based on the application priority TLV received on the configuration-source port. When an application priority TLV is received on the configuration-source port, the auto-upstream and auto-downstream ports use the internally propagated PFC priorities to match against the received application priority. Otherwise, these ports use their locally configured PFC priorities in application priority TLVs.
 - If no configuration source is configured, auto-upstream and auto-downstream ports check to see that the locally configured PFC priorities match the priorities in a received application priority TLV.
- On manual ports: An application priority TLV is advertised only if the priorities in the TLV match the PFC priorities configured on the port.

DCB Configuration Exchange

On an MXL Switch, the DCBX protocol supports the exchange and propagation of configuration information for the following DCB features.

- Enhanced transmission selection (ETS)
- Priority-based flow control (PFC)

DCBX uses the following methods to exchange DCB configuration parameters:

- **Asymmetric:** DCB parameters are exchanged between a DCBX-enabled port and a peer port without requiring that a peer port and the local port use the same configured values for the configurations to be compatible. For example, ETS uses an asymmetric exchange of parameters between DCBX peers.
- **Symmetric:** DCB parameters are exchanged between a DCBX-enabled port and a peer port with the requirement that each configured parameter value is the same for the configurations to be compatible. For example, PFC uses an symmetric exchange of parameters between DCBX peers.

Configuration Source Election

When an auto-upstream or auto-downstream port receives a DCB configuration from a peer, the port first checks to see if there is an active configuration source on the switch.

- If a configuration source already exists, the received peer configuration is checked against the local port configuration. If the received configuration is compatible, the DCBX marks the port as DCBX-enabled. If the configuration received from the peer is not compatible, a warning message is logged and the DCBX frame error counter is incremented. Although DCBX is operationally disabled, the port keeps the peer link up and continues to exchange DCBX packets. If a compatible peer configuration is later received, DCBX is enabled on the port.
- If there is no configuration source, a port may elect itself as the configuration source. A port may become the configuration source if the following conditions exist:
 - No other port is the configuration source.
 - The port role is auto-upstream.
 - The port is enabled with link up and DCBX enabled.
 - The port has performed a DCBX exchange with a DCBX peer.
 - The switch is capable of supporting the received DCB configuration values through either a symmetric or asymmetric parameter exchange.

A newly elected configuration source propagates configuration changes received from a peer to the other auto-configuration ports. Ports receiving auto-configuration information from the configuration source ignore their current settings and use the configuration source information.

Propagation of DCB Information

When an auto-upstream or auto-downstream port receives a DCB configuration from a peer, the port acts as a DCBX client and checks if a DCBX configuration source exists on the switch.

- If a configuration source is found, the received configuration is checked against the currently configured values that are internally propagated by the configuration source. If the local configuration is compatible with the received configuration, the port is enabled for DCBX operation and synchronization.
- If the configuration received from the peer is not compatible with the internally propagated configuration used by the configuration source, the port is disabled as a client for DCBX operation and synchronization and a syslog error message is generated. The port keeps the peer link up and continues to exchange DCBX packets. If a compatible configuration is later received from the peer, the port is enabled for DCBX operation.



Note: DCB configurations internally propagated from a configuration source do not overwrite the configuration on a DCBX port in a manual role.

When a configuration source is elected, all auto-upstream ports other than the configuration source are marked as *willing disabled*. The internally propagated DCB configuration is refreshed on all auto-configuration ports and each port may begin configuration negotiation with a DCBX peer again.

Auto-Detection and Manual Configuration of the DCBX Version

When operating in Auto-Detection mode (**dcbx version auto** command in [DCBX Configuration Procedure](#)), a DCBX port automatically detects the DCBX version on a peer port. Legacy CIN and CEE versions are supported in addition to the standard IEEE version 2.5 DCBX.

A DCBX port detects a peer version after receiving a valid frame for that version. The local DCBX port reconfigures to operate with the peer version and maintains the peer version on the link until one of the following conditions occurs:

- The switch reboots.
- The link is reset (goes down and up).
- User-configured CLI commands require the version negotiation to restart.
- The peer times out.
- Multiple peers are detected on the link.

If you configure a DCBX port to operate with a specific version (**dcbx version {cee | cin | ieee-v2.5}** command in [DCBX Configuration Procedure](#)), DCBX operations are performed according to the configured version, including fast and slow transmit timers and message formats. If a DCBX frame with a different version is received, a syslog message is generated and the peer version is recorded in the peer status table. If the frame cannot be processed, it is discarded and the discard counter is incremented.



Note: Because DCBX TLV processing is best effort, it is possible that CIN frames may be processed when DCBX is configured to operate in CEE mode and vice versa. In this case, the unrecognized TLVs cause the unrecognized TLV counter to be incremented, but the frame is processed and is not discarded.

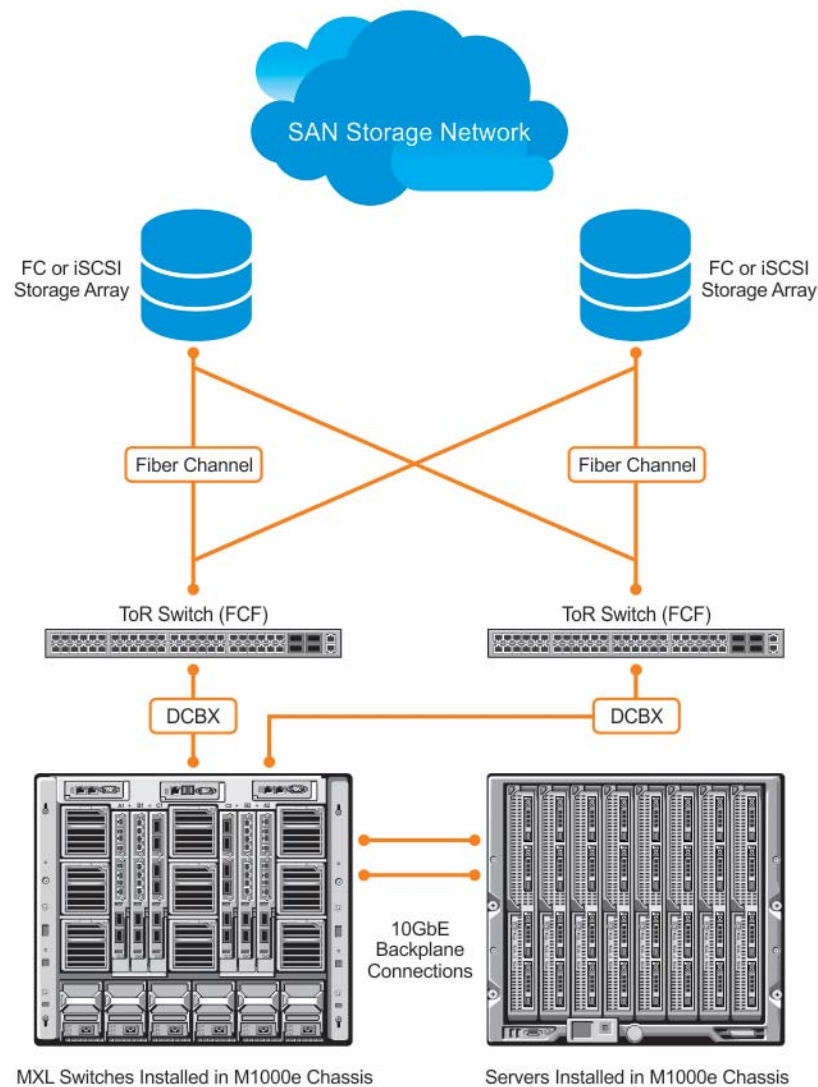
Legacy DCBX (CIN and CEE) supports the DCBX control state machine that is defined to maintain the sequence number and acknowledge number sent in the DCBX control TLVs.

DCBX Example

[Figure 8-4](#) shows how DCBX is used on an MXL Switch installed in a PowerEdge M1000e chassis in which servers are also installed.

- The external 40GbE ports on the base module (ports 33 and 37) of two switches are used for uplinks configured as DCBX auto-upstream ports. The MXL Switch is connected to third-party, top-of-rack (ToR) switches through 40GbE uplinks. The ToR switches are part of a Fibre Channel storage network.
- The internal ports (ports 1-32) connected to the 10GbE backplane are configured as auto-downstream ports.
- On the MXL Switch, PFC and ETS use DCBX to exchange link-level configuration with DCBX peer devices.

Figure 8-4. DCBX Sample Topology



DCBX Prerequisites and Restrictions

The following prerequisites and restrictions apply when you configure DCBX operation on a port:

- DCBX requires LLDP in both send (TX) and receive (RX) mode to be enabled on a port interface (**protocol lldp mode** command; refer to [Figure 18-7](#)). If a multiple DCBX peer ports are detected on a local DCBX interface, LLDP is shut down.
- The CIN version of DCBX supports only PFC, ETS, and FCOE; it does not support iSCSI, backward congestion management (BCN), logical link down (LLDF), and network interface virtualization (VIV).

DCBX Configuration Procedure

To configure an MXL Switch for DCBX operation in a data center network, you must:

1. Configure ToR- and FCF-facing interfaces as auto-upstream ports.
2. Configure server-facing interfaces as auto-downstream ports.
3. Configure a port to operate in a configuration-source role.
4. Configure ports to operate in a manual role.

To verify the DCBX configuration on a port, use the **show interface dcbx detail** command (Figure 8-16).

Configure DCBX operation at the interface level on a switch or globally on the switch.

Prerequisite: DCBX requires LLDP to be enabled to advertise DCBX TLVs to peers. For more information, refer to [Link Layer Discovery Protocol \(LLDP\)](#).

Configuring DCBX on an Interface

To configure DCBX operation on an interface, follow these steps:

Step	Task	Command	Command Mode
1	Enter INTERFACE Configuration mode.	<code>interface <i>type slot/port</i></code>	CONFIGURATION
2	Enter LLDP Configuration mode to enable DCBX operation.	<code>[no] protocol lldp</code>	INTERFACE
3	Configure the DCBX version used on the interface, where: auto configures the port to operate using the DCBX version received from a peer. <ul style="list-style-type: none"> • <code>cee</code> configures the port to use CEE (Intel 1.01). • <code>cin</code> configures the port to use Cisco-Intel-Nuova (DCBX 1.0). • <code>ieee-v2.5</code> configures the port to use IEEE 802.1Qaz (Draft 2.5). Default: Auto.	<code>[no] dcbx version {auto cee cin ieee-v2.5}</code>	PROTOCOL LLDP

Step	Task	Command	Command Mode
4	<p>Configure the DCBX port role used by the interface to exchange DCB information, where:</p> <ul style="list-style-type: none"> • auto-upstream configures the port to receive a peer configuration. The configuration source is elected from auto-upstream ports. • auto-downstream configures the port to accept the internally propagated DCB configuration from a configuration source. • config-source configures the port to serve as the configuration source on the switch. • manual configures the port to operate only on administer-configured DCB parameters. The port does not accept a DCB configuration received from a peer or a local configuration source. <p>Default: Manual.</p>	[no] dcbx port-role {config-source auto-downstream auto-upstream manual}	PROTOCOL LLDP
5	<p>On manual ports only: Configure the PFC and ETS TLVs advertised to DCBX peers, where:</p> <ul style="list-style-type: none"> • ets-conf enables the advertisement of ETS Configuration TLVs. • ets-reco enables the advertisement of ETS Recommend TLVs. • pfc enables the advertisement of PFC TLVs. Default: All PFC and ETS TLVs are advertised. <p>Note: You can configure the transmission of more than one TLV type at a time; for example: advertise dcbx-tlv ets-conf ets-reco. You can enable ETS recommend TLVs (ets-reco) only if ETS configuration TLVs (ets-conf) are enabled. To disable TLV transmission, use the no form of the command; for example, no advertise dcbx-tlv pfc ets-reco.</p>	[no] advertise dcbx-tlv {ets-conf ets-reco pfc} [ets-conf ets-reco pfc] [ets-conf ets-reco pfc]	PROTOCOL LLDP
6	<p>On manual ports only: Configure the Application Priority TLVs advertised on the interface to DCBX peers, where:</p> <ul style="list-style-type: none"> • fcoe enables the advertisement of FCoE in Application Priority TLVs. • iscsi enables the advertisement of iSCSI in Application Priority TLVs. <p>Default: Application Priority TLVs are enabled to advertise FCoE and iSCSI.</p> <p>Note: To disable TLV transmission, enter the no form of the command; for example, no advertise dcbx-appln-tlv iscsi.</p> <p>For information about how to use FCoE and iSCSI, refer to FIP Snooping and iSCSI Optimization.</p>	[no] advertise dcbx-appln-tlv {fcoe iscsi}	PROTOCOL LLDP

Configuring DCBX Globally on the Switch

To globally configure DCBX operation on a switch, follow these steps:

Step	Task	Command	Command Mode
1	Enter Global Configuration mode.	configure	EXEC PRIVILEGE
2	Enter LLDP Configuration mode to enable DCBX operation.	[no] protocol lldp	CONFIGURATION
3	<p>Configure the DCBX version used on all interfaces not already configured to exchange DCB information, where:</p> <ul style="list-style-type: none"> • auto configures all ports to operate using the DCBX version received from a peer. • cee configures a port to use CEE (Intel 1.01). cin configures a port to use Cisco-Intel-Nuova (DCBX 1.0). • ieee-v2.5 configures a port to use IEEE 802.1Qaz (Draft 2.5). <p>Default: Auto.</p>	[no] dcbx version {auto cee cin ieee-v2.5}	PROTOCL LLDP
<p>Note: You can configure the DCBX port role used by interfaces to exchange DCB information by using the dcbx port-role command in INTERFACE Configuration mode (see Step 3 in Configuring DCBX Globally on the Switch).</p>			
4	<p>(OPTIONAL) Configure the PFC and ETS TLVs to be advertised on unconfigured interfaces with a manual port-role, where:</p> <ul style="list-style-type: none"> • ets-conf enables transmission of ETS Configuration TLVs. • ets-reco enables transmission of ETS Recommend TLVs. • pfc enables transmission of PFC TLVs. <p>Note: You can configure the transmission of more than one TLV type at a time. ETS recommend TLVs (ets-reco) can be enabled only if ETS configuration TLVs (ets-conf) are enabled. To disable TLV transmission, use the no form of the command; for example, no advertise dcbx-tlv pfc ets-reco.</p> <p>Default: All TLV types are enabled.</p>	[no] advertise dcbx-tlv {ets-conf ets-reco pfc} [ets-conf ets-reco pfc] [ets-conf ets-reco pfc]	PROTOCOL LLDP

Step	Task	Command	Command Mode
5	<p>Configure the Application Priority TLVs to be advertised on unconfigured interfaces with a manual port-role, where:</p> <ul style="list-style-type: none"> • <code>fcoe</code> enables the advertisement of FCoE in Application Priority TLVs. • <code>iscsi</code> enables the advertisement of iSCSI in Application Priority TLVs. <p>Default: Application Priority TLVs are enabled and advertise FCoE and iSCSI.</p> <p>Note: To disable TLV transmission, use the <code>no</code> form of the command; for example, <code>no advertise dcbx-appln-tlv iscsi</code>.</p> <p>For information about how to use FCoE and iSCSI, refer to FIP Snooping and iSCSI Optimization.</p>	<code>[no] advertise dcbx-appln-tlv {fcoe iscsi}</code>	PROTOCOL LLDP
6	<p>Configure the FCoE priority advertised for the FCoE protocol in Application Priority TLVs. The priority-bitmap range is from 1 to FF. Default: 0x8.</p>	<code>[no] fcoe priority-bits priority-bitmap</code>	PROTOCOL LLDP
7	<p>Configure the iSCSI priority advertised for the iSCSI protocol in Application Priority TLVs. The priority-bitmap range is from 1 to FF. Default: 0x10.</p>	<code>[no] iscsi priority-bits priority-bitmap</code>	PROTOCOL LLDP

DCBX Error Messages

An error in DCBX operation is displayed using the following syslog messages:

`LLDP_MULTIPLE_PEER_DETECTED`: DCBX is operationally disabled after detecting more than one DCBX peer on the port interface.

`LLDP_PEER_AGE_OUT`: DCBX is disabled as a result of LLDP timing out on a DCBX peer interface.

`DSM_DCBX_PEER_VERSION_CONFLICT`: A local port expected to receive the IEEE, CIN, or CEE version in a DCBX TLV from a remote peer but received a different, conflicting DCBX version.

`DSM_DCBX_PFC_PARAMETERS_MATCH` and `DSM_DCBX_PFC_PARAMETERS_MISMATCH`: A local DCBX port received a compatible (match) or incompatible (mismatch) PFC configuration from a peer.

`DSM_DCBX_ETS_PARAMETERS_MATCH` and `DSM_DCBX_ETS_PARAMETERS_MISMATCH`: A local DCBX port received a compatible (match) or incompatible (mismatch) ETS configuration from a peer.

`LLDP_UNRECOGNISED_DCBX_TLV_RECEIVED`: A local DCBX port received an unrecognized DCBX TLV from a peer.

Debugging DCBX on an Interface

To enable DCBX debug traces for all or a specific control path, use the following command:

Task	Command	Command Mode
Enable DCBX debugging, where: <ul style="list-style-type: none"> • all: Enables all DCBX debugging operations. • auto-detect-timer: Enables traces for DCBX auto-detect timers. • config-exchng: Enables traces for DCBX configuration exchanges. • fail: Enables traces for DCBX failures. • mgmt: Enables traces for DCBX management frames. • resource: Enables traces for DCBX system resource frames. • sem: Enables traces for the DCBX state machine. • tlv: Enables traces for DCBX TLVs. 	<code>debug dcbx {all auto-detect-timer config-exchng fail mgmt resource sem tlv}</code>	EXEC PRIVILEGE

Verifying DCB Configuration

Use the **show** commands in [Table 8-2](#) to display DCB configurations.

Table 8-2. Displaying DCB Configurations

Command	Output
<code>show dot1p-queue mapping</code> (Figure 8-5)	Displays the current 802.1p priority-queue mapping.
<code>show dcb [stack-unit <i>unit-number</i>]</code> (Figure 8-6)	Displays data center bridging status, number of PFC-enabled ports, and number of PFC-enabled queues. On the master switch in a stack, you can specify a stack-unit number. Valid values: 0 to 5.
<code>show qos dcb-input [pfc-profile]</code> (Figure 8-7)	Displays the PFC configuration in a DCB input policy.
<code>show qos dcb-output [ets-profile]</code> (Figure 8-8)	Displays the ETS configuration in a DCB output policy.
<code>show qos priority-groups</code> (Figure 8-9)	Displays the ETS priority groups configured on the switch, including the 802.1p priority classes and ID of each group.
<code>show interface <i>port-type slot/port</i> pfc {summary detail}</code> (Figure 8-10)	Displays the PFC configuration applied to ingress traffic on an interface, including priorities and link delay. To clear PFC TLV counters, use the <code>clear pfc counters interface <i>port-type slot/port</i></code> command.
<code>show interface <i>port-type slot/port</i> pfc statistics</code> (Figure 8-11)	Displays counters for the PFC frames received and transmitted (by dot1p priority class) on an interface.
<code>show interface <i>port-type slot/port</i> ets {summary detail}</code> (Figure 8-12 and Figure 8-13)	Displays the ETS configuration applied to egress traffic on an interface, including priority groups with priorities and bandwidth allocation. To clear ETS TLV counters, enter the <code>clear ets counters interface <i>port-type slot/port</i></code> command.

Figure 8-5. show dot1p-queue mapping Command Example

```
FTOS(conf)# show dot1p-queue-mapping
Dot1p Priority: 0 1 2 3 4 5 6 7
Queue          : 0 0 0 1 2 3 3 3
```

Figure 8-6. show dcb Command Example

```
FTOS# show dcb
stack-unit 0 port-set 0
      DCB Status : Enabled
      PFC Port Count : 56 (current), 56 (configured)
      PFC Queue Count : 2 (current), 2 (configured)
```

Figure 8-7. show qos dcb-input Command Example

```
FTOS(conf)# show qos dcb-input
dcb-input pfc-profile
  pfc link-delay 32
  pfc priority 0-1
dcb-input pfc-profile1
  no pfc mode on
  pfc priority 6-7
```

Figure 8-8. show qos dcb-output Command Example

```
FTOS# show qos dcb-output
dcb-output ets
priority-group san qos-policy san
priority-group ipc qos-policy ipc
priority-group lan qos-policy lan
```

Figure 8-9. show qos priority-groups Command Example

```
FTOS#show qos priority-groups
priority-group ipc
priority-list 4
set-pgid 2
```

Figure 8-10. show interfaces pfc summary Command Example

```
FTOS# show interfaces tengigabitethernet 0/49 pfc summary
Interface TenGigabitEthernet 0/49
Admin mode is on
Admin is enabled
Remote is enabled, Priority list is 4
Remote Willing Status is enabled
Local is enabled
Oper status is Recommended
PFC DCBX Oper status is Up
State Machine Type is Feature
TLV Tx Status is enabled
PFC Link Delay 45556 pause quantams
Application Priority TLV Parameters :
-----
FCOE TLV Tx Status is disabled
ISCSI TLV Tx Status is disabled
Local FCOE PriorityMap is 0x8
Local ISCSI PriorityMap is 0x10
Remote FCOE PriorityMap is 0x8
Remote ISCSI PriorityMap is 0x8

FTOS# show interfaces tengigabitethernet 0/49 pfc detail
Interface TenGigabitEthernet 0/49
Admin mode is on
Admin is enabled
Remote is enabled
Remote Willing Status is enabled
Local is enabled
Oper status is recommended
PFC DCBX Oper status is Up
State Machine Type is Feature
TLV Tx Status is enabled
PFC Link Delay 45556 pause quanta
Application Priority TLV Parameters :
-----
FCOE TLV Tx Status is disabled
ISCSI TLV Tx Status is disabled
Local FCOE PriorityMap is 0x8
Local ISCSI PriorityMap is 0x10
Remote FCOE PriorityMap is 0x8
Remote ISCSI PriorityMap is 0x8

0 Input TLV pkts, 1 Output TLV pkts, 0 Error pkts, 0 Pause Tx pkts, 0 Pause Rx pkts
```

Table 8-3. show interface pfc summary Command Description

Field	Description
Interface	Interface type with stack-unit and port number.
Admin mode is on Admin is enabled	PFC Admin mode is on or off with a list of the configured PFC priorities. When PFC admin mode is on, PFC advertisements are enabled to be sent and received from peers; received PFC configuration takes effect. The admin operational status for a DCBX exchange of PFC configuration is enabled or disabled.
Remote is enabled, Priority list Remote Willing Status is enabled	Operational status (enabled or disabled) of peer device for DCBX exchange of PFC configuration with a list of the configured PFC priorities. Willing status of peer device for DCBX exchange (Willing bit received in PFC TLV): enabled or disabled.
Local is enabled	DCBX operational status (enabled or disabled) with a list of the configured PFC priorities.
Operational status (local port)	Port state for current operational PFC configuration: Init: Local PFC configuration parameters were exchanged with peer. Recommend: Remote PFC configuration parameters were received from peer. Internally propagated: PFC configuration parameters were received from configuration source.
PFC DCBX Oper status	Operational status for exchange of PFC configuration on local port: match (up) or mismatch (down).
State Machine Type	Type of state machine used for DCBX exchanges of PFC parameters: Feature - for legacy DCBX versions; Symmetric - for an IEEE version.
TLV Tx Status	Status of PFC TLV advertisements: enabled or disabled.
PFC Link Delay	Link delay (in quanta) used to pause specified priority traffic.
Application Priority TLV: FCOE TLV Tx Status	Status of FCoE advertisements in application priority TLVs from local DCBX port: enabled or disabled.
Application Priority TLV: ISCSI TLV Tx Status	Status of ISCSI advertisements in application priority TLVs from local DCBX port: enabled or disabled.
Application Priority TLV: Local FCOE Priority Map	Priority bitmap used by local DCBX port in FCoE advertisements in application priority TLVs.
Application Priority TLV: Local ISCSI Priority Map	Priority bitmap used by local DCBX port in ISCSI advertisements in application priority TLVs.
Application Priority TLV: Remote FCOE Priority Map	Status of FCoE advertisements in application priority TLVs from remote peer port: enabled or disabled.
Application Priority TLV: Remote ISCSI Priority Map	Status of iSCSI advertisements in application priority TLVs from remote peer port: enabled or disabled.
PFC TLV Statistics: Input TLV pkts	Number of PFC TLVs received.

Table 8-3. show interface pfc summary Command Description

Field	Description
PFC TLV Statistics: Output TLV pkts	Number of PFC TLVs transmitted.
PFC TLV Statistics: Error pkts	Number of PFC error packets received.
PFC TLV Statistics: Pause Tx pkts	Number of PFC pause frames transmitted.
PFC TLV Statistics: Pause Rx pkts	Number of PFC pause frames received

Figure 8-11. show interface pfc statistics Command Example

```

FTOS#show interfaces te 0/0 pfc statistics
Interface TenGigabitEthernet 0/0
Priority Received PFC Frames Transmitted PFC Frames
-----
0          0          0
1          0          0
2          0          0
3          0          0
4          0          0
5          0          0
6          0          0
7          0          0

```

Figure 8-12. show interface ets summary Command Example

```
FTOS(conf)# show interfaces te 0/0 ets summary
Interface TenGigabitEthernet 0/0
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters:
-----
Admin is enabled
TC-grp  Priority#      Bandwidth    TSA
0         0,1,2,3,4,5,6,7  100%         ETS
1         0%                ETS
2         0%                ETS
3         0%                ETS
4         0%                ETS
5         0%                ETS
6         0%                ETS
7         0%                ETS
Priority#      Bandwidth    TSA
0              13%         ETS
1              13%         ETS
2              13%         ETS
3              13%         ETS
4              12%         ETS
5              12%         ETS
6              12%         ETS
7              12%         ETS
Remote Parameters:
-----
Remote is disabled
Local Parameters:
-----
Local is enabled
TC-grp  Priority#      Bandwidth    TSA
0         0,1,2,3,4,5,6,7  100%         ETS
1         0%                ETS
2         0%                ETS
3         0%                ETS
4         0%                ETS
5         0%                ETS
6         0%                ETS
7         0%                ETS
Priority#      Bandwidth    TSA
0              13%         ETS
1              13%         ETS
2              13%         ETS
3              13%         ETS
4              12%         ETS
5              12%         ETS
6              12%         ETS
7              12%         ETS
Oper status is init
Conf TLV Tx Status is disabled
Traffic Class TLV Tx Status is disabled
```

Figure 8-13. show interface ets detail Command Example

```

FTOS(conf)# show interfaces tengigabitethernet 0/0 ets detail
Interface TenGigabitEthernet 0/0
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters :
-----
Admin is enabled
TC-grp   Priority#       Bandwidth      TSA
0        0,1,2,3,4,5,6,7  100%          ETS
1                0%            ETS
2                0%            ETS
3                0%            ETS
4                0%            ETS
5                0%            ETS
6                0%            ETS
7                0%            ETS

Priority#           Bandwidth      TSA
0                   13%           ETS
1                   13%           ETS
2                   13%           ETS
3                   13%           ETS
4                   12%           ETS
5                   12%           ETS
6                   12%           ETS
7                   12%           ETS

Remote Parameters:
-----
Remote is disabled

Local Parameters :
-----
Local is enabled
TC-grp   Priority#       Bandwidth      TSA
0        0,1,2,3,4,5,6,7  100%          ETS
1                0%            ETS
2                0%            ETS
3                0%            ETS
4                0%            ETS
5                0%            ETS
6                0%            ETS
7                0%            ETS

Priority#           Bandwidth      TSA
0                   13%           ETS
1                   13%           ETS
2                   13%           ETS
3                   13%           ETS
4                   12%           ETS
5                   12%           ETS
6                   12%           ETS
7                   12%           ETS

Oper status is init
Conf TLV Tx Status is disabled
Traffic Class TLV Tx Status is disabled
0 Input Conf TLV Pkts, 0 Output Conf TLV Pkts, 0 Error Conf TLV Pkts
0 Input Traffic Class TLV Pkts, 0 Output Traffic Class TLV Pkts, 0 Error Traffic Class TLV
Pkts

```

Table 8-4. show interface ets detail Command Description

Field	Description
Interface	Interface type with stack-unit and port number.
Max Supported TC Group	Maximum number of priority groups supported.
Number of Traffic Classes	Number of 802.1p priorities currently configured.
Admin mode	ETS mode: on or off. When on, the scheduling and bandwidth allocation configured in an ETS output policy or received in a DCBX TLV from a peer can take effect on an interface.
Admin Parameters	ETS configuration on local port, including priority groups, assigned dot1p priorities, and bandwidth allocation.
Remote Parameters	ETS configuration on remote peer port, including Admin mode (enabled if a valid TLV was received or disabled), priority groups, assigned dot1p priorities, and bandwidth allocation. If the ETS Admin mode is enabled on the remote port for DCBX exchange, the Willing bit received in ETS TLVs from the remote peer is included.
Local Parameters	ETS configuration on local port, including Admin mode (enabled when a valid TLV is received from a peer), priority groups, assigned dot1p priorities, and bandwidth allocation.
Operational status (local port)	Port state for current operational ETS configuration: Init: Local ETS configuration parameters were exchanged with peer. Recommend: Remote ETS configuration parameters were received from peer. Internally propagated: ETS configuration parameters were received from configuration source.
ETS DCBX Oper status	Operational status of ETS configuration on local port: match or mismatch.
State Machine Type	Type of state machine used for DCBX exchanges of ETS parameters: Feature - for legacy DCBX versions; Asymmetric - for an IEEE version.
Conf TLV Tx Status	Status of ETS Configuration TLV advertisements: enabled or disabled.
ETS TLV Statistic: Input Conf TLV pkts	Number of ETS Configuration TLVs received.
ETS TLV Statistic: Output Conf TLV pkts	Number of ETS Configuration TLVs transmitted.
ETS TLV Statistic: Error Conf TLV pkts	Number of ETS Error Configuration TLVs received.

Figure 8-14. show stack-unit all stack-ports all pfc details Command Example

```

FTOS(conf)# show stack-unit all stack-ports all pfc details

stack unit 0 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts

stack unit 1 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts

```

Figure 8-15. show stack-unit all stack-ports all ets details Command Example

```

FTOS(conf)# show stack-unit all stack-ports all ets details

Stack unit 0 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
Admin mode is on

Admin Parameters:
-----
Admin is enabled
TC-grp      Priority#      Bandwidth      TSA
-----
0           0,1,2,3,4,5,6,7  100%           ETS
1           -                -              -
2           -                -              -
3           -                -              -
4           -                -              -
5           -                -              -
6           -                -              -
7           -                -              -
8           -                -              -

Stack unit 1 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
Admin mode is on
Admin Parameters:
-----
Admin is enabled
TC-grp      Priority#      Bandwidth      TSA
-----
0           0,1,2,3,4,5,6,7  100%           ETS
1           -                -              -
2           -                -              -
3           -                -              -
4           -                -              -
5           -                -              -
6           -                -              -
7           -                -              -
8           -                -              -

```


Figure 8-16. show interface dcbx detail Command Example

```

FTOS(conf)# show interface tengigabitethernet 0/49 dcbx detail
FTOS#show interface te 0/49 dcbx detail

E-ETS Configuration TLV enabled           e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled          r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled           p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled    f-Application Priority for FCOE disabled
I-Application priority for iSCSI enabled    i-Application Priority for iSCSI disabled
-----

Interface TenGigabitEthernet 0/49
  Remote Mac Address 00:00:00:00:00:11
  Port Role is Auto-Upstream
  DCBX Operational Status is Enabled
  Is Configuration Source? TRUE

Local DCBX Compatibility mode is CEE
Local DCBX Configured mode is CEE
Peer Operating version is CEE
Local DCBX TLVs Transmitted: ErPfi

Local DCBX Status
-----
  DCBX Operational Version is 0
  DCBX Max Version Supported is 0
  Sequence Number: 2
  Acknowledgment Number: 2
  Protocol State: In-Sync

Peer DCBX Status:
-----
  DCBX Operational Version is 0
  DCBX Max Version Supported is 255
  Sequence Number: 2
  Acknowledgment Number: 2
  Total DCBX Frames transmitted 27
  Total DCBX Frames received 6
  Total DCBX Frame errors 0
  Total DCBX Frames unrecognized 0
  
```

Table 8-5. show interface dcbx detail Command Description

Field	Description
Interface	Interface type with chassis slot and port number.
Port-Role	Configured DCBX port role: auto-upstream, auto-downstream, config-source, or manual.
DCBX Operational Status	Operational status (enabled or disabled) used to elect a configuration source and internally propagate a DCB configuration. The DCBX operational status is the combination of PFC and ETS operational status.
Configuration Source	Specifies whether the port serves as the DCBX configuration source on the switch: true (yes) or false (no).

Table 8-5. show interface dcbx detail Command Description

Field	Description
Local DCBX Compatibility mode	DCBX version accepted in a DCB configuration as compatible. In auto-upstream mode, a port can only received a DCBX version supported on the remote peer.
Local DCBX Configured mode	DCBX version configured on the port: CEE, CIN, IEEE v2.5, or Auto (port auto-configures to use the DCBX version received from a peer).
Peer Operating version	DCBX version that the peer uses to exchange DCB parameters.
Local DCBX TLVs Transmitted	Transmission status (enabled or disabled) of advertised DCB TLVs (see TLV code at the top of the show command output).
Local DCBX Status: DCBX Operational Version	DCBX version advertised in Control TLVs.
Local DCBX Status: DCBX Max Version Supported	Highest DCBX version supported in Control TLVs.
Local DCBX Status: Sequence Number	Sequence number transmitted in Control TLVs.
Local DCBX Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs
Local DCBX Status: Protocol State	Current operational state of DCBX protocol: ACK or IN-SYNC.
Peer DCBX Status: DCBX Operational Version	DCBX version advertised in Control TLVs received from peer device.
Peer DCBX Status: DCBX Max Version Supported	Highest DCBX version supported in Control TLVs received from peer device.
Peer DCBX Status: Sequence Number	Sequence number transmitted in Control TLVs received from peer device.
Peer DCBX Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs received from peer device.
Total DCBX Frames transmitted	Number of DCBX frames sent from local port.
Total DCBX Frames received	Number of DCBX frames received from remote peer port.
Total DCBX Frame errors	Number of DCBX frames with errors received.
Total DCBX Frames unrecognized	Number of unrecognizable DCBX frames received.

PFC and ETS Configuration Examples

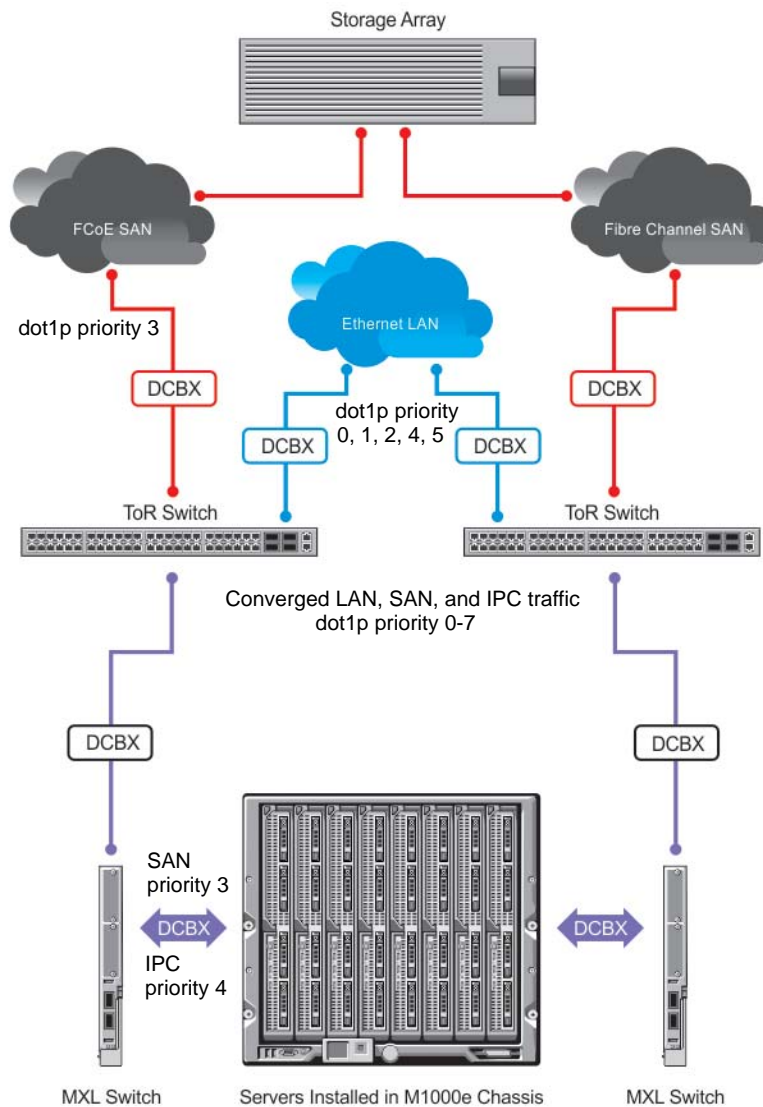
This section contains examples of how to configure and apply DCB input and output policies on an interface.

Using PFC and ETS to Manage Data Center Traffic

In the example shown in [Figure 8-17](#) for an MXL 10/40GbE Switch:

- Incoming SAN traffic is configured for priority-based flow control.
- Outbound LAN, IPC, and SAN traffic is mapped into three ETS priority groups and configured for enhanced traffic selection (bandwidth allocation and scheduling).
- One lossless queue is used.

Figure 8-17. Example: PFC and ETS Applied to LAN, IPC, and SAN Priority Traffic



QoS Traffic Classification: On the MXL Switch, the service-class dynamic dot1p command has been used in Global Configuration mode to map ingress dot1p frames to the queues shown in [Table 8-6](#). For more information, refer to [QoS dot1p Traffic Classification and Queue Assignment](#).

Table 8-6. Example: dot1p-Queue Assignment

dot1p Value in Incoming Frame	Queue Assignment
0	0
1	0
2	0
3	1
4	2
5	3
6	3
7	3

Lossless SAN traffic with dot1p priority 3 is assigned to queue 1. Other traffic types are assigned the 802.1p priorities shown in [Table 8-7](#) and the bandwidth allocations shown in [Table 8-8](#).

Table 8-7. Example: dot1p-priority class group Assignment

dot1p Value in Incoming Frame	Priority Group Assignment
0	LAN
1	LAN
2	LAN
3	SAN
4	IPC
5	LAN
6	LAN
7	LAN

Table 8-8. Example: priority group-bandwidth Assignment

Priority Group	Bandwidth Assignment
IPC	5%
SAN	50%
LAN	45%

Figure 8-18. PFC and ETS Configuration Command Example

Configure QoS priority-queue assignment to honor dot1p priorities or use L2 class maps to mark and map ingress traffic to output queues; for example:

```
FTOS(conf)# service-class dynamic dot1p
```

Or

```
FTOS(conf)# interface tengigabitethernet 0/1
FTOS(conf-if-te-0/1)# service-class dynamic dot1p
```

Configure a DCB input policy for applying PFC to lossless SAN priority traffic:

```
FTOS(conf)# dcb-input ipc_san_lan
FTOS(conf-qos-policy-in)# pfc mode on
FTOS(conf-qos-policy-in)# pfc priority 3
```

Configure an ETS priority group:

```
FTOS(conf)# priority-group san
FTOS(conf-pg)# priority-list 3
FTOS(conf-pg)# set-pgid 1
FTOS(conf-pg)# exit
FTOS(conf)# priority-group ipc
FTOS(conf-pg)# priority-list 4
FTOS(conf-pg)# set-pgid 2
FTOS(conf-pg)# exit
FTOS(conf)# priority-group lan
FTOS(conf-pg)# priority-list 0-2,5-7
FTOS(conf-pg)# set-pgid 3
FTOS(conf-pg)# exit
```

Configure an ETS output policy for egress traffic:

```
FTOS(conf)# qos-policy-output san ets
FTOS(conf-qos-policy-out)# bandwidth-percentage 50
FTOS(conf-qos-policy-out)# exit
FTOS(conf)# qos-policy-output lan ets
FTOS(conf-qos-policy-out)# bandwidth-percentage 45
FTOS(conf-qos-policy-out)# exit
FTOS(conf)# qos-policy-output ipc ets
FTOS(conf-qos-policy-out)# bandwidth-percentage 5
FTOS(conf-qos-policy-out)# exit
```

Figure 8-19. Example: DCB PFC and ETS Configuration (Continued)**Configure a DCB output policy for applying ETS (bandwidth allocation and scheduling) to IPC, SAN, and LAN priority traffic:**

```
FTOS(conf)# dcb-output ets
FTOS(conf-dcb-out)# priority-group san qos-policy san
FTOS(conf-dcb-out)# priority-group lan qos-policy lan
FTOS(conf-dcb-out)# priority-group ipc qos-policy ipc
```

Apply DCB input and output policies to a port interface:

```
FTOS(conf)# interface tengigabitethernet 0/1
FTOS(conf-if-te-0/1)# dcb-policy input pfc
FTOS(conf-if-te-0/1)# dcb-policy output ets
```

If the DCBX version is CIN, configure a QoS output policy to specify bandwidth allocation to different traffic types:

```
FTOS(conf)#qos-policy-output lan-q0
FTOS(conf-qos-policy-out)#bandwidth-percentage 30
FTOS(conf-qos-policy-out)#exit
FTOS(conf)#qos-policy-output lan-q3
FTOS(conf-qos-policy-out)#bandwidth-percentage 70
FTOS(conf-qos-policy-out)#exit
```

Create a QoS policy map for DCBX CIN bandwidth allocation:

```
FTOS(conf)# policy-map-output ets-queues
FTOS(conf-policy-map-out)# service-queue 0 qos-policy lan-q0
FTOS(conf-policy-map-out)# service-queue 3 qos-policy lan-q3
```

Apply the QoS policy map for DCBX CIN bandwidth allocation to an interface:

```
FTOS(conf-if-te-0/1)# service-policy output ets-queues
```

Using PFC and ETS to Manage Converged Ethernet Traffic in a Switch Stack

Figure 8-20 shows how to apply the DCB PFC input policy (ipc_san_lan) and ETS output policy (ets) configured in Figure 8-18 and Figure 8-19 on all ports on all MXL Switches in a switch stack.

Figure 8-20. Apply DCB PFC Input Policy and ETS Output Policy in a Switch Stack Example**On the stack master, apply DCB PFC input and ETS output policies to all port interfaces on stacked switches:**

```
FTOS(conf)# dcb-policy output stack-unit all stack-ports all ets
FTOS(conf)# dcb-policy input stack-unit all stack-ports all pfc
```

Hierarchical Scheduling in ETS Output Policies

On an MXL Switch, ETS supports up to three levels of hierarchical scheduling. For example, you can apply ETS output policies with the following configurations:

- Priority group 1 assigns traffic to one priority queue with 20% of the link bandwidth and strict-priority scheduling.
- Priority group 2 assigns traffic to one priority queue with 30% of the link bandwidth.
- Priority group 3 assigns traffic to two priority queues with 50% of the link bandwidth and strict-priority scheduling.

In this example, the configured ETS bandwidth allocation and scheduler behavior is as follows:

- Unused bandwidth usage: Normally, if there is no traffic or unused bandwidth for a priority group, the bandwidth allocated to the group is distributed to the other priority groups according to the bandwidth percentage allocated to each group. However, when three priority groups with different bandwidth allocations are used on an interface:
 - If priority group 3 has free bandwidth, it is distributed as follows: 20% of the free bandwidth to priority group 1 and 30% of the free bandwidth to priority group 2.
 - If priority group 1 or 2 has free bandwidth, (20 + 30)% of the free bandwidth is distributed to priority group 3. Priority groups 1 and 2 retain whatever free bandwidth remains up to the (20+ 30)%.
- Strict-priority groups: If two priority groups have strict-priority scheduling, traffic assigned from the priority group with the higher priority-queue number is scheduled first. However, when three priority groups are used and two groups have strict-priority scheduling (such as groups 1 and 3 in the example), the strict priority group whose traffic is mapped to one queue takes precedence over the strict priority group whose traffic is mapped to two queues.
Therefore, in the example, scheduling traffic to priority group 1 (mapped to one strict-priority queue) takes precedence over scheduling traffic to priority group 3 (mapped to two strict-priority queues).

Dynamic Host Configuration Protocol (DHCP)

This chapter contains the following sections:

- [Overview](#)
- [Implementation Information](#)
- [Configuration Tasks](#)
- [Configure the System to be a DHCP Server](#)
- [Configure the System to be a Relay Agent](#)
- [Configure the System to be a DHCP Client](#)
- [Configure Secure DHCP](#)

Overview

Dynamic host configuration protocol (DHCP) is an application layer protocol that dynamically assigns IP addresses and other configuration parameters to network end-stations (hosts) based on configuration policies determined by network administrators. DHCP:

- relieves network administrators of manually configuring hosts, which can be a tedious and error-prone process when hosts often join, leave, and change locations on the network.
- reclaims IP addresses that are no longer in use to prevent address exhaustion.

DHCP is based on a client-server model. A host discovers the DHCP server and requests an IP address, and the server either leases or permanently assigns one. There are three types of devices that are involved in DHCP negotiation:

- **DHCP Server**—a network device offering configuration parameters to the client.
- **DHCP Client**—a network device requesting configuration parameters from the server.
- **Relay agent**—an intermediary network device that passes DHCP messages between the client and server when the server is not on the same subnet as the host.

DHCP Packet Format and Options

DHCP uses the user datagram protocol (UDP) as its transport protocol. The server listens on port 67 and transmits to port 68; the client listens on port 68 and transmits to port 67. The configuration parameters are carried as options in the DHCP packet in type, length, value (TLV) format; many options are specified in RFC 2132. To limit the number parameters that servers must provide, hosts specify the parameters that they require, and the server sends only those; some common options are given in [Table 9-1](#).

Figure 9-1. DHCP Packet Format



Table 9-1. Common DHCP Options

Option	Code	Description
Subnet Mask	1	Specifies the clients subnet mask.
Router	3	Specifies the router IP addresses that may serve as the client's default gateway.
Domain Name Server	6	Specifies the DNS servers that are available to the client.
Domain Name	15	Specifies the domain name that client should use when resolving hostnames via DNS.
IP Address Lease Time	51	Specifies the amount of time that the client is allowed to use an assigned IP address.
DHCP Message Type	53	1: DHCPDISCOVER 2: DHCPOFFER 3: DHCPREQUEST 4: DHCPDECLINE 5: DHCPACK 6: DHCPNACK 7: DHCPRELEASE 8: DHCPINFORM
Parameter Request List	55	Clients use this option to tell the server which parameters it requires. It is a series of octets where each octet is DHCP option code.
Renewal Time	58	Specifies the amount of time after the IP address is granted that the client attempts to renew its lease with the <i>original</i> server.
Rebinding Time	59	Specifies the amount of time after the IP address is granted that the client attempts to renew its lease with <i>any</i> server, if the original server does not respond.
End	255	Signals the last option in the DHCP packet.

Assigning an IP Address Using DHCP


When a client joins a network:

1. The client initially broadcasts a **DHCPDISCOVER** message on the subnet to discover available DHCP servers. This message includes the parameters that the client requires and might include suggested values for those parameters.
2. Servers unicast or broadcast a **DHCPOFFER** message in response to the DHCPDISCOVER that offers to the client values for the requested parameters. Multiple servers might respond to a single DHCPDISCOVER; the client might wait a period of time and then act on the most preferred offer.
3. The client broadcasts a **DHCPREQUEST** message in response to the offer, requesting the offered values.
4. After receiving a DHCPREQUEST, the server binds the clients' unique identifier (the hardware address plus IP address) to the accepted configuration parameters and stores the data in a database called a *binding table*. The server then broadcasts a **DHCPACK** message, which signals to the client that it may begin using the assigned parameters.
5. When the client leaves the network, or the lease time expires, returns its IP address to the server in a **DHCPRELEASE** message.

There are additional messages that are used in case the DHCP negotiation deviates from the process previously described and shown in [Figure 9-2](#).

- **DHCPDECLINE**—A client sends this message to the server in response to a DHCPACK if the configuration parameters are unacceptable, for example, if the offered address is already in use. In this case, the client starts the configuration process over by sending a DHCPDISCOVER.
- **DHCPINFORM**—A client uses this message to request configuration parameters when it assigned an IP address manually rather than with DHCP. The server responds by unicast.
- **DHCPNAK**—A server sends this message to the client if it is not able to fulfill a DHCPREQUEST, for example, if the requested address is already in use. In this case, the client starts the configuration process over by sending a DHCPDISCOVER.

Figure 9-2. Assigning Network Parameters using DHCP



Implementation Information

- The Dell Force10 implementation of DHCP is based on RFC 2131 and RFC 3046.
- IP source address validation is a sub-feature of DHCP snooping; FTOS uses ACLs internally to implement this feature and as such, you cannot apply ACLs to an interface which has IP source address validation. If you configure IP source address validation on a member port of a VLAN, and then attempt to apply a access list to the VLAN, FTOS displays the first line in [Message 1](#). If you first apply an ACL to a VLAN, and then attempt enable IP source address validation on one of its member ports, FTOS displays the second line in [Message 1](#).

Message 1 DHCP Snooping with VLAN ACL Compatibility Error

```
% Error: Vlan member has access-list configured.
% Error: Vlan has an access-list configured.
```

- FTOS provides 40K entries that you can divide between leased addresses and excluded addresses. By extension, the maximum number of pools you can configure depends on the on the subnet mask that you give to each pool. FTOS displays an error message for configurations that exceed the allocated memory.
- Supports 4K DHCP snooping entries.
- Supports DAI on 16 VLANs per system.

Configuration Tasks

- [Configure the System to be a DHCP Server](#)
- [Configure the System to be a Relay Agent](#)
- [Configure Secure DHCP](#)

Configure the System to be a DHCP Server

A DHCP server is a network device that has been programmed to provide network configuration parameters to clients upon request. Servers typically serve many clients, making host management much more organized and efficient.

The key responsibilities of DHCP servers are:

1. **Address storage and management:** DHCP servers are the owners of the addresses used by DHCP clients. The server stores the addresses and manages their use, keeping track of which addresses have been allocated and which are still available.
2. **Configuration parameter storage and management:** DHCP servers also store and maintain other parameters that are sent to clients when requested. These parameters specify in detail how a client is to operate.

3. **Lease Management:** DHCP servers use leases to allocate addresses to clients for a limited time. The DHCP server maintains information about each of the leases, including lease length.
4. **Responding To client requests:** DHCP servers respond to different types of requests from clients, primarily, granting, renewing, and terminating leases.
5. **Providing administration services:** DHCP servers include functionality that allows an administrator to implement policies that govern how the DHCP performs its other tasks.

Configuration Tasks

To configure DHCP, an administrator must first set up a DHCP server and provide it with configuration parameters and policy information including IP address ranges, lease length specifications, and configuration data that DHCP hosts need.

Configuring the Dell Force10 system to be a DHCP server is a three-step process:

1. [Configure the Server for Automatic Address Allocation](#)
2. [Specify a Default Gateway](#)
3. [Enable DHCP Server](#)

Related Configuration Tasks

- [Configure a Method of Hostname Resolution](#)
- [Create Manual Binding Entries](#)
- [Debug DHCP Server](#)
- [DHCP Clear Commands](#)

Configure the Server for Automatic Address Allocation

Automatic address allocation is an address assignment method by which the DHCP server leases an IP address to a client from a pool of available addresses.

Create an IP Address Pool

An address pool is a range of IP addresses that may be assigned by the DHCP server. Address pools are indexed by subnet number.

To create an address pool, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Access the DHCP server CLI context.	<code>ip dhcp server</code>	CONFIGURATION
2	Create an address pool and give it a name.	<code>pool <i>name</i></code>	DHCP

Step	Task	Command Syntax	Command Mode
3	Specify the range of IP addresses from which the DHCP server may assign addresses. <ul style="list-style-type: none"> <i>network</i> is the subnet address. <i>prefix-length</i> specifies the number of bits used for the network portion of the address you specify. 	<i>network network /prefix-length</i> Prefix-length Range: 17 to 31	DHCP <POOL>
4	Display the current pool configuration.	show config	DHCP <POOL>

After an IP address is leased to a client, only that client may release the address. FTOS performs a IP + MAC source address validation to ensure that no client can release another clients address. This is a default behavior, and is separate from IP+MAC source address validation. For more information, refer to [IP+MAC Source Address Validation on page 185](#).

Exclude Addresses from the Address Pool

The DHCP server assumes that all IP addresses in a DHCP address pool are available for assigning to DHCP clients. You must specify the IP address that the DHCP server should not assign to clients.

Task	Command Syntax	Command Mode
Exclude an address range from DHCP assignment. The exclusion applies to all configured pools.	excluded-address	DHCP

Specify an Address Lease Time

Task	Command Syntax	Command Mode
Specify an address lease time for the addresses in a pool.	lease {days [hours] [minutes] infinite} Default: 24 hours	DHCP <POOL>

Specify a Default Gateway

The IP address of the default router must be on the same subnet as the client. To specify the default gateway:

Task	Command Syntax	Command Mode
Specify default gateway(s) for the clients on the subnet, in order of preference.	default-router <i>address</i>	DHCP <POOL>


Enable DHCP Server

DHCP server is disabled by default. To enable the DHCP server, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Enter the DHCP command-line context.	ip dhcp server	CONFIGURATION
2	Enable DHCP server.	no disable Default: Disabled	DHCP
3	Display the current DHCP configuration.	show config	DHCP

In [Figure 9-3](#), an IP phone is powered by power over ethernet (PoE) and has acquired an IP address from the Dell Force10 system, which is advertising link layer discover protocol (LLDP)-media endpoint discovery (MED). The leased IP address is displayed using the show ip dhcp binding command and confirmed with the show lldp neighbors command.

Figure 9-3. Configuring DHCP Server



Configure a Method of Hostname Resolution

Dell Force10 systems are capable of providing DHCP clients with parameters for two methods of hostname resolution.

Address Resolution using DNS

A domain is a group of networks. DHCP clients query domain name server (DNS) IP servers when they need to correlate host names to IP addresses. To specify the order of preference for the DNS servers, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Create a domain.	domain-name <i>name</i>	DHCP <POOL>
2	Specify in order of preference for the DNS servers that are available to a DHCP client.	dns-server <i>address</i>	DHCP <POOL>

Address Resolution using NetBIOS WINS

Windows internet naming service (WINS) is a name resolution service that Microsoft DHCP clients use to correlate host names to IP addresses within a group of networks. Microsoft DHCP clients can be one of four types of NetBIOS nodes: broadcast, peer-to-peer, mixed, or hybrid. To specify the NetBIOS WINS, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Specify the NetBIOS WINS name servers, in order of preference, that are available to Microsoft DHCP clients.	<code>netbios-name-server address</code>	DHCP <POOL>
2	Specify the NetBIOS node type for a Microsoft DHCP client. Dell Force10 recommends specifying clients as hybrid.	<code>netbios-node-type type</code>	DHCP <POOL>

Create Manual Binding Entries

An address binding is a mapping between the IP address and media access control (MAC) address of a client. The DHCP server assigns the client an available IP address automatically, and then creates an entry in the binding table. However, the administrator can manually create an entry for a client. Manual bindings are useful when you want to guarantee that particular network device receives a particular IP address. Manual bindings can be considered single-host address pools. There is no limit on the number of manual bindings, but you can only configure one manual binding per host.



Note: FTOS does not prevent you from using a network IP as a host IP; be sure to NOT use a network IP as a host IP.

To create a manual binding, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Create an address pool	<code>pool name</code>	DHCP
2	Specify the client IP address.	<code>host address</code>	DHCP <POOL>
3	Specify the client hardware address. <ul style="list-style-type: none"> • <i>hardware-address</i> is the client MAC address. • <i>type</i> is the protocol of the hardware platform. The default protocol is Ethernet. 	<code>hardware-address hardware-address type</code>	DHCP <POOL>

Debug DHCP Server

To display debug information, follow this step:

Task	Command Syntax	Command Mode
Display debug information for DHCP server.	debug ip dhcp server [events packets]	EXEC Privilege

DHCP Clear Commands

To clear DHCP binding entries, follow these steps:

Task	Command Syntax	Command Mode
Clear DHCP binding entries for the entire binding table.	clear ip dhcp binding	EXEC Privilege
Clear a DHCP binding entry for an individual IP address.	clear ip dhcp binding <i>ip address</i>	EXEC Privilege
Clear a DHCP address conflict.	clear ip dhcp conflict	EXEC Privilege
Clear DHCP server counters.	clear ip dhcp server statistics	EXEC Privilege

Configure the System to be a Relay Agent

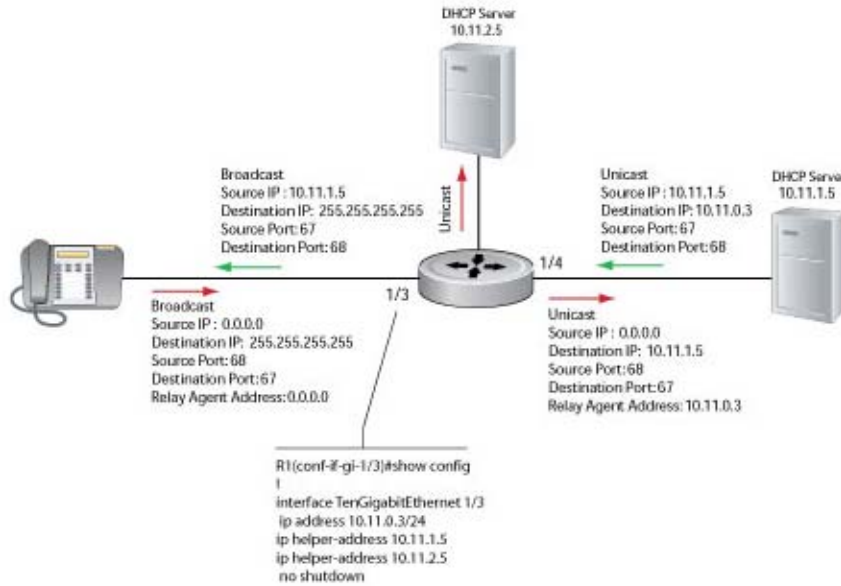
DHCP clients and servers request and offer configuration information using broadcast DHCP messages. Routers do not forward broadcasts, so if there are no DHCP servers on the subnet, the client does not receive a response to its request and therefore cannot access the network.

You can configure an interface on the Dell Force10 system to relay the DHCP messages to a specific DHCP server using the `ip helper-address dhcp-address` command from INTERFACE mode (Figure 9-4). Specify multiple DHCP servers by entering the `ip helper-address dhcp-address` command multiple times.

When you configure the `ip helper-address` command, the system listens for DHCP broadcast messages on port 67. The system rewrites packets received from the client and forwards it using unicast; the system rewrites the destination IP address and writes its own address as the relay device. Responses from the server are unicast back to the relay agent on port 68, and the relay agent rewrites the destination address and forwards the packet to the client subnet using broadcast.



Note: DHCP relay is not available on Layer 2 interfaces and VLANs.

Figure 9-4. Configuring Dell Force10 MXL 10/40GbE Switch IO Module system as a DHCP Relay Device

To view the ip helper-address configuration for an interface, use the show ip interface command from EXEC privilege mode (Figure 9-5).

Figure 9-5. Displaying the Helper Address Configuration

```
FTOS#show ip int tengig 1/3
TenGigabitEthernet 1/3 is up, line protocol is down
Internet address is 10.11.0.1/24
Broadcast address is 10.11.0.255
Address determined by user input
IP MTU is 1500 bytes
Helper address is 192.168.0.1
                  192.168.0.2
Directed broadcast forwarding is disabled
Proxy ARP is enabled
Split Horizon is enabled
Poison Reverse is disabled
ICMP redirects are not sent
ICMP unreachable are not sent
```

Configure the System to be a DHCP Client

A DHCP client is a network device that requests an IP address and configuration parameters from a DHCP server. On an MXL Switch, the DHCP client functionality is implemented as follows:

- The switch can obtain a dynamically-assigned IP address from a DHCP server. The switch does not receive a start-up configuration. To receive configuration parameters (FTOS version and a configuration file), you must use bare metal provisioning (BMP) on the switch (see [Bare Metal Provisioning \(BMP\) on page 101](#)). BMP is enabled on a switch as a factory-default setting.
A switch cannot operate with BMP and as a DHCP client at the same time. You can disable BMP by entering the stop jump-start command in EXEC mode. After BMP is stopped, the switch can act as a DHCP client.
- The dynamic IP address acquired by a DHCP client is for a limited period of time or until the client releases the address.
- A DHCP server manages and assigns IP addresses to clients from an address pool stored on the server (see [Create an IP Address Pool on page 163](#)).
- Dynamically-assigned IP addresses are supported only on Ethernet interfaces: 10-Gigabit, 40-Gigabit, and 100/1000/10000 Ethernet interfaces. DHCP client is supported on VLAN and port-channel interfaces
- The public out-of-band management interface and default VLAN 1 are configured, by default, as a DHCP client to acquire a dynamic IP address from a DHCP server.

To configure an interface to operate as a DHCP client to receive an IP address, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Enter interface configuration mode on an Ethernet interface.	<code>interface type slot/port</code>	CONFIGURATION
2	Configure the Ethernet interface to acquire its IP address from a DHCP network server. To release a DHCP-assigned IP address and remove the interface from being a DHCP client, enter the no ip address dhcp command.	<code>ip address dhcp</code>	INTERFACE

On an MXL switch configured as a DHCP client, you can release a dynamically-assigned IP address without removing the DHCP client operation on the interface. You can later manually acquire a new IP address from the DHCP server as follows:

Task	Command Syntax	Command Mode
Release a dynamically-acquired IP address while retaining the DHCP client configuration on the interface.	<code>release dhcp interface type slot/port</code>	EXEC Privilege
Acquire a new IP address with renewed lease time from a DHCP server.	<code>renew dhcp interface type slot/port</code>	EXEC Privilege

To display DHCP client information, enter the following **show** commands:

Task	Command Syntax	Command Mode
Display statistics about DHCP client interfaces (Figure 9-6).	show ip dhcp client statistics interface <i>type slot/port</i>	EXEC Privilege
Clear DHCP client statistics on a specified or on all interfaces.	clear ip dhcp client statistics {all interface <i>type slot/port</i> }	EXEC Privilege
Display lease information about the dynamic IP address currently assigned to a DHCP client interface (Figure 9-7).	show ip dhcp lease [interface <i>type slot/port</i>]	EXEC Privilege

Figure 9-6. show ip dhcp client statistics

```
FTOS# show ip dhcp client statistics interface tengigabitethernet 0/1
Message                Received
DHCPPOFFER             0
DHCPACK                0
DHCPNAK                0

Message                Sent
DHCPDISCOVER           0
DHCPREQUEST            0
DHCPDECLINE            0
DHCPRELEASE            0
DHCPREBIND             0
DHCPRENEW              0
```

Figure 9-7. show ip dhcp lease

```
FTOS# show ip dhcp lease interface tengigabitethernet 4/37

Interface  Lease-IP      Def-Router    ServerId      State      Lease Obtnd At      Lease Expires At
=====  =====
Te 4/37   189.17.9.2/30  0.0.0.0      189.17.9.1   BOUND     06-12-2012 07:35   01-18-2038 11:14

Renew Time      Rebind Time
=====
09-05-2023 04:56  11-06-2034 13:46
```

To enable debug messages for DHCP client operation, enter the following **debug** commands:

Task	Command Syntax	Command Mode
Enable the display of log messages for all DHCP packets sent and received on DHCP client interfaces.	[no] debug ip dhcp client packets [interface <i>type slot/port</i>]	EXEC Privilege
Enable the display of log messages for the following events on DHCP client interfaces: <ul style="list-style-type: none">• IP address acquisition• IP address release• Renewal of IP address and lease time• Release of an IP address	[no] debug ip dhcp client events [interface <i>type slot/port</i>]	EXEC Privilege

Figure 9-8 shows an example of the packet- and event-level debug messages displayed for the packet transmissions and state transitions on a DHCP client interface when you enable and disable a DHCP client.

Figure 9-8. DHCP Client: Debug Messages Logged during DHCP Client Enabling/Disabling

```

FTOS (conf-if-te-0/1)# ip address dhcp
May 27 15:52:46: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
DHCP ENABLE CMD Received in state START
May 27 15:52:48: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
Transitioned to state SELECTING
May 27 15:52:48: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP DISCOVER sent in Interface Te 0/1
May 27 15:52:48: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
Received DHCP OFFER packet in Interface Te 0/1 with Lease-Ip:10.16.134.250,
Mask:255.255.0.0,Server-Id:10.16.134.249
May 27 15:52:51: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
IP STATUS MESSAGE Received in state SELECTING status: 0
May 27 15:52:51: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
Transitioned to state REQUESTING
May 27 15:52:51: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP REQUEST sent in Interface Te 0/1
May 27 15:52:51: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
Received DHCPACK packet in InterfaceGi 0/1 with Lease-IP:10.16.134.250, Mask:255.255.0.0,
May 27 15:53:01: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
IP STATUS MESSAGE Received in state REQUESTING status: 0
May 27 15:53:01: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
Transitioned to state BOUND,IP Address: 10.16.134.250 Renewal in 2582 seconds

FTOS (conf-if-te-0/1)# no ip address dhcp
May 27 15:53:40: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
DHCP DISABLE CMD Received in state BOUND
May 27 15:53:40: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP RELEASE sent in Interface Te 0/1
May 27 15:53:40: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
Transitioned to state START
May 27 15:53:40: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
DHCP DISABLED CMD sent to FTOS in state START

FTOS#release dhcp int Te 0/1
FTOS#May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te
0/1 :DHCP RELEASE CMD Received in state BOUND
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT: DHCP RELEASE sent in
Interface Te 0/1
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1
:Transitioned to state STOPPED
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :DHCP IP
RELEASED CMD sent to FTOS in state STOPPED

FTOS#renew dhcp int te 0/1
FTOS#May 27 15:55:28: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :DHCP
RENEW CMD Received in state STOPPED
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1
:Transitioned to state SELECTING
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT: DHCP DISCOVER sent in
Interface Te 0/1
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT: Received DHCP OFFER packet
in Interface Te 0/1 with Lease-Ip:10.16.134.250, Mask:255.255.0.0,Server-Id:10.16.134.249

```

Figure 9-9 shows an example of the packet- and event-level debug messages displayed for the packet transmissions and state transitions on a DHCP client interface when you release and renew a DHCP client.

Figure 9-9. DHCP Client: Debug Messages Logged during DHCP Client Release/Renew

```
FTOS# release dhcp interface tengigabitethernet 0/1
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
DHCP RELEASE CMD Received in state BOUND
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP RELEASE sent in Interface Te 0/1
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
Transitioned to state STOPPED
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
DHCP IP RELEASED CMD sent to FTOS in state STOPPED

FTOS# renew dhcp interface tengigabitethernet 0/1
FTOS#May 27 15:55:28: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
DHCP RENEW CMD Received in state STOPPED
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT: Interface Te 0/1 :
Transitioned to state SELECTING
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP DISCOVER sent in Interface Te 0/1
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
Received DHCP OFFER packet in Interface Te 0/1 with Lease-Ip:10.16.134.250,
Mask:255.255.0.0,Server-Id:10.16.134.249
```



FTOS Behavior:

The **ip address dhcp** command enables DHCP server-assigned dynamic IP addresses on an interface. This setting persists after a switch reboot. If you enter the **shutdown** command on the interface, DHCP transactions are stopped and the dynamically-acquired IP address is saved. Use the **show interface type slot/port** command to display the dynamic IP address and DHCP as the mode of IP address assignment. If you later enter the **no shutdown** command and the lease timer for the dynamic IP address has expired, the IP address is unconfigured and the interface tries to acquire a new dynamic address from DHCP server.

If you later enter the **no shutdown** command and the lease timer for the dynamic IP address has expired, the IP address is released.

You cannot configure a secondary (backup) IP address on an interface using the **ip address dhcp** command; you must use the **ip address** command at the interface configuration level.

When you enter the **no ip address dhcp** command:

- The IP address dynamically acquired from a DHCP server is released from the interface.
- The DHCP client is disabled on the interface; it can no longer acquire a dynamic IP address from a DHCP server.
- DHCP packet transactions on the interface are stopped.

When you enter the **release dhcp** command, although the IP address that was dynamically-acquired from a DHCP server is released from an interface, the ability to acquire a new DHCP server-assigned address remains in the running configuration for the interface. To acquire a new IP address, enter either the **renew dhcp** command at the EXEC privilege level or the **ip address dhcp** command at the interface configuration level.

When you manually configures a static IP address on an interface (**ip address** command), you are prompted to release a dynamically-acquired IP address that already exists. If you confirm, the ability to receive a DHCP server-assigned IP address is also removed.

On an interface already configured with a static IP address:

- If you enter the **ip address dhcp** command to enable the acquisition of a dynamic IP address from a DHCP server, you are prompted to confirm the IP address reconfiguration. If you confirm, the statically-configured IP address is released.
- An error message is displayed if you enter the **release dhcp** or **renew dhcp** command.

On an interface already configured with a dynamic IP address:

- If you enter **renew dhcp** command, the lease time of the dynamically acquired IP address is renewed.

Important: To verify the currently configured dynamic IP address on an interface, enter the **show ip dhcp lease** command. The **show running-configuration** command output only displays `ip address dhcp`; the currently assigned dynamic IP address is not displayed.

DHCP Client on a Management Interface

When you enable a management interface to operate as a DHCP client, the following conditions apply:

- The management default route is added with the gateway as the router IP address received in the DHCP ACK packet. This is required to send and receive traffic to and from other subnets on the external network. This route is added irrespective both when the DHCP client and server are in the same or different subnets.

The management default route is deleted if the management IP address is released like other management routes added by the DHCP client.

- If "ip route for 0.0.0.0" is present or added later, it will take precedence.
- Management routes added by a DHCP client are displayed with Route Source as DHCP in **show ip management route** and **show ip management-route dynamic** command output.
- If a static IP route configured with the **ip route** command replaces a management route added by the DHCP client and then if the statically-configured IP route is removed (**no ip route** command), the management route added by DHCP is automatically re-installed. The management routes added by the DHCP client must be manually deleted.
- If a management route added by the DHCP client is removed or replaced by the same statically-configured management route, it is not re-installed unless you release the DHCP IP address and renew it on the management interface.
- A management route added by the DHCP client has higher precedence over the same statically-configured management route. If a dynamically-acquired management route added by the DHCP client overwrites a static management route, the static route is not removed from the running configuration.
- Management routes added by the DHCP client are not added to the running configuration.



Note: Management routes added by the DHCP client include the specific routes to reach a DHCP server in a different subnet and the default management route.

DHCP Client Operation with other Features

Stacking

The DHCP client daemon runs only on the master unit and handles all DHCP packet transactions. The DHCP client running on the master unit periodically synchronizes the lease file with the standby unit.

When a stack failover occurs, the new master requests the same DHCP server-assigned IP address on DHCP client interfaces. On non-bound interfaces, the new master re-initiates a DHCP packet transaction by sending a DHCP discovery packet.

VLT

A DHCP client is not supported on VLT interfaces.

VLAN and Port Channels

DHCP client configuration and behavior is the same on port-channel (LAG) and VLAN interfaces as on a physical interface.

DHCP Snooping

A DHCP client can run at the same on a switch with the DHCP snooping feature as follows:

- If you enable DHCP snooping globally on the switch and DHCP client on an interface, the trust port, source MAC address, and snooping table validations are not performed on the interface by DHCP snooping for packets destined to the DHCP client daemon.

Packets destined for the DHCP client are determined by following criteria:

- DHCP is enabled on the interface
 - UDP destination port in the packet is 68.
 - chaddr in the DHCP header of the packet is the same as the interface's MAC address.
- An entry in the DHCP snooping table is not added for a DHCP client interface.

DHCP Server

A switch can operate as a DHCP client and a DHCP server with the exception that a DHCP client interface does not acquire a dynamic IP address from the DHCP server running on the switch. A dynamic IP address must be acquired from another DHCP server.

VRRP

You cannot enable DHCP client on an interface and set the priority to 255 or assign the same IP address acquired by DHCP to a VRRP virtual group. Setting the priority to 255 or assigning an interface IP address to a VRRP virtual group guarantees that this router becomes the VRRP group owner.

If DHCP client is enabled in an interface, if the interface is added to a VRRP group, and if you want to use this router as the VRRP owner, you must assign a priority that is less than 255 and that is the highest priority among all the priorities assigned in the group.

Configure Secure DHCP

DHCP as defined by RFC 2131 provides no authentication or security mechanisms. Secure DHCP is a suite of features that protects networks that use dynamic address allocation from spoofing and attacks.

- [Option 82](#)
- [DHCP Snooping](#)
- [Dynamic ARP Inspection](#)
- [Source Address Validation](#)

Option 82

RFC 3046 (the relay agent information option, or Option 82) is used for class-based IP address assignment.

The code for the relay agent information option is 82 and is comprised of two sub-options, circuit ID and remote ID.

- **Circuit ID** is the interface on which the client-originated message is received.
- **Remote ID** identifies the host from which the message is received. The value of this sub-option is the MAC address of the relay agent that adds Option 82.

The DHCP relay agent inserts Option 82 before forwarding DHCP packets to the server. The server can use this information to:

- track the number of address requests per relay agent; restricting the number of addresses available per relay agent that can harden a server against address exhaustion attacks.
- associate client MAC addresses with a relay agent to prevent offering an IP address to a client spoofing the same MAC address on a different relay agent.
- assign IP addresses according to the relay agent. This prevents generating DHCP offers in response to requests from an unauthorized relay agent.

The server echoes the option back to the relay agent in its response, and the relay agent can use the information in the option to forward a reply out the interface on which the request was received rather than flooding it on the entire VLAN.

The relay agent strips Option 82 from DHCP responses before forwarding them to the client:

Task	Command Syntax	Command Mode
Insert Option 82 into DHCP packets. For routers between the relay agent and the DHCP server, enter the trust-downstream option.	ip dhcp relay information-option [trust-downstream]	CONFIGURATION

DHCP Snooping

DHCP snooping protects networks from spoofing. In the context of DHCP snooping, all ports are either trusted or untrusted. By default, all ports are untrusted. Trusted ports are ports through which attackers cannot connect. Manually configure ports connected to legitimate servers and relay agents as trusted.

When you enable DHCP snooping, the relay agent builds a binding table—using DHCPACK messages—containing the client MAC address, IP addresses, IP address lease time, port, VLAN ID, and binding type. Every time the relay agent receives a DHCPACK on an trusted port, it adds an entry to the table.

The relay agent then checks all subsequent DHCP client-originated IP traffic (DHCPRELEASE, DHCPNACK, and DHCPDECLINE) against the binding table to ensure that the MAC-IP address pair is legitimate, and that the packet arrived on the correct port. Packets that do not pass this check are dropped. This check-point prevents an attacker from spoofing a client and declining or releasing the real client's address. Server-originated packets (DHCPOFFER, DHCPACK, DHCPNACK) that arrive on an untrusted port are also dropped. This check-point prevents an attacker from impersonating as a DHCP server to facilitate a man-in-the-middle (MITM) attack.

Binding table entries are deleted when a lease expires, or the relay agent encounters a DHCPRELEASE, DHCPNACK, DHCPDECLINE.



FTOS Behavior: Introduced in FTOS version 7.8.1.0, DHCP snooping was available for Layer 3 only and dependent on DHCP relay agent (ip helper-address). FTOS version 8.2.1.0 extends DHCP snooping to Layer 2. You do not have to enable relay agent to snoop on Layer 2 interfaces.

FTOS Behavior: Binding table entries are deleted when a lease expires or when the relay agent encounters a DHCPRELEASE. The switch maintains a list of snooped VLANs. When the binding table is exhausted, DHCP packets are dropped on snooped VLANs, while these packets are forwarded across non-snooped VLANs. Because DHCP packets are dropped, no new IP address assignments are made. However, DHCPRELEASE and DHCPDECLINE packets are allowed so that the DHCP snooping table can decrease in size. After the table usage falls below the maximum limit of 4000 entries, new IP address assignments are allowed.



Note: DHCP server packets are dropped on all untrusted interfaces of a system configured for DHCP snooping. To prevent these packets from being dropped, configure ip dhcp snooping trust on the server-connected port.

Enable DHCP Snooping

To enable DHCP snooping, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Enable DHCP snooping globally.	ip dhcp snooping	CONFIGURATION
2	Specify ports connected to DHCP servers as trusted.	ip dhcp snooping trust	INTERFACE
3	Enable DHCP snooping on a VLAN.	ip dhcp snooping vlan	CONFIGURATION

Add a Static Entry in the Binding Table

To add a static entry in the binding table, follow this step:

Task	Command Syntax	Command Mode
Add a static entry in the binding table.	ip dhcp snooping binding mac	EXEC Privilege

Clear the Binding Table

To clear the binding table, follow this step:

Task	Command Syntax	Command Mode
Delete all of the entries in the binding table	clear ip dhcp snooping binding	EXEC Privilege

Display the Contents of the Binding Table

To display the contents of the binding table, follow this step:

Task	Command Syntax	Command Mode
Display the contents of the binding table.	show ip dhcp snooping	EXEC Privilege

To view the DHCP snooping statistics, use the `show ip dhcp snooping` command (Figure 9-10).

Figure 9-10. Command example: show ip dhcp snooping

```
FTOS#show ip dhcp snooping

IP DHCP Snooping                : Disabled.
IP DHCP Snooping Mac Verification : Disabled.
IP DHCP Relay Information-option  : Disabled.
IP DHCP Relay Trust Downstream   : Enabled.

Database write-delay (In minutes) : 0

DHCP packets information
Relay Information-option packets  : 0
Relay Trust downstream packets   : 0
Snooping packets                 : 0

Packets received on snooping disabled L3 Ports : 0
Snooping packets processed on L2 vlans        : 0

DHCP Binding File Details
Invalid File                               : 0
Invalid Binding Entry                     : 0
Binding Entry lease expired               : 0
FTOS#
```

Drop DHCP Packets on Snooped VLANs Only

Binding table entries are deleted when a lease expires or the relay agent encounters a DHCPRELEASE.

Starting with FTOS Release 8.2.1.1, line cards maintain a list of snooped VLANs. When the binding table fills, DHCP packets are dropped only on snooped-VLANs, while such packets are forwarded across non-snooped VLANs. Because DHCP packets are dropped, no new IP address assignments are made. However, DHCP release and decline packets are allowed so that the DHCP snooping table can decrease in size. After the table usage falls below the max limit of 4000 entries, new IP address assignments are allowed.

To view the number of entries in the table, use the `show ip dhcp snooping binding` command. This output displays the snooping binding table created using the ACK packets from the trusted port (Figure 9-11).

Figure 9-11. Command example: show ip dhcp snooping binding

```
FTOS#show ip dhcp snooping binding

Codes : S - Static D - Dynamic

IP Address      MAC Address      Expires(Sec)  Type  VLAN  Interface
=====
10.1.1.251      00:00:4d:57:f2:50  172800        D     V1 10  Te 0/2
10.1.1.252      00:00:4d:57:e6:f6  172800        D     V1 10  Te 0/1
10.1.1.253      00:00:4d:57:f8:e8  172740        D     V1 10  Te 0/3
10.1.1.254      00:00:4d:69:e8:f2  172740        D     V1 10  Te 0/50

Total number of Entries in the table : 4
```

Dynamic ARP Inspection

Dynamic address resolution protocol (ARP) inspection prevents ARP spoofing by forwarding only ARP frames that have been validated against the DHCP binding table.

ARP is a stateless protocol that provides no authentication mechanism. Network devices accept ARP request and replies from any device. ARP replies are accepted even when no request was sent. If a client receives an ARP message for which a relevant entry already exists in its ARP cache, it overwrites the existing entry with the new information.

The lack of authentication in ARP makes it vulnerable to spoofing. ARP spoofing is a technique attackers use to inject false IP to MAC mappings into the ARP cache of a network device. It is used to launch man-in-the-middle (MITM), and denial-of-service (DoS) attacks, among others.

A spoofed ARP message is one in which the MAC address in the sender hardware address field and the IP address in the sender protocol field are strategically chosen by the attacker. For example, in an MITM attack, the attacker sends a client an ARP message containing the attacker's MAC address and the gateway's IP address. The client then thinks that the attacker is the gateway and sends all internet-bound packets to it. Likewise, the attacker sends the gateway an ARP message containing the attacker's MAC address and the client's IP address. The gateway then thinks that the attacker is the client and forwards all packets addressed to the client to it. As a result, the attacker is able to sniff all packets to and from the client.

Other attacks using ARP spoofing include:

- broadcast—an attacker can broadcast an ARP reply that specifies FF:FF:FF:FF:FF:FF as the gateway's MAC address, resulting in all clients broadcasting all internet-bound packets.
- MAC flooding—an attacker can send fraudulent ARP messages to the gateway until the ARP cache is exhausted, after which, traffic from the gateway is broadcast.

- denial of service—an attacker can send fraudulent ARP messages to a client to associate a false MAC address with the gateway address, which blackholes all internet-bound packets from the client.



Note: Dynamic ARP inspection (DAI) uses entries in the L2SysFlow CAM region, a sub-region of SystemFlow. One CAM entry is required for every DAI-enabled VLAN. You can enable DAI on up to 16 VLANs on a system. You can configure 10 to 16 DAI-enabled VLANs by allocating more CAM space to the L2SysFlow region before enabling DAI.

Note: SystemFlow has 102 entries by default. This region is comprised of two sub-regions: L2Protocol and L2SystemFlow. L2Protocol has 87 entries; L2SystemFlow has 15 entries. Six L2SystemFlow entries are used by Layer 2 protocols, leaving 9 for DAI. L2Protocol can have a maximum of 100 entries. This region must be expanded to capacity before you can increase the size of L2SystemFlow. This is relevant when you are enabling DAI on VLANs. If, for example, you want to enable DAI on 16 VLANs, you need seven more entries; in this case, reconfigure the SystemFlow region for 122 entries:

```
layer-2 eg-acl value fib value frp value ing-acl value learn value l2pt value qos value system-flow 122
```

Note: The logic is as follows:

L2Protocol has 87 entries by default and must be expanded to its maximum capacity, 100 entries, before L2SystemFlow can be increased; therefore 13 more L2Protocol entries are required. L2SystemFlow has 15 entries by default, but only nine are for DAI; to enable DAI on 16 VLANs, seven more entries are required:

87 L2Protocol + 13 additional L2Protocol + 15 L2SystemFlow + 7 additional L2SystemFlow equals 122.

Step	Task	Command Syntax	Command Mode
1	Enable DHCP snooping.		
2	Validate ARP frames against the DHCP snooping binding table.	arp inspection	INTERFACE VLAN



Note: Dynamic ARP Inspection (DAI) may sometimes filter ARP traffic from valid clients in the DHCP snooping binding table.

To view the number of entries in the ARP database, use the show arp inspection database command (Figure 9-12).

Figure 9-12. Command example: show arp inspection database

```
FTOS#show arp inspection database
```

Protocol	Address	Age(min)	Hardware Address	Interface	VLAN	CPU
Internet	10.1.1.251	-	00:00:4d:57:f2:50	Te 0/2	V1 10	CP
Internet	10.1.1.252	-	00:00:4d:57:e6:f6	Te 0/1	V1 10	CP
Internet	10.1.1.253	-	00:00:4d:57:f8:e8	Te 0/3	V1 10	CP
Internet	10.1.1.254	-	00:00:4d:69:e8:f2	Te 0/50	V1 10	CP

FTOS#

To see how many valid and invalid ARP packets have been processed, use the `show arp inspection statistics` command (Figure 9-13).

Figure 9-13. Command example: show arp inspection database

```
FTOS#show arp inspection statistics

Dynamic ARP Inspection (DAI) Statistics
-----
Valid ARP Requests           : 0
Valid ARP Replies           : 1000
Invalid ARP Requests        : 1000
Invalid ARP Replies         : 0
FTOS#
```

Bypass the ARP Inspection

You can configure a port to skip ARP inspection by defining the interface as trusted, which is useful in multi-switch environments. ARPs received on trusted ports bypass validation against the binding table. All ports are untrusted by default.

Task	Command Syntax	Command Mode
Specify an interface as trusted so that ARPs are not validated against the binding table.	<code>arp inspection-trust</code>	INTERFACE



FTOS Behavior: Introduced in FTOS version 8.2.1.0, DAI was available for Layer 3 only. FTOS version 8.2.1.1 extends DAI to Layer 2.

Source Address Validation

Using the DHCP binding table, FTOS can perform three types of source address validation (SAV):

- [IP Source Address Validation on page 185](#): prevents IP spoofing by forwarding only IP packets that have been validated against the DHCP binding table.
- [DHCP MAC Source Address Validation on page 185](#): verifies a DHCP packet's source hardware address matches the client hardware address field (CHADDR) in the payload.
- [IP+MAC Source Address Validation on page 185](#): verifies that the IP source address and MAC source address are a legitimate pair.

IP Source Address Validation

IP source address validation prevents IP spoofing by forwarding only IP packets that have been validated against the DHCP binding table. A spoofed IP packet is one in which the IP source address is strategically chosen to disguise the attacker. For example, using ARP spoofing, an attacker can assume a legitimate client's identity and receive traffic addressed to it. Then the attacker can spoof the client's IP address to interact with other clients.

The DHCP binding table associates addresses assigned by the DHCP servers, with the port on which the requesting client is attached. When you enable IP source address validation on a port, the system verifies that the source IP address is one that is associated with the incoming port. If an attacker is impersonating a legitimate client, the source address appears on the wrong ingress port and the system drops the packet. Likewise, if the IP address is fake, the address will not be on the list of permissible addresses for the port, and the packet is dropped.

To enable IP source address validation, follow this step:

Task	Command Syntax	Command Mode
Enable IP Source Address Validation	<code>ip dhcp source-address-validation</code>	INTERFACE

DHCP MAC Source Address Validation

DHCP MAC source address validation validates a DHCP packet's source hardware address against the client hardware address field (CHADDR) in the payload.

FTOS Release 8.2.1.1 ensures that the packet's source MAC address is checked against the CHADDR field in the DHCP header only for packets from snooped VLANs.

To enable DHCP MAC source address validation, follow this step:

Task	Command Syntax	Command Mode
Enable DHCP MAC Source Address Validation.	<code>ip dhcp snooping verify mac-address</code>	CONFIGURATION

IP+MAC Source Address Validation

IP source address validation validates the IP source address of an incoming packet against the DHCP snooping binding table. IP+MAC source address validation ensures that the IP source address and MAC source address are a legitimate pair, rather than validating each attribute individually. You cannot configure IP+MAC source address validation with IP source address validation.

To enable IP+MAC source address validation, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Allocate at least one FP block to the ipmacacl CAM region.	cam-acl l2acl	CONFIGURATION
2	Save the running-config to the startup-config.	copy running-config startup-config	EXEC Privilege
3	Reload the system.	reload	EXEC Privilege
4	Enable IP+MAC Source Address Validation.	ip dhcp source-address-validation ipmac	INTERFACE

FTOS creates an ACL entry for each IP+MAC address pair in the binding table and applies it to the interface.

To display the IP+MAC ACL, follow this step:

Task	Command Syntax	Command Mode
Display the IP+MAC ACL for an interface for the entire system.	show ip dhcp snooping source-address-validation [interface]	EXEC Privilege

FIP Snooping

FIP snooping is supported on the MXL 10/40GbE Switch.

This chapter describes the FIP snooping concepts and configuration procedures:

- [Fibre Channel over Ethernet](#)
- [Ensuring Robustness in a Converged Ethernet Network](#)
- [FIP Snooping on Ethernet Bridges](#)
- [FIP Snooping in a Switch Stack](#)
- [Configuring FIP Snooping](#)
- [Displaying FIP Snooping Information](#)
- [FIP Snooping Configuration Example](#)

Fibre Channel over Ethernet

Fibre Channel over Ethernet (FCoE) provides a converged Ethernet network that allows the combination of storage-area network (SAN) and LAN traffic on a Layer 2 link by encapsulating Fibre Channel data into Ethernet frames.

FCoE works with the Ethernet enhancements provided in data center bridging (DCB) to support lossless (no-drop) SAN and LAN traffic. In addition, DCB provides flexible bandwidth sharing for different traffic types, such as LAN and SAN, according to 802.1p priority classes of service. For more information, refer to the [Data Center Bridging \(DCB\)](#) chapter.

Ensuring Robustness in a Converged Ethernet Network

Fibre Channel networks used for SAN traffic employ switches that operate as trusted devices. End devices log into the switch to which they are attached in order to communicate with other end devices attached to the Fibre Channel network. Because Fibre Channel links are point-to-point, a Fibre Channel switch controls all storage traffic that an end device sends and receives over the network. As a result, the switch can enforce zoning configurations, ensure that end devices use their assigned addresses, and secure the network from unauthorized access and denial-of-service attacks.

To ensure similar Fibre Channel robustness and security with FCoE in an Ethernet cloud network, the Fibre Channel over Ethernet initialization protocol (FIP) establishes virtual point-to-point links between FCoE end-devices (server ENodes and target storage devices) and FCoE forwarders (FCFs) over transit FCoE-enabled bridges.

Ethernet bridges commonly provide access control list (ACLs) that can emulate a point-to-point link by providing the traffic enforcement required to create a Fibre Channel-level of robustness. In addition, FIP serves as a Layer 2 protocol to:

- Operate between FCoE end-devices and FCFs over intermediate Ethernet bridges to prevent unauthorized access to the network and achieve the required security.
- Allow transit Ethernet bridges to efficiently monitor FIP frames passing between FCoE end-devices and an FCF, and use the FIP snooping data to dynamically configure ACLs on the bridge to only permit traffic authorized by the FCF.

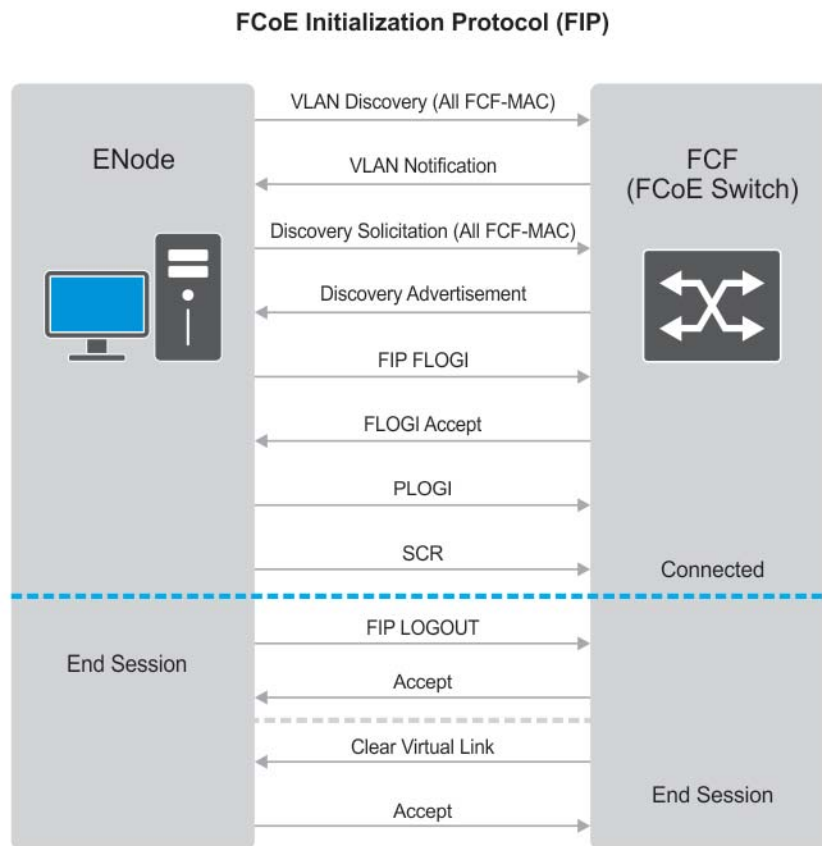
FIP enables FCoE devices to discover one another, initialize and maintain virtual links over an Ethernet network, and access storage devices in a storage area network. FIP satisfies the Fibre Channel requirement for point-to-point connections by creating a unique virtual link for each connection between an FCoE end-device and an FCF via a transit switch.

FIP provides functionality for discovering and logging in to an FCF. After discovering and logging in, FIP allows FCoE traffic to be sent and received between FCoE end-devices (ENodes) and the FCF. FIP uses its own EtherType and frame format. [Figure 10-1](#) shows the communication that occurs between an ENode server and an FCoE switch (FCF).

FIP performs the following functions:

- FIP virtual local area network (VLAN) discovery: FCoE devices (ENodes) discover the FCoE VLANs on which to transmit and receive FIP and FCoE traffic.
- FIP discovery: FCoE end-devices and FCFs are automatically discovered.
- Initialization: FCoE devices perform fabric login (FLOGI) and fabric discovery (FDISC) to create a virtual link with an FCoE switch.
- Maintenance: A valid virtual link between an FCoE device and an FCoE switch is maintained and the link termination logout (LOGO) functions properly.

Figure 10-1. FIP discovery and login between an ENode and an FCF



FIP Snooping on Ethernet Bridges

In a converged Ethernet network, intermediate Ethernet bridges can snoop on FIP packets during the login process on an FCF. Then, using ACLs, a transit bridge can permit only authorized FCoE traffic to be transmitted between an FCoE end-device and an FCF. An Ethernet bridge that provides these functions is called a FIP snooping bridge (FSB).

On a FIP snooping bridge, ACLs are created dynamically as FIP login frames are processed. The ACLs are installed on switch ports configured for the following port modes:

- ENode mode for server-facing ports
- FCF mode for a trusted port directly connected to an FCF

You must enable FIP snooping on an MXL Switch and configure the FIP snooping parameters. When you enable FIP snooping, all ports on the switch by default become ENode ports.

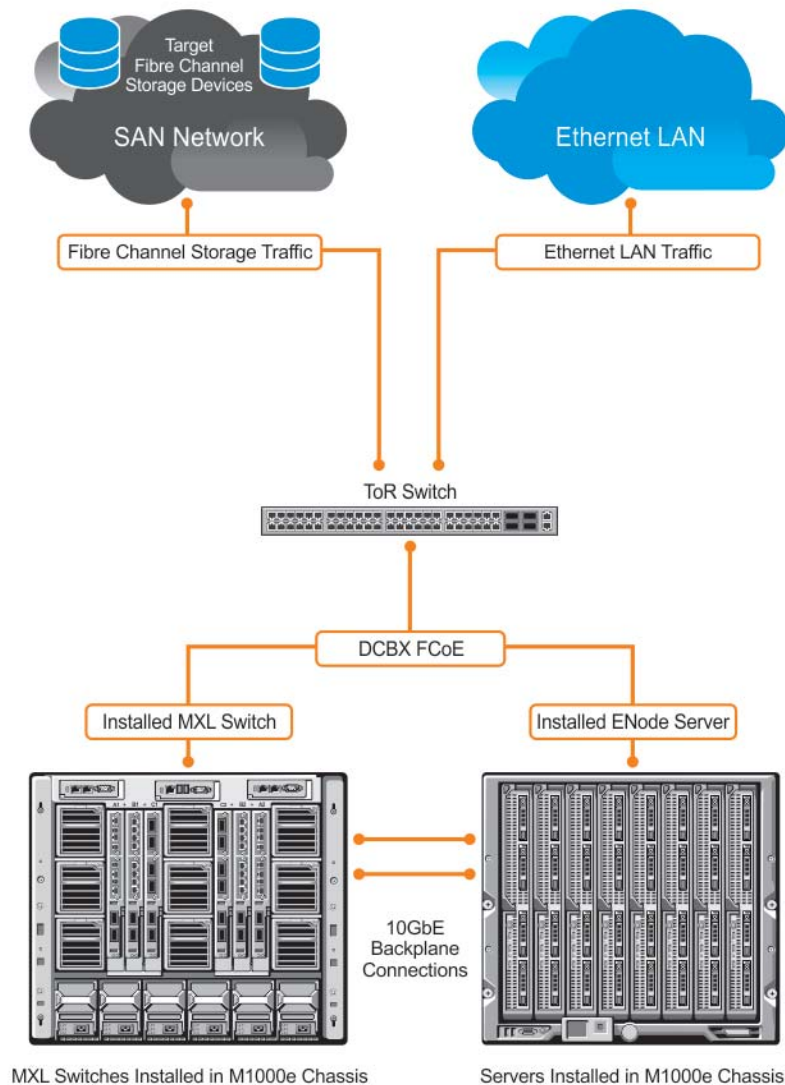
Dynamic ACL generation on an MXL Switch operating as a FIP snooping bridge functions as follows:

- Global ACLs are applied on server-facing ENode ports.

- Port-based ACLs are applied on ports directly connected to an FCF and on server-facing ENode ports.
- Port-based ACLs take precedence over global ACLs.
- FCoE-generated ACLs take precedence over user-configured ACLs. A user-configured ACL entry cannot deny FCoE and FIP snooping frames.

Figure 10-2 shows an MXL 10/40GbE Switch used as a FIP snooping bridge in a converged Ethernet network. The ToR switch operates as an FCF for FCoE traffic. Converged LAN and SAN traffic is transmitted between the ToR switch and an MXL switch. The MXL switch operates as a lossless FIP snooping bridge to transparently forward FCoE frames between the ENode servers and the FCF switch.

Figure 10-2. FIP Snooping on an MXL 10/40GbE Switch



The following sections describe how to configure the FIP snooping feature on a switch that functions as a FIP snooping bridge so that it can perform the following functions:

- Perform FIP snooping (allowing and parsing FIP frames) globally on all VLANs or on a per-VLAN basis.
- Set the FCoE MAC address prefix (FC-MAP) value used by an FCF to assign a MAC address to an FCoE end-device (server ENode or storage device) after a server successfully logs in.
- Set the FCF mode to provide additional port security on ports that are directly connected to an FCF.
- Check FIP snooping-enabled VLANs to ensure that they are operationally active.
- Process FIP VLAN discovery requests and responses, advertisements, solicitations, FLOGI/FDISC requests and responses, FLOGO requests and responses, keep-alive packets, and clear virtual-link messages.

FIP Snooping in a Switch Stack

FIP snooping supports switch stacking as follows:

- A switch stack configuration is synchronized with the standby stack unit.
- Dynamic population of the FCoE database (ENode, Session, and FCF tables) is synchronized with the standby stack unit. The FCoE database is maintained by snooping FIP keep-alive messages.
- In case of a failover, the new master switch starts the required timers for the FCoE database tables. Timers run only on the master stack unit.



Note: As a best practice, Dell Force10 recommends that you do not configure FIP Snooping on a stacked MXL Switch.

Configuring FIP Snooping

The configuration of FIP snooping consists of the following tasks:

1. Configure VLAN membership for all FCoE ports so that each port is a tagged member of the FCoE VLAN and an untagged member of the default VLAN.
2. Enable the FIP snooping feature on a switch to maintain FIP snooping information on the switch.
3. Enable FIP snooping on all VLANs (globally) or individual VLANs on a FIP snooping bridge.
4. Configure the FC-Map value applied globally by the switch on all VLANs or an individual VLAN.
5. Configure FCF mode for a FIP snooping bridge-to-FCF link.

For a sample FIP snooping configuration, refer to [Figure 10-11](#).

Enabling the FIP Snooping Feature

As soon as you enable the FIP snooping feature on a switch-bridge, existing VLAN-specific and FIP snooping configurations are applied. By default, all FCoE and FIP frames are dropped unless specifically permitted by existing FIP snooping-generated ACLs. You can reconfigure any of the FIP snooping settings.

If you disable FIP snooping, FIP and FCoE traffic are handled as normal Ethernet frames and no FIP snooping ACLs are generated. The VLAN-specific and FIP snooping configuration is disabled and stored until you re-enable FIP snooping and the configurations are re-applied.



Note: When you disable FIP snooping, the switch acts as pure Layer 2 switch that switches FCoE and FIP packets.

When you enable FIP snooping, the switch snoops FIP packets on VLANs enabled for FIP snooping and allows legitimate sessions. On VLANs disabled for FIP snooping, the switch drops FCoE and FIP packets.

Enabling FIP Snooping on VLANs

You can enable FIP snooping globally on a switch on all VLANs or on a specified VLAN. When you enable FIP snooping on VLANs:

- FIP frames are allowed to pass through the switch on the enabled VLANs and are processed to generate FIP snooping ACLs.
- FCoE traffic is allowed on VLANs only after a successful virtual-link initialization (fabric login FLOGI) between an ENode and an FCF. All other FCoE traffic is dropped.
- You must configure at least one interface for FCF (FIP snooping bridge-FCF) mode on a FIP snooping-enabled VLAN. You can configure multiple FCF trusted interfaces in a VLAN.
- A maximum of eight VLANs are supported for FIP snooping on the switch. When enabled globally, FIP snooping processes FIP packets in traffic only from the first eight incoming VLANs. When enabled on a per-VLAN basis, FIP snooping is supported on up to eight VLANs.

Configuring the FC-MAP Value

You can configure the FC-MAP value to be applied globally by the switch on all or individual FCoE VLANs to authorize FCoE traffic.

The configured FC-MAP value is used to check the FC-MAP value for the MAC address assigned to ENodes in incoming FCoE frames. If the FC-MAP value does not match, FCoE frames are dropped. A session between an ENode and an FCF is established by the switch-bridge only when the FC-MAP value on the FCF matches the FC-MAP value on the FIP snooping bridge.

Configuring a Port for a Bridge-to-FCF Link

If a port is directly connected to an FCF, configure the port mode as FCF. Initially, all FCoE traffic is blocked; only FIP frames are allowed to pass.

FCoE traffic is allowed on the port only after a successful FLOGI request/response and confirmed use of the configured FC-MAP value for the VLAN.

Impact on other Software Features

When you enable FIP snooping on a switch, other software features are impacted as follows:

- **MAC address learning:** MAC address learning is not performed on FIP and FCoE frames, which are denied by ACLs dynamically created by FIP snooping on server-facing ports in ENode mode.
- **MTU auto-configuration:** MTU size is set to mini-jumbo (2500 bytes) when a port is in Switchport mode, the FIP snooping feature is enabled on the switch, and FIP snooping is enabled on all or individual VLANs.
- **Link aggregation group (LAG):** FIP snooping is supported on port channels on ports on which PFC mode is on (PFC is operationally up).
- **STP:** If you enable an STP protocol (STP, RSTP, PVSTP, or MSTP) on the switch and ports enter a blocking state, when the state change occurs, the corresponding port-based ACLs are deleted. If a port is enabled for FIP snooping in ENode or FCF mode, the ENode/FCF MAC-based ACLs are deleted.

FIP Snooping Prerequisites

Before you configure FIP snooping on an MXL switch, ensure that the following conditions are met:

- A FIP snooping bridge requires DCBX and PFC to be enabled on the switch for lossless Ethernet connections (refer to [Data Center Bridging \(DCB\)](#)). Dell recommends that you also enable ETS; ETS is recommended but not required.
If you enable DCBX and PFC mode is on (PFC is operationally up) in a port configuration, FIP snooping is operational on the port. If the PFC parameters in a DCBX exchange with a peer are not synchronized, FIP and FCoE frames are dropped on the port after you enable the FIP snooping feature.
- **VLAN membership:**
 - You must create the VLANs on the switch which handles FCoE traffic (`interface vlan` command).
 - You must configure each FIP snooping port to operate in Hybrid mode so that it accepts both tagged and untagged VLAN frames (`portmode hybrid` command).
 - You must configure tagged VLAN membership on each FIP snooping port that sends and receives FCoE traffic and has links with an FCF, ENode server, or another FIP snooping bridge (`tagged port-type slot/port` command).
The default VLAN membership of the port should continue to operate with untagged frames. FIP snooping is not supported on a port that is configured for non-default untagged VLAN membership.

FIP Snooping Restrictions

The following restrictions apply when you configure FIP snooping on an MXL switch:

- The maximum number of FCoE VLANs supported on the switch is eight.
- The maximum number of FIP snooping sessions (including NPIV sessions) supported per ENode server is 16.

In a full FCoE N_port ID virtualization (NPIV) configuration, 16 sessions (one FLOGI + fifteen NPIV sessions) are supported per ENode. In an FCoE NPV configuration, only one session is supported per ENode.

- The maximum number of FCFs supported per FIP snooping-enabled VLAN is four.
- Links to other FIP snooping bridges on a FIP snooping-enabled port (bridge-to-bridge links) are not supported on the MXL Switch.

Configuration Procedure

You can enable FIP snooping globally on all FCoE VLANs on a switch or on an individual FCoE VLAN.

To enable FIP snooping on the switch and configure FIP snooping parameters on ports, follow these steps:

Step	Task	Command	Command Mode
1	Enable the FIP snooping feature on a switch. Default: FIP snooping is disabled.	<code>feature fip-snooping</code>	CONFIGURATION
2	Enable FIP snooping on all VLANs or on a specified VLAN. Default: FIP snooping is disabled on all VLANs.	<code>fip-snooping enable</code>	CONFIGURATION Or VLAN INTERFACE
3	Configure the FC-MAP value used by FIP snooping on all VLANs. Default: 0x0EFC00. Valid values are from 0EFC00 to 0EFCFF.	<code>fip-snooping fc-map fc-map-value</code>	CONFIGURATION Or VLAN INTERFACE
4	Enter interface configuration mode to configure the port for FIP snooping links. Note: By default, a port is configured for bridge-to-ENode links.	<code>interface port-type slot/port</code>	CONFIGURATION
5	Configure the port for bridge-to-FCF links.	<code>fip-snooping port-mode fcf</code>	INTERFACE CONFIGURATION



Note: To disable the FIP snooping feature or FIP snooping on VLANs, use the no version of a command; for example, `no feature fip-snooping` or `no fip-snooping enable`.

Displaying FIP Snooping Information

Use the show commands in [Table 10-1](#) to display information on FIP snooping.

Table 10-1. Displaying FIP Snooping Information

Command	Output
show fip-snooping sessions [interface vlan <i>vlan-id</i>] (Figure 10-3)	Displays information on FIP-snooped sessions on all VLANs or a specified VLAN, including the ENode interface and MAC address, the FCF interface and MAC address, VLAN ID, FCoE MAC address and FCoE session ID number (FC-ID), worldwide node name (WWNN) and the worldwide port name (WWPN). Information on NPIV sessions is also displayed.
show fip-snooping config (Figure 10-4)	Displays the FIP snooping status and configured FC-MAP values.
show fip-snooping enode [<i>enode-mac-address</i>] (Figure 10-5)	Displays information on the ENodes in FIP-snooped sessions, including the ENode interface and MAC address, FCF MAC address, VLAN ID and FC-ID.
show fip-snooping fcf [<i>fcf-mac-address</i>] (Figure 10-6)	Displays information on the FCFs in FIP-snooped sessions, including the FCF interface and MAC address, FCF interface, VLAN ID, FC-MAP value, FKA advertisement period, and number of ENodes connected.
clear fip-snooping database interface vlan <i>vlan-id</i> { <i>fcoe-mac-address</i> <i>enode-mac-address</i> <i>fcf-mac-address</i> }	Clears FIP snooping information on a VLAN for a specified FCoE MAC address, ENode MAC address, or FCF MAC address, and removes the corresponding ACLs generated by FIP snooping.
show fip-snooping statistics [interface vlan <i>vlan-id</i> interface <i>port-type port/slot</i> interface port-channel <i>port-channel-number</i>] (Figure 10-7 and Figure 10-8)	Displays statistics on the FIP packets snooped on all interfaces, including VLANs, physical ports, and port channels.
clear fip-snooping statistics [interface vlan <i>vlan-id</i> interface <i>port-type port/slot</i> interface port-channel <i>port-channel-number</i>]	Clears the statistics on the FIP packets snooped on all VLANs, a specified VLAN, or a specified port interface.
show fip-snooping system (Figure 10-9)	Display information on the status of FIP snooping on the switch (enabled or disabled), including the number of FCoE VLANs, FCFs, ENodes, and currently active sessions.
show fip-snooping vlan (Figure 10-10)	Display information on the FCoE VLANs on which FIP snooping is enabled.

Figure 10-3. show fip-snooping sessions Command Example

```

FTOS#show fip-snooping sessions
ENode MAC           ENode Intf      FCF MAC          FCF Intf        VLAN
aa:bb:cc:00:00:00  Te 0/42        aa:bb:cd:00:00:00 Te 0/43         100
aa:bb:cc:00:00:00  Te 0/42        aa:bb:cd:00:00:00 Te 0/43         100
aa:bb:cc:00:00:00  Te 0/42        aa:bb:cd:00:00:00 Te 0/43         100
aa:bb:cc:00:00:00  Te 0/42        aa:bb:cd:00:00:00 Te 0/43         100
aa:bb:cc:00:00:00  Te 0/42        aa:bb:cd:00:00:00 Te 0/43         100

FCoE MAC           FC-ID          Port WWPN          Port WWNN
0e:fc:00:01:00:01  01:00:01      31:00:0e:fc:00:00:00:00  21:00:0e:fc:00:00:00:00
0e:fc:00:01:00:02  01:00:02      41:00:0e:fc:00:00:00:00  21:00:0e:fc:00:00:00:00
0e:fc:00:01:00:03  01:00:03      41:00:0e:fc:00:00:00:01  21:00:0e:fc:00:00:00:00
0e:fc:00:01:00:04  01:00:04      41:00:0e:fc:00:00:00:02  21:00:0e:fc:00:00:00:00
0e:fc:00:01:00:05  01:00:05      41:00:0e:fc:00:00:00:03  21:00:0e:fc:00:00:00:00

```

Table 10-2. show fip-snooping sessions Command Description

Field	Description
ENode MAC	MAC address of the ENode.
ENode Interface	Slot/ port number of the interface connected to the ENode.
FCF MAC	MAC address of the FCF.
FCF Interface	Slot/ port number of the interface to which the FCF is connected.
VLAN	VLAN ID number used by the session.
FCoE MAC	MAC address of the FCoE session assigned by the FCF.
FC-ID	Fibre Channel ID assigned by the FCF.
Port WWPN	Worldwide port name of the CNA port.
Port WWNN	Worldwide node name of the CNA port.

Figure 10-4. show fip-snooping config Command Example

```

FTOS# show fip-snooping config
FIP Snooping Feature enabled Status: Enabled
FIP Snooping Global enabled Status: Enabled
Global FC-MAP Value: 0X0EFC00

FIP Snooping enabled VLANs
VLAN   Enabled   FC-MAP
----   -
100    TRUE      0X0EFC00

```

Figure 10-5. show fip-snooping enode Command Example

```

FTOS# show fip-snooping enode
Enode MAC           Enode Interface     FCF MAC             VLAN                FC-ID
-----
d4:ae:52:1b:e3:cd   Te 0/11             54:7f:ee:37:34:40  100                 62:00:11
    
```

Table 10-3. show fip-snooping enode Command Description

Field	Description
ENode MAC	MAC address of the ENode.
ENode Interface	Slot/ port number of the interface connected to the ENode.
FCF MAC	MAC address of the FCF.
VLAN	VLAN ID number used by the session.
FC-ID	Fibre Channel session ID assigned by the FCF.

Figure 10-6. show fip-snooping fcf Command Example

```

FTOS# show fip-snooping fcf
FCF MAC           FCF Interface     VLAN                FC-MAP              FKA_ADV_PERIOD      No. of Enodes
-----
54:7f:ee:37:34:40  Po 22             100                 0e:fc:00            4000                 2
    
```

Table 10-4. show fip-snooping fcf Command Description

Field	Description
FCF MAC	MAC address of the FCF.
FCF Interface	Slot/port number of the interface to which the FCF is connected.
VLAN	VLAN ID number used by the session.
FC-MAP	FC-Map value advertised by the FCF.
ENode Interface	Slot/ number of the interface connected to the ENode.
FKA_ADV_PERIOD	Period of time (in milliseconds) during which FIP keep-alive advertisements are transmitted.
No of ENodes	Number of ENodes connected to the FCF.
FC-ID	Fibre Channel session ID assigned by the FCF.

Figure 10-7. show fip-snooping statistics (VLAN and port) Command Example

```

FTOS# show fip-snooping statistics interface vlan 100
Number of Vlan Requests                :0
Number of Vlan Notifications           :0
Number of Multicast Discovery Solicits  :2
Number of Unicast Discovery Solicits    :0
Number of FLOGI                        :2
Number of FDISC                         :16
Number of FLOGO                         :0
Number of Enode Keep Alive              :9021
Number of VN Port Keep Alive            :3349
Number of Multicast Discovery Advertisement :4437
Number of Unicast Discovery Advertisement :2
Number of FLOGI Accepts                 :2
Number of FLOGI Rejects                 :0
Number of FDISC Accepts                 :16
Number of FDISC Rejects                 :0
Number of FLOGO Accepts                 :0
Number of FLOGO Rejects                 :0
Number of CVL                           :0
Number of FCF Discovery Timeouts         :0
Number of VN Port Session Timeouts      :0
Number of Session failures due to Hardware Config :0
FTOS(conf)#

FTOS# show fip-snooping statistics int tengigabitethernet 0/11
Number of Vlan Requests                :1
Number of Vlan Notifications           :0
Number of Multicast Discovery Solicits  :1
Number of Unicast Discovery Solicits    :0
Number of FLOGI                        :1
Number of FDISC                         :16
Number of FLOGO                         :0
Number of Enode Keep Alive              :4416
Number of VN Port Keep Alive            :3136
Number of Multicast Discovery Advertisement :0
Number of Unicast Discovery Advertisement :0
Number of FLOGI Accepts                 :0
Number of FLOGI Rejects                 :0
Number of FDISC Accepts                 :0
Number of FDISC Rejects                 :0
Number of FLOGO Accepts                 :0
Number of FLOGO Rejects                 :0
Number of CVL                           :0
Number of FCF Discovery Timeouts         :0
Number of VN Port Session Timeouts      :0
Number of Session failures due to Hardware Config :0

```


Figure 10-8. show fip-snooping statistics (port channel) Command Example

```
FTOS# show fip-snooping statistics interface port-channel 22
Number of Vlan Requests :0
Number of Vlan Notifications :2
Number of Multicast Discovery Solicits :0
Number of Unicast Discovery Solicits :0
Number of FLOGI :0
Number of FDISC :0
Number of FLOGO :0
Number of Enode Keep Alive :0
Number of VN Port Keep Alive :0
Number of Multicast Discovery Advertisement :4451
Number of Unicast Discovery Advertisement :2
Number of FLOGI Accepts :2
Number of FLOGI Rejects :0
Number of FDISC Accepts :16
Number of FDISC Rejects :0
Number of FLOGO Accepts :0
Number of FLOGO Rejects :0
Number of CVL :0
Number of FCF Discovery Timeouts :0
Number of VN Port Session Timeouts :0
Number of Session failures due to Hardware Config :0
```

Table 10-5. show fip-snooping statistics Command Descriptions

Field	Description
Number of Vlan Requests	Number of FIP-snooped VLAN request frames received on the interface.
Number of VLAN Notifications	Number of FIP-snooped VLAN notification frames received on the interface.
Number of Multicast Discovery Solicits	Number of FIP-snooped multicast discovery solicit frames received on the interface.
Number of Unicast Discovery Solicits	Number of FIP-snooped unicast discovery solicit frames received on the interface.
Number of FLOGI	Number of FIP-snooped FLOGI request frames received on the interface.
Number of FDISC	Number of FIP-snooped FDISC request frames received on the interface.
Number of FLOGO	Number of FIP-snooped FLOGO frames received on the interface.
Number of ENode Keep Alives	Number of FIP-snooped ENode keep-alive frames received on the interface.
Number of VN Port Keep Alives	Number of FIP-snooped VN port keep-alive frames received on the interface.
Number of Multicast Discovery Advertisements	Number of FIP-snooped multicast discovery advertisements received on the interface.
Number of Unicast Discovery Advertisements	Number of FIP-snooped unicast discovery advertisements received on the interface.
Number of FLOGI Accepts	Number of FIP FLOGI accept frames received on the interface.
Number of FLOGI Rejects	Number of FIP FLOGI reject frames received on the interface.
Number of FDISC Accepts	Number of FIP FDISC accept frames received on the interface.
Number of FDISC Rejects	Number of FIP FDISC reject frames received on the interface.
Number of FLOGO Accepts	Number of FIP FLOGO accept frames received on the interface.
Number of FLOGO Rejects	Number of FIP FLOGO reject frames received on the interface.
Number of CVLs	Number of FIP clear virtual link frames received on the interface.
Number of FCF Discovery Timeouts	Number of FCF discovery timeouts that occurred on the interface.
Number of VN Port Session Timeouts	Number of VN port session timeouts that occurred on the interface.
Number of Session failures due to Hardware Config	Number of session failures due to hardware configuration that occurred on the interface.

Figure 10-9. show fip-snooping system Command Example

```
FTOS# show fip-snooping system
Global Mode           : Enabled
FCOE VLAN List (Operational) : 1, 100
FCFs                  : 1
Enodes                 : 2
Sessions               : 17
```



Note: NPIV sessions are included in the number of FIP-snooped sessions displayed.

Figure 10-10. show fip-snooping vlan Command Example

```
FTOS# show fip-snooping vlan
* = Default VLAN

VLAN    FC-MAP      FCFs    Enodes  Sessions
-----  -
*1      -            -       -       -
100     0X0EFC00    1       2       17
```




Note: NPIV sessions are included in the number of FIP-snooped sessions displayed.

FIP Snooping Configuration Example

Figure 10-11 shows an MXL Switch used as a FIP snooping bridge for FCoE traffic between an ENode (server blade) and an FCF (ToR switch). The ToR switch operates as an FCF and FCoE gateway.

Figure 10-11. Configuration Example: FIP Snooping on an MXL 10/40GbE Switch



In Figure 10-11, DCBX and PFC are enabled on the MXL Switch (FIP snooping bridge) and on the FCF ToR switch. On the FIP snooping bridge, DCBX is configured as follows:

- A server-facing port is configured for DCBX in an auto-downstream role.
- An FCF-facing port is configured for DCBX in an auto-upstream or configuration-source role.

The DCBX configuration on the FCF-facing port is detected by the server-facing port and the DCB PFC configuration on both ports is synchronized. For more information about how to configure DCBX and PFC on a port, refer to [Data Center Bridging \(DCB\)](#).

Figure 10-12 shows how to configure FIP snooping on FCoE VLAN 10, an FCF-facing port (0/50), and an ENode server-facing port (0/1), and to configure the FIP snooping ports as tagged members of the FCoE VLAN enabled for FIP snooping.

Figure 10-12. FIP Snooping Configuration Example

Enable the FIP snooping feature on the switch (FIP snooping bridge):

```
FTOS(conf)# feature fip-snooping
```

Enable FIP snooping on FCoE VLAN 10:

```
FTOS(conf)# interface vlan 10
FTOS(conf-if-vl-10)# fip-snooping enable
```

Enable an FC-MAP value on VLAN 10:

```
FTOS(conf-if-vl-10)# fip-snooping fc-map 0x0EFC01
```

Note: Configuring an FC-MAP value is only required if you do not use the default FC-MAP value (0x0EFC00).

Configure the ENode server-facing port:

```
FTOS(conf)# interface tengigabitethernet 0/1
FTOS(conf-if-te-0/1)# portmode hybrid
FTOS(conf-if-te-0/1)# switchport
```

Note: A port is enabled by default for bridge-ENode links.

Configure the FCF-facing port:

```
FTOS(conf)# interface tengigabitethernet 0/50
FTOS(conf-if-te-0/50)# portmode hybrid
FTOS(conf-if-te-0/50)# switchport
FTOS(conf-if-te-0/50)# fip-snooping port-mode fcf
```

Configure FIP snooping ports as tagged members of FCoE VLAN:

```
FTOS(conf)# interface vlan 10
FTOS(conf-if-vl-10)# tagged tengigabitethernet 0/1
FTOS(conf-if-vl-10)# tagged tengigabitethernet 0/50
FTOS(conf-if-te-0/1)# no shut
FTOS(conf-if-te-0/50)# no shut
FTOS(conf-if-vl-10)# no shut
```

After FIP packets are exchanged between the ENode and the switch, a FIP snooping session is established. ACLS are dynamically generated for FIP snooping on the FIP snooping bridge/switch.

GARP VLAN Registration Protocol (GVRP)

This chapter contains the following sections:

- [Configuring GVRP](#)
- [Enabling GVRP Globally](#)
- [Enabling GVRP on a Layer 2 Interface](#)
- [Configuring GVRP Registration](#)
- [Configuring a GARP Timer](#)

Overview

Typical virtual local area network (VLAN) implementation involves manually configuring each Layer 2 switch that participates in a given VLAN. GARP VLAN registration protocol (GVRP), defined by the IEEE 802.1q specification, is a Layer 2 network protocol that provides for automatic VLAN configuration of switches. GVRP-compliant switches use the generic attribute registration protocol (GARP) to register and de-register attribute values, such as VLAN IDs, with each other.

GVRP exchanges network VLAN information to allow switches to dynamically forward frames for one or more VLANs. Consequently, GVRP spreads this information and configures the needed VLAN(s) on any additional switches in the network. Data propagates using the exchange of GVRP protocol data units (PDUs).

The purpose of GVRP is to simplify (but not eliminate) static configuration. The idea is to configure switches at the edge and have the information dynamically propagate into the core. As such, the edge ports must still be statically configured with VLAN membership information, and they do not run GVRP. It is this information that is propagated to create dynamic VLAN membership in the core of the network.

Important Points to Remember

- GVRP propagates VLAN membership throughout a network. GVRP allows end stations and switches to issue and revoke declarations relating to VLAN membership.
- VLAN registration is made in the context of the port that receives the GARP PDU and is propagated to the other active ports.
- GVRP is disabled by default; you must enable GVRP for the switch and then for individual ports.

- Dynamic VLANs are aged out after the LeaveAll timer expires three times without receipt of a Join message. Use the `show gvrp statistics {interface interface | summary}` command to display status.
- On the MXL Switch, per-VLAN spanning tree+ (PVST+) and GVRP cannot be enabled at the same time (Figure 11-1). If spanning tree and GVRP are both required, implement either rapid spanning tree protocol (RSTP), spanning tree protocol (STP), or multiple spanning tree protocol (MSTP). The MXL 10/40GbE Switch IO Module system does support enabling GVRP and MSTP at the same time.

Figure 11-1. GVRP Compatibility Error Message

```
FTOS(conf)#protocol spanning-tree pvst
FTOS(conf-pvst)#no disable
% Error: GVRP running. Cannot enable PVST.


FTOS(conf)#protocol gvrp
FTOS(conf-gvrp)#no disable
% Error: PVST running. Cannot enable GVRP.
```

Configuring GVRP

Globally, enable GVRP on each switch to facilitate GVRP communications. Then, GVRP configuration is per interface on a switch-by-switch basis. Enable GVRP on each port that connects to a switch where you want GVRP information exchanged. In Figure 11-2, that kind of port is referred to as a “VLAN trunk port,” but it is not necessary to specifically identify to FTOS that the port is a trunk port.

Figure 11-2. GVRP Configuration Overview

GVRP is configured globally and on all VLAN trunk ports for the edge and core switches.



NOTES:

- VLAN 1 mode is always fixed and cannot be configured
- All VLAN trunk ports must be configured for GVRP
- All VLAN trunk ports must be configured as 802.1Q

Basic GVRP configuration is a two-step process:

1. [Enabling GVRP Globally](#)
2. [Enabling GVRP on a Layer 2 Interface.](#)

Related Configuration Tasks

- [Configuring GVRP Registration](#)
- [Configuring a GARP Timer](#)

Enabling GVRP Globally

Enable GVRP for the entire switch using the `gvrp enable` command in CONFIGURATION mode (Figure 11-3). Use the `show gvrp brief` command to inspect the global configuration.

Figure 11-3. Enabling GVRP Globally

```
FTOS(conf)#protocol gvrp
FTOS(conf-gvrp)#no disable
FTOS(conf-gvrp)#show config
!
protocol gvrp
no disable
FTOS(conf-gvrp)#
```

Enabling GVRP on a Layer 2 Interface

Enable GVRP on a Layer 2 interface using the `gvrp enable` command in INTERFACE mode (Figure 11-4). Use the `show config` command from INTERFACE mode to inspect the interface configuration (Figure 11-4), or use the `show gvrp interface` command in EXEC or EXEC Privilege mode.

Figure 11-4. Enabling GVRP on a Layer 2 Interface

```
FTOS(conf-if-te-1/21)#switchport
FTOS(conf-if-te-1/21)#gvrp enable
FTOSFTOS(conf-if-te-1/21)#no shutdown
FTOS(conf-if-te-1/21)#show config
!
interface TenGigabitEthernet 1/21
no ip address
switchport
gvrp enable
no shutdown
```

Configuring GVRP Registration

There are three types of GVRP registration:

- 1) Normal Registration
- 2) Fixed Registration
- 3) Forbidden Registration.

- **Normal Registration:** Allows dynamic creation, registration, and de-registration of VLANs (if you enabled dynamic VLAN creation). By default, the registration mode is set to normal when you enable GVRP on a port. This default mode enables the port to dynamically register and de-register VLANs, and to propagate both dynamic and static VLAN information.
- **Fixed Registration Mode:** Configuring a port in fixed registration mode allows for manual creation and registration of VLANs, prevents VLAN de-registration, and registers all VLANs known on other ports on the port. For example, if an interface is statically configured using the command line interface (CLI) to belong to a VLAN, it should not be un-configured when it receives a Leave PDU. So, the registration mode on that interface is FIXED.
- **Forbidden Mode:** Disables the port to dynamically register VLANs, and to propagate VLAN information except information about VLAN 1. A port with forbidden registration type thus allows only VLAN 1 to pass through even though the PDU carries information for more VLANs. So, set the interface to registration mode of FORBIDDEN if you do not want the interface to advertise or learn about particular VLANS.

Based on the configuration in the example shown in [Figure 11-5](#), the interface 1/21 is not removed from VLAN 34 or VLAN 35 despite receiving a GVRP Leave message. Additionally, the interface is not dynamically added to VLAN 45 or VLAN 46, even if a GVRP Join message is received.

Figure 11-5. Configuring GVRP Registration

```

FTOS(conf-if-te-1/21)#gvrp registration fixed 34,35
FTOS(conf-if-te-1/21)#gvrp registration forbidden 45,46
FTOS(conf-if-te-1/21)#show conf
!
interface TenGigabitEthernet 1/21
 no ip address
 switchport
 gvrp enable
 gvrp registration fixed 34-35
 gvrp registration forbidden 45-46
 no shutdown
FTOS(conf-if-te-1/21)#

```

Configuring a GARP Timer

GARP timers must be set to the same values on all devices that are exchanging information using GVRP:

- **Join:** A GARP device reliably transmits Join messages to other devices by sending each Join message two times. Use this parameter to define the interval between the two sending operations of each Join message. The FTOS default is 200ms.
- **Leave:** When a GARP device expects to de-register a piece of attribute information, it sends out a Leave message and start this timer. If a Join message does not arrive before the timer expires, the information is de-registered. The Leave timer must be greater than or equal to 3x the Join timer. The FTOS default is 600ms.
- **LeaveAll:** After startup, a GARP device globally starts a LeaveAll timer. After expiration of this interval, it sends out a LeaveAll message so that other GARP devices can re-register all relevant attribute information. The device then restarts the LeaveAll timer to begin a new cycle. The LeaveAll timer must be greater than or equal to 5x of the Leave timer. The FTOS default is 10000ms.

Figure 11-6 shows GVRP registration.

Figure 11-6. Configuring GVRP Registration

```
FTOS(conf)#garp timer leav 1000
FTOS(conf)#garp timers leave-all 5000
FTOS(conf)#garp timer join 300
```

Verification:

```
FTOS(conf)#do show garp timer
GARP Timers      Value (milliseconds)
-----
Join Timer       300
Leave Timer       1000
LeaveAll Timer    5000
FTOS(conf)#
```

FTOS displays [Message 1](#) if an attempt is made to configure an invalid GARP timer.

Message 1 GARP Timer Error

```
FTOS(conf)#garp timers join 300
% Error: Leave timer should be >= 3*Join timer.
```

Internet Group Management Protocol (IGMP)

Multicast is based on identifying many hosts by a single destination IP address. Hosts represented by the same IP address are a *multicast group*. The internet group management protocol (IGMP) is a Layer 3 multicast protocol that hosts use to join or leave a multicast group. Multicast routing protocols (such as protocol-independent multicast [PIM]) use the information in IGMP messages to discover which groups are active and to populate the multicast routing table.

This chapter contains the following sections:

- [IGMP Snooping](#)
- [Fast Convergence after MSTP Topology Changes](#)
- [Designating a Multicast Router Interface](#)

Overview


IGMP has three versions. Version 3 obsoletes and is backwards-compatible with version 2; version 2 obsoletes version 1.

IGMP Version 2

IGMP version 2 improves upon version 1 by specifying IGMP Leave messages, which allows hosts to notify routers that they no longer care about traffic for a particular group. Leave messages reduce the amount of time that the router takes to stop forwarding traffic for a group to a subnet (leave latency) after the last host leaves the group. In version 1 hosts quietly leave groups, and the router waits for a query response timer several times the value of the query interval to expire before it stops forwarding traffic.

To receive multicast traffic from a particular source, a host must join the multicast group to which the source is sending traffic. A host that is a member of a group is called a “receiver.” A host may join many groups, and may join or leave any group at any time. A host joins and leaves a multicast group by sending an IGMP message to its IGMP querier. The querier is the router that surveys a subnet for multicast receivers and processes survey responses to populate the multicast routing table.

IGMP messages are encapsulated in IP packets ([Figure 12-1](#)).

Figure 12-1. IGMP Version 2 Packet Format

Joining a Multicast Group

There are two ways that a host may join a multicast group: it may respond to a general query from its querier, or it may send an unsolicited report to its querier.

- Responding to an IGMP Query
 - One router on a subnet is elected as the querier. The querier periodically multicasts (to all-multicast-systems address 224.0.0.1) a general query to all hosts on the subnet.
 - A host that wants to join a multicast group responds with an IGMP membership report that contains the multicast address of the group it wants to join (the packet is addressed to the same group). If multiple hosts want to join the same multicast group, only the report from the first host to respond reaches the querier, and the remaining hosts suppress their responses (for how the delay timer mechanism works, refer to [IGMP Snooping](#)).
 - The querier receives the report for a group and adds the group to the list of multicast groups associated with its outgoing port to the subnet. Multicast traffic for the group is then forwarded to that subnet.
- Sending an Unsolicited IGMP Report
 - A host does not have to wait for a general query to join a group. It may send an unsolicited IGMP membership report, also called an IGMP Join message, to the querier.

Leaving a Multicast Group

- A host sends a membership report of type 0x17 (IGMP Leave message) to the all routers multicast address 224.0.0.2 when it no longer cares about multicast traffic for a particular group.
- The querier sends a group-specific query to determine whether there are any remaining hosts in the group. There must be at least one receiver in a group on a subnet for a router to forward multicast traffic for that group to the subnet.
- Any remaining hosts respond to the query according to the delay timer mechanism (refer to [IGMP Snooping](#)). If no hosts respond (because there are none remaining in the group), the querier waits a specified period and sends another query. If it still receives no response, the querier removes the group from the list associated with forwarding port and stops forwarding traffic for that group to the subnet.

IGMP Version 3

Conceptually, IGMP version 3 behaves the same as version 2. However, there are differences:

- Version 3 adds the ability to filter by multicast source, which helps the multicast routing protocols avoid forwarding traffic to subnets where there are no interested receivers.
- To enable filtering, routers must keep track of more state information, that is, the list of sources that must be filtered. An additional query type, the group-and-source-specific query, keeps track of state changes, while the group-specific and general queries still refresh existing state.
- Reporting is more efficient and robust. Hosts do not suppress query responses (non-suppression helps track state and enables the immediate-leave and IGMP snooping features), state-change reports are retransmitted to insure delivery, and a single membership report bundles multiple statements from a single host, rather than sending an individual packet for each statement.

To accommodate these protocol enhancements, the IGMP version 3 packet structure is different from version 2. Queries (Figure 12-2) are still sent to the all-systems address 224.0.0.1, but reports (Figure 12-3) are sent to all the IGMP version 3-capable multicast routers address 244.0.0.22.

Figure 12-2. IGMP version 3 Membership Query Packet Format

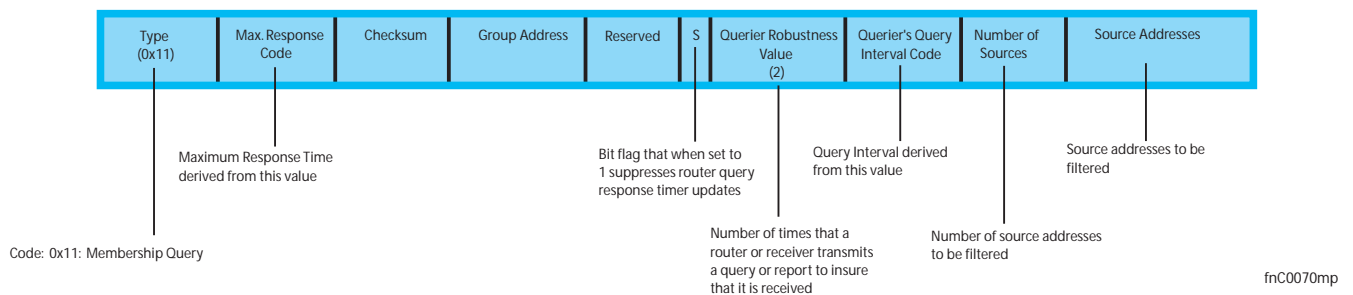
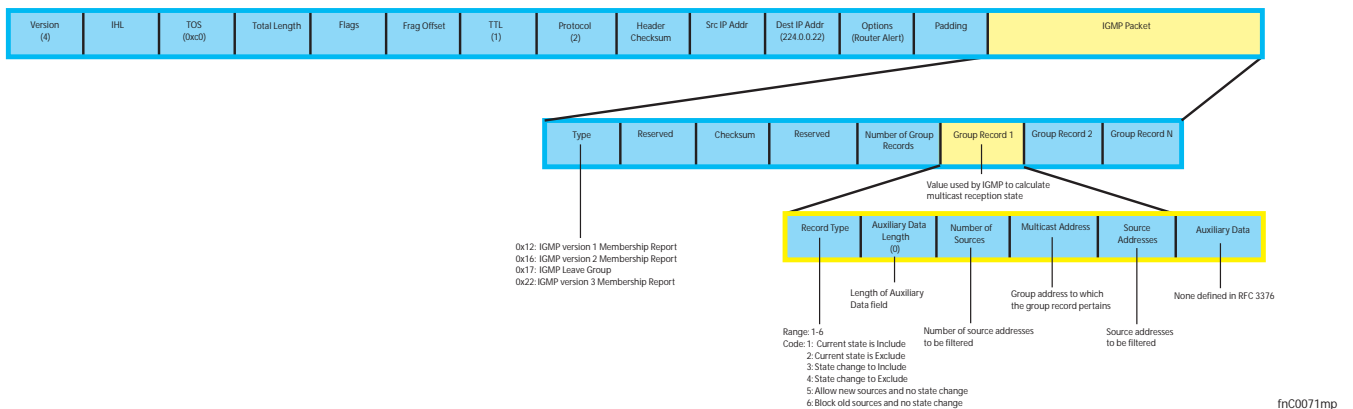


Figure 12-3. IGMP version 3 Membership Report Packet Format

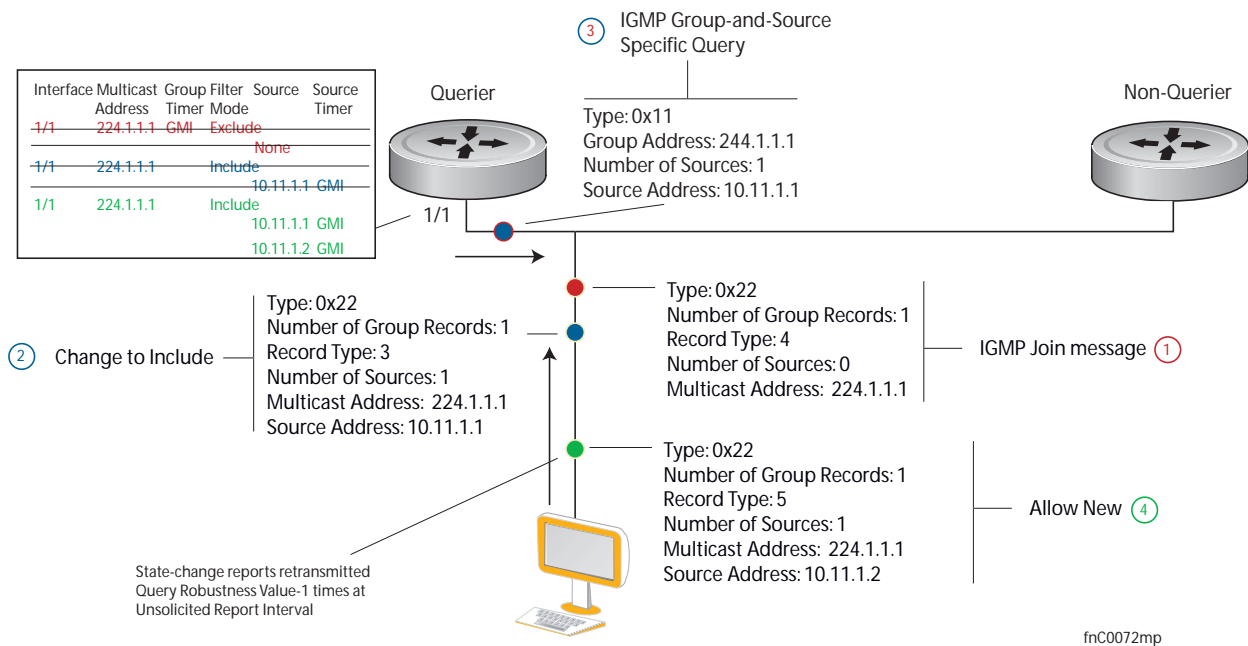


Joining and Filtering Groups and Sources

Figure 12-4 shows how multicast routers maintain the group and source information from unsolicited reports.

1. The first unsolicited report from the host indicates that it wants to receive traffic for group 224.1.1.1.
2. The host's second report indicates that it is only interested in traffic from group 224.1.1.1, source 10.11.1.1. Include messages prevent traffic from all other sources in the group from reaching the subnet, so before recording this request, the querier sends a group-and-source query to verify that there are no hosts interested in any other sources. The multicast router must satisfy all hosts if they have conflicting requests. For example, if another host on the subnet is interested in traffic from 10.11.1.3, the router cannot record the include request. There are no other interested hosts, so the request is recorded. At this point, the multicast routing protocol prunes the tree to all but the specified sources.
3. The host's third message indicates that it is only interested in traffic from sources 10.11.1.1 and 10.11.1.2. Because this request again prevents all other sources from reaching the subnet, the router sends another group-and-source query so that it can satisfy all other hosts. There are no other interested hosts, so the request is recorded.

Figure 12-4. IGMP Membership Reports: Joining and Filtering
Membership Reports: Joining and Filtering

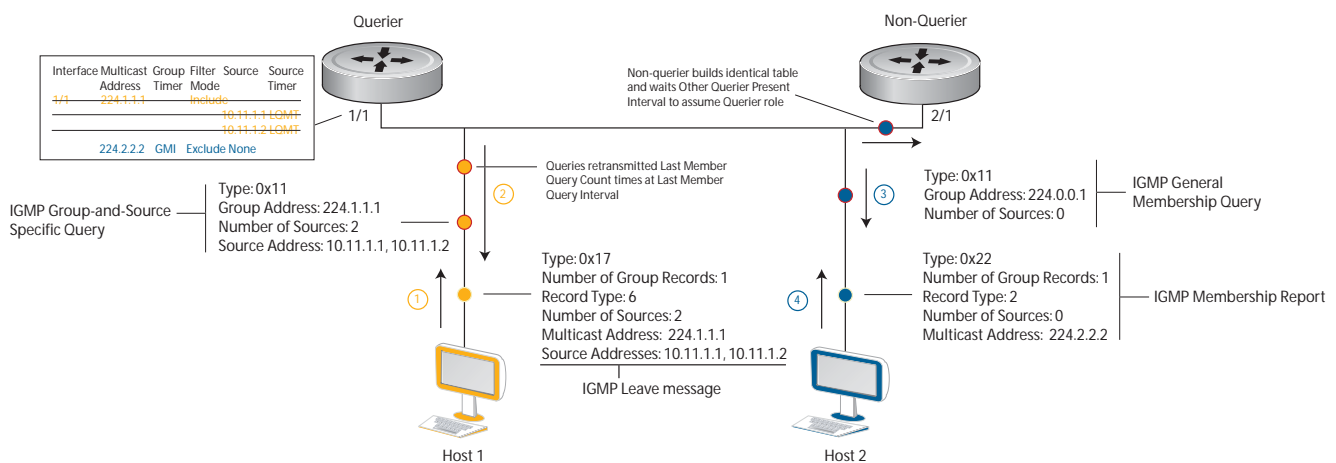


Leaving and Staying in Groups

Figure 12-5 shows how multicast routers track and refresh state changes in response to group-and-specific and general queries.

1. Host 1 sends a message indicating it is leaving group 224.1.1.1 and that the included filter for 10.11.1.1 and 10.11.1.2 are no longer necessary.
2. The querier, before making any state changes, sends a group-and-source query to see if any other host is interested in these two sources; queries for state-changes are retransmitted multiple times. If any are interested, they respond with their current state information and the querier refreshes the relevant state information.
3. Separately in Figure 12-5, the querier sends a general query to 224.0.0.1.
4. Host 2 responds to the periodic general query so the querier refreshes the state information for that group.

Figure 12-5. IGMP Membership Queries: Leaving and Staying in Groups
Membership Queries: Leaving and Staying



IGMP Snooping

Multicast packets are addressed with multicast MAC addresses, which represent a group of devices, rather than one unique device. Switches forward multicast frames out of all ports in a VLAN by default, even though there may be only some interested hosts, which is a waste of bandwidth. IGMP snooping enables switches to use information in IGMP packets to generate a forwarding table that associates ports with multicast groups so that when they receive multicast frames, they can forward them only to interested receivers.

IGMP Snooping Implementation Information

- IGMP snooping on the Dell Force 10 operating system (FTOS) uses IP multicast addresses not MAC addresses.
- IGMP snooping is not supported on stacked VLANs.
- IGMP snooping is supported on all MXL 10/40GbE stack members.
- IGMP snooping reacts to STP and MSTP topology changes by sending a general query on the interface that transitions to the forwarding state.

Configuring IGMP Snooping

Configuring IGMP snooping is a one-step process. That is, you enable it on a switch using the `ip igmp snooping enable` command from CONFIGURATION mode.

To view the configuration, use the `show running-config` command from CONFIGURATION mode (Figure 12-6). To disable snooping for a VLAN, use the `no ip igmp snooping` command from INTERFACE VLAN mode.

Figure 12-6. Enabling IGMP Snooping

```
FTOS(conf)#ip igmp snooping enable
FTOS(conf)#do show running-config igmp
ip igmp snooping enable
FTOS(conf)#
```

Related Configuration Tasks

- [Enabling IGMP Immediate-leave](#)
- [Disabling Multicast Flooding](#)
- [Specifying a Port as Connected to a Multicast Router](#)
- [Configuring the Switch as Querier](#)

Enabling IGMP Immediate-leave

To configure the switch to remove a group-port association after receiving an IGMP Leave message, use the `ip igmp fast-leave` command from INTERFACE VLAN mode.

To view the configuration, use the `show config` command from INTERFACE VLAN mode (Figure 12-7).

Figure 12-7. Enabling IGMP Snooping

```
FTOS(conf-if-vl-100)#show config
!
interface Vlan 100
no ip address
ip igmp snooping fast-leave
shutdown
FTOS(conf-if-vl-100)#
```

Disabling Multicast Flooding

If the switch receives a multicast packet that has an IP address of a group it has not learned (unregistered frame), the switch floods that packet out of all ports on the VLAN.

On the MXL Switch, when you configure `no ip igmp snooping flood`, the system forwards the frames on mrouter ports for first 96 IGMP snooping enabled VLANs. For all other VLANs, the unregistered multicast packets are dropped.

Specifying a Port as Connected to a Multicast Router

To statically specify a port in a VLAN as connected to a multicast router, use the `ip igmp snooping mrouter` command from `INTERFACE VLAN` mode.

To view the ports that are connected to multicast routers, use the `show ip igmp snooping mrouter` command from EXEC Privilege mode.

Configuring the Switch as Querier

Hosts that do not support unsolicited reporting wait for a general query before sending a membership report. When the multicast source and receivers are in the same VLAN, multicast traffic is not routed, so there is no querier. You must configure the switch to be the querier for a VLAN so that hosts send membership reports and the switch can generate a forwarding table by snooping.

To configure the switch to be the querier for a VLAN, first assign an IP address to the VLAN interface, and then use the `ip igmp snooping querier` command from `INTERFACE VLAN` mode.

- IGMP snooping querier does not start if there is a statically configured multicast router interface in the VLAN.
- The switch may lose the querier election if it does not have the lowest IP address of all potential queriers on the subnet.
- When enabled, the 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.

Adjusting the Last Member Query Interval

When the querier receives a Leave message from a receiver, it sends a group-specific query out of the ports specified in the forwarding table. If no response is received, it sends another. The amount of time that the querier waits to receive a response to the initial query before sending a second one is the last member query interval (LMQI). The switch waits one LMQI after the second query before removing the group-port entry from the forwarding table.

To adjust the LMQI, use the `ip igmp snooping last-member-query-interval` command from `INTERFACE VLAN` mode.

Fast Convergence after MSTP 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.

Designating a Multicast Router Interface

You can designate an interface 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 originating from receivers is forwarded to multicast router interfaces.

Interfaces

This chapter describes 100/1000/10000 Mbps Ethernet, 10 Gigabit Ethernet, and 40 Gigabit Ethernet interface types, both physical and logical, and how to configure them with the Dell Force10 operating software (FTOS).

Basic Interface Configuration:

- [Interface Types](#)
- [View Basic Interface Information](#)
- [Enable a Physical Interface](#)
- [Physical Interfaces](#)
- [Management Interfaces](#)
- [VLAN Interfaces](#)
- [Loopback Interfaces](#)
- [Null Interfaces](#)
- [Port Channel Interfaces](#)

Advanced Interface Configuration:

- [Bulk Configuration](#)
- [Interface Range Macros](#)
- [Monitor and Maintain Interfaces](#)
- [Splitting QSFP Ports to SFP+ Ports](#)
- [MTU Size on an Interface](#)
- [Layer 2 Flow Control Using Ethernet Pause Frames](#)
- [Configure MTU Size on an Interface](#)
- [Port-Pipes](#)
- [Auto-Negotiation on Ethernet Interfaces](#)
- [View Advanced Interface Information](#)

Interface Types

The following lists the different interface types.

Interface Type	Modes Possible	Default Mode	Requires Creation	Default State
Physical	L2, L3	Unset	No	Shutdown (disabled)
Management	N/A	N/A	No	No Shutdown (enabled)
Loopback	L3	L3	Yes	No Shutdown (enabled)
Null	N/A	N/A	No	Enabled
Port Channel	L2, L3	L3	Yes	Shutdown (disabled)
VLAN	L2, L3	L2	Yes (except default)	L2 - No Shutdown (enabled) L3 - Shutdown (disabled)

View Basic Interface Information

You have several options for viewing interface status and configuration parameters. The `show interfaces` command in EXEC mode lists all configurable interfaces on the chassis and has options to display 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 can list the interfaces configured in the port channel.



Note: To end output from the system, such as the output from the `show interfaces` command, enter CTRL+C and FTOS returns to the command prompt.



Note: The command line interface (CLI) output may be incorrectly displayed as 0 (zero) for the Rx/Tx power values. Perform a simple network management protocol (SNMP) query to obtain the correct power information.

Figure 13-1 shows the configuration and status information for one interface.

Figure 13-1. show interfaces Command Example (Partial)

```
FTOS#show interfaces tengigabitethernet 0/16
TenGigabitEthernet 0/16 is up, line protocol is up
Hardware is DellForce10Eth, address is 00:1e:c9:f1:00:05
    Current address is 00:1e:c9:f1:00:05
Server Port AdminState is Up
Pluggable media not present
Interface index is 38080769
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :tenG145001ec9f10005
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 10000 Mbit
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 5dlh18m
Queueing strategy: fifo
Input Statistics:
    34561 packets, 6266197 bytes
    38 64-byte pkts, 4373 over 64-byte pkts, 21491 over 127-byte pkts
    8659 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
    21984 Multicasts, 12577 Broadcasts
    0 runts, 0 giants, 0 throttles
    0 CRC, 0 overrun, 0 discarded
Output Statistics:
    44329 packets, 4722779 bytes, 0 underruns
    0 64-byte pkts, 44329 over 64-byte pkts, 0 over 127-byte pkts
    0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
    44329 Multicasts, 0 Broadcasts, 0 Unicasts
    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: 4d0h28m
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
FTOS#
```

Use the show ip interfaces brief command in EXEC Privilege mode to view which interfaces are enabled for Layer 3 data transmission. In [Figure 13-2](#), the TenGigabitEthernet interface 1/5 is in Layer 3 mode because an IP address has been assigned to it and the interface's status is operationally up.

Figure 13-2. show ip interfaces brief Command Example (Partial)

```
FTOS#show ip interface brief
Interface                IP-Address      OK? Method Status              Protocol
TenGigabitEthernet 1/0      unassigned     NO Manual administratively down down
TenGigabitEthernet 1/1      unassigned     NO Manual administratively down down
TenGigabitEthernet 1/2      unassigned     YES Manual up                up
TenGigabitEthernet 1/3      unassigned     YES Manual up                up
TenGigabitEthernet 1/4      unassigned     YES Manual up                up
TenGigabitEthernet 1/5      10.10.10.1    YES Manual up                up
TenGigabitEthernet 1/6      unassigned     NO Manual administratively down down
TenGigabitEthernet 1/7      unassigned     NO Manual administratively down down
TenGigabitEthernet 1/8      unassigned     NO Manual administratively down down
```

Use the show interfaces configured command in EXEC Privilege mode to view only configured interfaces.

To determine which physical interfaces are available, use the show running-config command in EXEC mode. This command displays all physical interfaces available on the switch ([Figure 13-3](#)).

Figure 13-3. show running-config Command Example (Partial)

```
FTOS#show running
Current Configuration ...
!
interface TenGigabitEthernet 9/6
 no ip address
 shutdown
!
interface TenGigabitEthernet 9/7
 no ip address
 shutdown
!
interface TenGigabitEthernet 9/8
 no ip address
 shutdown
!
interface TenGigabitEthernet 9/9
 no ip address
 shutdown
```


Enable a Physical Interface

After determining the type of physical interfaces available, you can enter INTERFACE mode by entering the interface *interface slot/port* command to enable and configure the interface.

To enter INTERFACE mode, follow these steps, starting in CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	<code>interface interface</code>	CONFIGURATION	Enter the keyword <code>interface</code> followed by the type of interface and slot/port information: <ul style="list-style-type: none">• For the Management interface on the stack-unit, enter the keyword <code>ManagementEthernet</code> followed by the slot/port information.• For a 10 Gigabit Ethernet interface, enter the keyword <code>TenGigabitEthernet</code> followed by the slot/port information.• For a 40-Gigabit Ethernet interface, enter the keyword <code>fortyGigE</code> followed by the slot/part information.
2	<code>no shutdown</code>	INTERFACE	Enter the <code>no shutdown</code> command to enable the interface.

To confirm that the interface is enabled, use the `show config` command in INTERFACE mode.

To leave INTERFACE mode, use the `exit` command or `end` command.

You cannot delete a physical interface.

Physical Interfaces

The management IP address on the D-fabric provides a dedicated management access to the system.

The switch interfaces support Layer 2 and Layer 3 traffic over the 100/1000/10000, 10-Gigabit, and 40-Gigabit Ethernet interfaces. These interfaces can also become part of virtual interfaces such as virtual local area networks (VLANs) or port channels.

For more information about VLANs, refer to [Bulk Configuration](#). For more information about port channels, refer to [Port Channel Interfaces](#).



FTOS Behavior: The MXL 10/40GbE switch systems use a single MAC address for all physical interfaces.

Configuration Task List for Physical Interfaces

By default, all interfaces are operationally disabled and traffic does not pass through them.

The following section includes information about optional configurations for physical interfaces:

- [Overview of Layer Modes](#)
- [Configure Layer 2 \(Data Link\) Mode](#)
- [Management Interfaces](#)
- [Auto-Negotiation on Ethernet Interfaces](#)
- [Adjust the Keepalive Timer](#)
- [Clear Interface Counters](#)

Overview of Layer Modes

On all systems running FTOS, you can place physical interfaces, port channels, and VLANs in Layer 2 mode or Layer 3 mode. By default, VLANs are in Layer 2 mode.

[Table 13-1](#) lists the different interface types and modes.

Table 13-1. Interface Types

Type of Interface	Possible Modes	Requires Creation	Default State
100/1000/10000 Ethernet, 40 Gigabit Ethernet, 10 Gigabit Ethernet	Layer 2 Layer 3	No	Shutdown (disabled)
Management	n/a	No	Shutdown (disabled)
Loopback	Layer 3	Yes	No shutdown (enabled)
Null interface	n/a	No	Enabled
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)

Configure Layer 2 (Data Link) Mode

To enable Layer 2 data transmissions through an individual interface, use the `switchport` command in `INTERFACE` mode. You cannot configure switching or Layer 2 protocols, such as the spanning tree protocol (STP), on an interface unless the interface has been set to Layer 2 mode.

[Figure 13-4](#) shows the basic configuration found in a Layer 2 interface.

Figure 13-4. `show config` Command Example of a Layer 2 Interface

```
FTOS(conf-if)#show config
!
interface Port-channel 1
  no ip address
  switchport
  no shutdown
FTOS(conf-if)#
```

To configure an interface in Layer 2 mode, use these commands in 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 about enabling and configuring STP, refer to [Layer 2 on page 305](#).

To view the interfaces in Layer 2 mode, use the command `show interfaces switchport` in EXEC mode.

Configure Layer 3 (Network) Mode

When you assign an IP address to a physical interface, you place it in Layer 3 mode. To enable Layer 3 mode on an individual interface, use the `ip address` command and `no shutdown` command in INTERFACE mode. You must configure the IP address and one of the interfaces in the system must be in Layer 3 mode before you configure or enter a Layer 3 protocol mode (for example, `open shortest path first [OSPF]`).

In all interface types except 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. Layer 2 traffic is unaffected by the `shutdown` command.

To view an example of a Layer 3 interface, use the `show config` command ([Figure 13-5](#)).

Figure 13-5. `show config` Command Example of a Layer 3 Interface

```
FTOS(conf-if)#show config
!
interface TenGigabitEthernet 1/5
 ip address 10.10.10.1 /24
 no shutdown
FTOS(conf-if)#
```

If an interface is in the incorrect layer mode for a given command, an error message is displayed. For example, in [Figure 13-6](#), the `ip address` command triggered an error message because the interface is in Layer 2 mode and the `ip address` command is a Layer 3 command only.

Figure 13-6. Error Message When Trying to Add an IP Address to a Layer 2 Interface

```
FTOS(conf-if)#show config
!
interface TenGigabitEthernet 1/2
 no ip address
 switchport
 no shutdown
FTOS(conf-if)#ip address 10.10.1.1 /24
% Error: Port is in Layer 2 mode TenGig 1/2.
FTOS(conf-if)#
```

← Error message

To determine the configuration of an interface, use the `show config` command in INTERFACE mode or the various `show interface` commands in EXEC mode.

To assign an IP address, use the following commands in INTERFACE mode:

Command Syntax	Command Mode	Purpose
no shutdown	INTERFACE	Enable the interface.
ip address <i>ip-address mask</i> [secondary]	INTERFACE	Configure a primary IP address and mask on the interface. The <i>ip-address</i> must be in dotted-decimal format (A.B.C.D) and the <i>mask</i> must be decimal and should be mentioned in slash format (/xx). Add the keyword <i>secondary</i> if the IP address is the interface's backup IP address.

You can only configure one (1) primary IP address per interface. You can configure up to 255 secondary IP addresses on a single interface.

To view all interfaces with an IP address assigned, use the show ip interfaces brief command in EXEC mode.

To view IP information on an interface in Layer 3 mode, use the show ip interface command in EXEC Privilege mode (Figure 13-7).

Figure 13-7. show ip interface Command Example

```
FTOS(conf-if-vl-10)#do sh int vl 10
Vlan 10 is up, line protocol is up
Address is 00:1e:c9:f1:03:38, Current address is 00:1e:c9:f1:03:38
Interface index is 1107787786
Internet address is 5.5.5.1/24
Mode of IP Address Assignment : MANUAL
DHCP Client-ID: vlan10001ec9f10338
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 1000 Mbit
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:01:09
Queueing strategy: fifo
Time since last interface status change: 00:00:46
```

Management Interfaces

The IOM management interface has both a public IP and private IP address on the internal Fabric D interface. The public IP address is exposed to the outside world for Web GUI configurations/WSMAN and other proprietary traffic. You can statically configure the public IP address or obtain the IP address dynamically using the dynamic host configuration protocol (DHCP).



Note: When you shut down a management interface, connectivity to the interface's private IP address is disabled.

You can access the full switch using:

- Internal RS-232 using the chassis management controller (CMC). Telnet into CMC and do a connect -b switch-id to get console access to corresponding IOM.
- External serial port with a universal serial bus (USB) connector (front panel): connect using the IOM front panel USB serial line to get console access (Labeled as USB B).
- Telnet/others using the public IP interface on the fabric D interface.
- CMC through the private IP interface on the fabric D interface.

The MXL Switch system supports the management ethernet interface as well as the standard interface on any front-end port. You can use either method to connect to the system.

Configure Management Interfaces on the MXL Switch

On the MXL Switch IO Module, the dedicated management interface provides management access to the system. You can configure this interface with FTOS, but the configuration options on this interface are limited. You cannot configure gateway addresses and IP addresses if it appears in the main routing table of FTOS. In addition, proxy address resolution protocol (ARP) is not supported on this interface.

For additional management access, IOM supports the default VLAN (VLAN 1) L3 interface in addition to the public fabric D management interface. You can assign the IP address for the VLAN 1 default management interface using the setup wizard (or) through the CLI.

If you do not configure the VLAN 1 default using the wizard or CLI presented in startup-config, by default, the VLAN 1 management interface gets its IP address using DHCP.

To configure a management interface, use the following command in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
interface Managementethernet <i>interface</i>	CONFIGURATION	Enter the slot and the port (0). Slot range: 0-0

To configure IP addresses on a management interface, use the following command in MANAGEMENT INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip address <i>ip-address mask</i>	INTERFACE	Configure an IP address and mask on the interface. <ul style="list-style-type: none">• <i>ip-address mask</i>: enter an address in dotted-decimal format (A.B.C.D), the mask must be in /prefix format (/x)

To access the management interface from another LAN, you must configure the management route command to point to the management interface.

There is only one management interface for the whole stack.

You can manage the MXL Switch from any port. Configure an IP address for the port using the `ip address` command. Enable the IP address for the port using the `no shutdown` command. You can use the `description` command from `INTERFACE` mode to note that the interface is the management interface. There is no separate management routing table, so you must configure all routes in the IP routing table (use the `ip route` command).

To display the configuration for a given port, use the `show interface` command from `EXEC Privilege mode` (Figure 13-8).

To display the routing table for a given port, use the `show ip route` command from `EXEC Privilege mode`.

Figure 13-8. Viewing Management Routes

```

FTOS#show int tengig 0/16
TenGigabitEthernet 0/16 is up, line protocol is down
Hardware is DellForce10Eth, address is 00:1e:c9:bb:02:c2
  Current address is 00:1e:c9:bb:02:c2
Server Port AdminState is Down
Pluggable media not present
Interface index is 38080769
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :tenG145001ec9bb02c2
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 2w4d2h
Queueing strategy: fifo
Input Statistics:
  0 packets, 0 bytes
  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:
  0 packets, 0 bytes, 0 underruns
  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 Unicasts
  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: 2w4d2h
FTOS#

```

VLAN Interfaces

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 about VLANs and Layer 2, refer to [Layer 2](#) and [Virtual LANs \(VLAN\)](#).



Note: To monitor VLAN interfaces, use the Management Information Base for Network Management of TCP/IP-based internets: MIB-II (RFC 1213).



Note: You cannot simultaneously use egress rate shaping and ingress rate policing on the same VLAN.

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 about configuring different routing protocols, refer to the chapters on the specific protocol.

A consideration for including VLANs in routing protocols is that you must configure the `no shutdown` command. (For routing traffic to flow, the VLAN must be enabled.)



Note: You cannot assign an IP address to the Default VLAN, which, by default, is VLAN 1. To assign another VLAN ID to the Default VLAN, use the `default vlan-id vlan-id` command.

Assign an IP address to an interface with the following command in the INTERFACE mode:

Command Syntax	Command Mode	Purpose
<code>ip address <i>ip-address mask</i> [secondary]</code>	INTERFACE	Configure an IP address and mask on the interface. <ul style="list-style-type: none"><i>ip-address mask</i>: enter an address in dotted-decimal format (A.B.C.D) and the mask must be in slash format (/24).<i>secondary</i>: the IP address is the interface's backup IP address.

[Figure 13-9](#) shows a sample configuration of a VLAN participating in an OSPF process.

Figure 13-9. Sample Layer 3 Configuration of a VLAN

```
interface Vlan 10
ip address 1.1.1.2/24
tagged TenGigabitEthernet 2/2-13
tagged TenGigabitEthernet 5/0
ip ospf authentication-key Dell Force10
ip ospf cost 1
ip ospf dead-interval 60
ip ospf hello-interval 15
no shutdown
!
```

Loopback Interfaces

A Loopback interface is a virtual interface in which the software emulates an interface. Packets routed to it are processed locally. Because 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 CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
interface loopback <i>number</i>	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 EXEC mode.

To delete a loopback interface, use the no interface loopback *number* command in CONFIGURATION mode.

Many of the same commands found in the physical interface are found in loopback interfaces. For more information, refer to [Access Control Lists \(ACLs\)](#).

Null Interfaces

The null interface is a virtual interface. There is only one null interface. It is always up, but no traffic is transmitted through this interface.

To enter INTERFACE mode of the null interface, use the following command in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
interface null 0	CONFIGURATION	Enter INTERFACE mode of null interface.

The only configurable command in INTERFACE mode of 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](#)
- [Port Channel Benefits](#)
- [Port Channel Implementation](#)
- [Configuration Task List for Port Channel Interfaces](#)

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.

A port channel provides redundancy by aggregating 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

A port channel interface provides many benefits, including easy management, link redundancy, and sharing.

Port channels are transparent to network configurations and can be modified 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 create larger-capacity interfaces by utilizing a group of lower-speed links. For example, you can build a 40-Gigabit interface by aggregating four 10-Gigabit Ethernet interfaces together. If one of the four interfaces fails, traffic is redistributed across the three 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 the link aggregation control protocol (LACP). For more information, refer to [Link Aggregation Control Protocol \(LACP\)](#).

Table 13-2 lists the number of port channels per platform.

Table 13-2. Number of Port Channels per Platform

Platform	Port-channels	Members/Channel
MXL 10/40GbE Switch IO Module	128	16

As soon as a port channel is configured, FTOS treats it like a physical interface. For example, IEEE 802.1Q tagging is maintained while the physical interface is in the port channel.

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 switch resets and chassis reloads.

A physical interface can belong to only one port channel at a time.

Each port channel must contain interfaces of the same interface type/speed.

Port channels can contain a mix of 100, 1000, or 10000 Mbps Ethernet interfaces and TenGigabit Ethernet interfaces. The interface speed (100, 1000, or 10000 Mbps) used by the port channel is determined by the first port channel member that is physically up. FTOS disables the interfaces that do not 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 TenGigabit Ethernet interface, all interfaces at 1000 Mbps are kept up, and all 100/1000/10000 interfaces that are not set to 1000 Mbps speed or auto negotiate are disabled.

100/1000/10000 Mbps Interfaces in Port Channels

When both 100/1000/10000 interfaces and TenGigabitEthernet 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 (TenGig 0/0, 0/1, 0/2, 0/3) in which TenGig 0/0 and TenGig 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 tengigabitethernet 0/0-3` while in port channel interface mode, and FTOS determines if the first interface specified (TenGig 0/0) is up. After it is up, the common speed of the port channel is 100 Mb/s. FTOS disables those interfaces configured with speed 1000 Mb/s 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 by setting the speed of the TenGig 0/0 interface to 1000 Mb/s.

Configuration Task List for Port Channel Interfaces

To configure a port channel (LAG), you use the commands similar to those found in physical interfaces. By default, no port channels are configured in the startup configuration.

- [Create a port channel](#) (mandatory)
- [Add a Physical Interface to a Port Channel](#) (mandatory)
- [Reassign an Interface to a New Port Channel](#) (optional)
- [Configure the Minimum oper up Links in a Port Channel \(LAG\)](#) (optional)
- [Add or Remove a Port Channel from a VLAN](#) (optional)
- [Assign an IP Address to a Port Channel](#) (optional)
- [Delete or Disable a Port Channel](#) (optional)

Create a Port Channel

You can create up to 128 port channels with 16 port members per group on an MXL Switch.

To configure a port channel, follow these steps, starting in CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	<code>interface port-channel <i>id-number</i></code>	CONFIGURATION	Create a port channel.
2	<code>no shutdown</code>	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. To place the port channel in Layer 2 mode or configure an IP address to place the port channel in Layer 3 mode, use the `switchport` command.

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



Note: Port channels can contain a mix of 100/1000/10000 Ethernet interfaces and 10 Gigabit Ethernet interfaces, but FTOS disables the interfaces that are not the same speed of the first channel member in the port channel (see [100/1000/10000 Mbps Interfaces in Port Channels](#)).

You can add any physical interface to a port channel if the interface configuration is minimal. You can configure only the following commands on an interface if it is a member of a port channel:

- `description`
- `shutdown/no shutdown`
- `mtu`

- ip mtu (if the interface is on a Jumbo-enabled by default.)



Note: The MXL Switch supports jumbo frames by default (the default maximum transmission unit [MTU] is 1554 bytes) You can configure the MTU using the mtu command from INTERFACE mode.

To view the interface's configuration, enter INTERFACE mode for that interface and use the show config command or from EXEC Privilege mode, use the show running-config interface *interface* command.

To add a physical interface to a port channel, use these commands in the following sequence in INTERFACE mode of a port channel:

Step	Command Syntax	Command Mode	Purpose
1	channel-member <i>interface</i>	INTERFACE PORT-CHANNEL	Add the interface to a port channel. The <i>interface</i> 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 command in EXEC Privilege mode ([Figure 13-10](#)).

Figure 13-10. show interfaces port-channel brief Command Example

```
FTOS#show int port brief
Codes: L - LACP Port-channel

LAG  Mode  Status      Uptime      Ports
 1   L3   down       00:00:00   Te 0/16   (Down)
FTOS#
```

Figure 13-11 shows 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.

Figure 13-11. show interface port-channel Command Example

```
FTOS#show int port-channel
Port-channel 1 is down, line protocol is down
Hardware address is 00:1e:c9:f1:00:05, Current address is 00:1e:c9:f1:00:05
Interface index is 1107755009
Minimum number of links to bring Port-channel up is 1
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :lag1001ec9f10005
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
Members in this channel:  Te 0/16(D)
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:05:44
Queueing strategy: fifo
Input Statistics:
  0 packets, 0 bytes
  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:
  0 packets, 0 bytes, 0 underruns
  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 Unicasts
  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: 00:05:44
```

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 data units (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 13-12 shows, the interface TenGigabitEthernet 1/6 is part of port channel 1, which is in Layer 2 mode, and an error message appeared when an IP address was configured.

Figure 13-12. Error Message

```

FTOS(conf-if-po-1)#show config
!
interface Port-channel 1
  no ip address
  channel-member TenGigabitEthernet 0/16
  shutdown
FTOS(conf-if-po-1)#
FTOS(conf-if-po-1)#int tengig 1/6
FTOS(conf-if)#ip address 10.56.4.4 /24
% Error: Te 1/6 Port is part of a LAG. ← Error message
FTOS(conf-if)#

```

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.

To reassign an interface to a new port channel, follow these steps, starting in INTERFACE mode of a port channel:

Step	Command Syntax	Command Mode	Purpose
1	<code>no channel-member <i>interface</i></code>	INTERFACE PORT-CHANNEL	Remove the interface from the first port channel.
2	<code>interface port-channel id <i>number</i></code>	INTERFACE PORT-CHANNEL	Change to the second port channel INTERFACE mode.
3	<code>channel-member <i>interface</i></code>	INTERFACE PORT-CHANNEL	Add the interface to the second port channel.

Figure 13-13 shows an example of moving the TenGigabitEthernet 1/8 interface from port channel 4 to port channel 3.

Figure 13-13. Command Example from Reassigning an Interface to a Different Port Channel

```

FTOS(conf-if-po-1)#show config
!
interface Port-channel 1
  no ip address
  channel-member TenGigabitEthernet 0/16
  shutdown
FTOS(conf-if-po-1)#no chann tengig 1/8
FTOS(conf-if-po-1)#int port 5
FTOS(conf-if-po-5)#channel tengig 1/8
FTOS(conf-if-po-5)#show conf
!
interface Port-channel 5
  no ip address
  channel-member TenGigabitEthernet 1/8
  shutdown
FTOS(conf-if-po-5)#

```

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 in “oper up” status. To configure the minimum links, use the following command in INTERFACE mode:

Command Syntax	Command Mode	Purpose
<code>minimum-links <i>number</i></code>	INTERFACE	Enter the number of links in a LAG that must be in “oper up” status. Default: 1

Figure 13-14 shows an example of configuring five minimum “oper up” links in a port channel.

Figure 13-14. Example of Using the minimum-links Command

```
FTOS#config t
FTOS(conf)#int po 1
FTOS(conf-if-po-1)#minimum-links 5
FTOS(conf-if-po-1)#
```

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 (use the switchport command).

To add a port channel to a VLAN, use either of the following commands:

Command Syntax	Command Mode	Purpose
<code>tagged port-channel <i>id number</i></code>	INTERFACE VLAN	Add the port channel to the VLAN as a tagged interface. An interface with tagging enabled can belong to multiple VLANs.
<code>untagged port-channel <i>id number</i></code>	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:

Command Syntax	Command Mode	Purpose
<code>no tagged port-channel <i>id number</i></code>	INTERFACE VLAN	Remove the port channel with tagging enabled from the VLAN.
<code>no untagged port-channel <i>id number</i></code>	INTERFACE VLAN	Remove the port channel without tagging enabled from the VLAN.

To see which port channels are members of VLANs, use the show vlan command in 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 INTERFACE mode:

Command Syntax	Command Mode	Purpose
<code>ip address <i>ip-address mask</i> [secondary]</code>	INTERFACE	Configure an IP address and mask on the interface. <ul style="list-style-type: none"> • <i>ip-address mask</i>: enter an address in dotted-decimal format (A.B.C.D) and the mask must be in slash format (/24). • <i>secondary</i>: the IP address is the interface's backup IP address.

Delete or Disable a Port Channel

To delete a port channel, use the `no interface portchannel channel-number` command in CONFIGURATION mode.

When you disable a port channel (using the shutdown command) all interfaces within the port channel are operationally down also.

Bulk Configuration

Bulk configuration allows 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 configuring 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.



Note: When creating an interface range, interfaces appear in the order they were entered and are not sorted.

To display all interfaces that have been validated under the interface range context, use the `show range` command in Interface Range mode.

To display the running configuration only for interfaces that are part of interface range, use the `show configuration` command in Interface Range mode.

Bulk Configuration Examples

The following are examples of using the interface range command for bulk configuration:

- [Create a Single-Range](#)
- [Create a Multiple-Range](#)
- [Exclude Duplicate Entries](#)
- [Exclude a Smaller Port Range](#)
- [Overlap Port Ranges](#)
- [Commas](#)
- [Add Ranges](#)

Create a Single-Range

Figure 13-15. Creating a Single-Range Bulk Configuration

```
FTOS(conf)# interface range tengigabitethernet 5/1 - 23
FTOS(conf-if-range-te-5/1-23)# no shutdown
FTOS(conf-if-range-te-5/1-23)#
```

Create a Multiple-Range

Figure 13-16. Creating a Multiple-Range Prompt

```
FTOS(conf)#interface range tengigabitethernet 3/0 , tengigabitethernet 2/1 - 47 , vlan 1000
FTOS(conf-if-range-te-2/1-47)#
```

Exclude Duplicate Entries

Duplicate single interfaces and port ranges are excluded from the resulting interface range prompt.

Figure 13-17. Interface Range Prompt Excluding Duplicate Entries

```
FTOS(conf)#interface range vlan 1 , vlan 1 , vlan 3 , vlan 3
FTOS(conf-if-range-vl-1,vl-3)#
FTOS(conf)#interface range tengigabitethernet 2/0 - 23 , tengigabitethernet 2/0 - 23 , tengigab 2/0 - 23
FTOS(conf-if-range-te-2/0-23)#
```

Exclude a Smaller Port Range

If the interface range has multiple port ranges, the smaller port range is excluded from the prompt.

Figure 13-18. Interface Range Prompt Excluding a Smaller Port Range

```
FTOS(conf)#interface range tengigabitethernet 2/0 - 23 , tengigab 2/1 - 10
FTOS(conf-if-range-te-2/0-23)#
```

Overlap Port Ranges

If overlapping port ranges are specified, the port range is extended to the smallest start port number and largest end port number.

Figure 13-19. Interface Range Prompt Including Overlapping Port Ranges

```
FTOS(conf)#inte ra tengig 2/1 - 11 , tengig 2/1 - 23
FTOS(conf-if-range-te-2/1-23)#
```

Commas

The example below shows how to use commas to add different interface types to the range, enabling all Ten Gigabit Ethernet interfaces in the range 5/1 to 5/23 and both Ten Gigabit Ethernet interfaces 1/1 and 1/2.

Figure 13-20. Multiple-Range Bulk Configuration Gigabit Ethernet and Ten-Gigabit Ethernet

```
FTOS(conf-if)# interface range tengigabitethernet 5/1 - 23, tengigabitethernet 1/1 - 2
FTOS(conf-if-range-te-5/1-23)# no shutdown
FTOS(conf-if-range-te-5/1-23)#
```

Add Ranges

The example below shows how to use commas to add VLAN and port-channel interfaces to the range.

Figure 13-21. Multiple-Range Bulk Configuration with VLAN, and Port-channel

```
FTOS(conf-ifrange-te-5/1-23-te-1/1-2)# interface range Vlan 2 - 100 , Port 1 - 25
FTOS(conf-if-range-te-5/1-23-te-1/1-2-vl-2-100-po-1-25)# no shutdown
FTOS(conf-if-range)#
```

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 the following command:

Command Syntax	Command Mode	Purpose
FTOS (config)# define <i>interface-range macro_name</i> {vlan <i>vlan_ID - vlan_ID</i> } {{tengigabitethernet fortyGigE} <i>slot/interface - interface</i> } [, {vlan <i>vlan_ID - vlan_ID</i> } {{tengigabitethernet fortyGigE} <i>slot/interface - interface</i> }]	CONFIGURATION	Defines the interface-range macro and saves it in the running configuration file.

Define the Interface Range

Figure 13-22 shows how to define an interface-range macro named “test” to select Ten Gigabit Ethernet interfaces 5/1 through 5/4.

Figure 13-22. Define an Interface Range Macro

```
FTOS(conf)# define interface-range test tengigabitethernet 5/1 - 4
```

Choose 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 <i>name</i>	CONFIGURATION	Selects the interfaces range to be configured using the values saved in a named interface-range macro.

Figure 13-23 shows how to change to the interface-range configuration mode using the interface-range macro named “test”.

Figure 13-23. Change the Interface Range Configuration

```
FTOS(conf)# interface range macro test  
FTOS(conf-if)#
```

Monitor and Maintain Interfaces

Monitor interface statistics with the monitor interface command. This command displays an ongoing list of the interface status (up/down), number of packets, traffic statistics, etc.

Command Syntax	Command Mode	Purpose
monitor interface <i>interface</i>	EXEC Privilege	View the interface's statistics. Enter the type of interface and slot/port information: <ul style="list-style-type: none">• For a 100/1000/10000 Ethernet interface, enter the keyword TenGigabitEthernet followed by the slot/port information.• For a 10 Gigabit Ethernet interface, enter the keyword TenGigabitEthernet followed by the slot/port information.• For a 40-Gigabit Ethernet interface, enter the keyword fortyGigE followed by the slot/port information.

The information displays in a continuous run, refreshes every two seconds by default (Figure 13-24). Use the following keys to manage the output.

m - Change mode

l - Page up

T - Increase refresh interval (by 1 second)

q - Quit

c - Clear screen

a - Page down

t - Decrease refresh interval (by 1 second)

Figure 13-24. monitor interface Command Example

```

FTOS#monitor interface tengig 3/1

Dell Force10 uptime is 1 day(s), 4 hour(s), 31 minute(s)
  Monitor time: 00:00:00   Refresh Intvl.: 2s

Interface: TenGig 3/1, Disabled, Link is Down, Linespeed is 1000 Mbit

Traffic statistics:
  Current          Rate          Delta
  Input bytes:    0          0 Bps          0
  Output bytes:   0          0 Bps          0
  Input packets:  0          0 pps          0
  Output packets: 0          0 pps          0
  64B packets:   0          0 pps          0
  Over 64B packets: 0        0 pps          0
  Over 127B packets: 0        0 pps          0
  Over 255B packets: 0        0 pps          0
  Over 511B packets: 0        0 pps          0
  Over 1023B packets: 0        0 pps          0
Error statistics:
  Input underruns: 0          0 pps          0
  Input giants:   0          0 pps          0
  Input throttles: 0          0 pps          0
  Input CRC:      0          0 pps          0
  Input IP checksum: 0        0 pps          0
  Input overrun:  0          0 pps          0
  Output underruns: 0          0 pps          0
  Output throttles: 0          0 pps          0

  m - Change mode          c - Clear screen
  l - Page up              a - Page down
  T - Increase refresh interval  t - Decrease refresh interval
  q - Quit

FTOS#

```

Maintenance Using TDR

The time domain reflectometer (TDR) is supported on all Dell Force10 switch/routers. 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 unterminated, 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. Do not use TDR 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 100/1000/10000 BASE-T modules, following these steps using the `tdr-cable-test` command.

Step	Command Syntax	Command Mode	Usage
1	<code>tdr-cable-test tengigabitethernet <slot>/<port></code>	EXEC Privilege	To test for cable faults on the TenGigabitEthernet cable. <ul style="list-style-type: none"> Between two ports, you must not start the test on both ends of the cable. Enable the interface before starting the test. The port must be enabled to run the test or the test prints an error message.
2	<code>show tdr tengigabitethernet <slot>/<port></code>	EXEC Privilege	Displays TDR test results.

Splitting QSFP Ports to SFP+ Ports

The MXL 10/40GbE switch supports splitting a 40GbE port on the base module or a 2-Port 40GbE QSFP+ module into four 10GbE SFP+ ports using a 4x10G breakout cable.



Note: By default, the 40GbE ports on a 2-Port 40GbE QSFP+ module come up in 4x10GbE (quad) mode as eight 10GbE ports. On the base module, you must convert the 40GbE ports to 4x10GbE mode as described below.

Command Syntax	Command Mode	Purpose
<code>stack-unit <stack-unit> port <number> portmode quad</code>	CONFIGURATION	Split a single 40G port into 4-10G ports. <p><i>stack-unit:</i> Enter the stack member unit identifier of the stack member to reset. Range: 0 to 5 To display the stack-unit number, enter the show system brief command.</p> <p><i>port <port number>:</i> Enter the port number of the 40G port to be split. Valid values on base module: 33 or 37; OPTM SLOT 0: 41 or 45; OPTM SLOT 1: 49 or 53.</p> <p><i>portmode quad:</i> Identifies the uplink port as a split 10GbE SFP+ port.</p> <p>Then save the configuration and reload the switch. FTOS# write memory FTOS#reload</p>

Merging SFP+ Ports to QSFP 40G Ports

To remove FANOUT mode in 40G QSFP Ports, use the following commands:

Command Syntax	Command Mode	Purpose
no stack-unit <i>stack-unit</i> port <i>number</i> portmode quad	CONFIGURATION	<p>Merge 4-10G ports to a single 40G port.</p> <p><i>stack-unit</i>: Enter the stack member unit identifier of the stack member to reset. Range: 0 to 5 To display the stack-unit number, enter the show system brief command.</p> <p>port <<i>port number</i>>: Enter the port number of the 40GbE QSFP+ port. Valid values on base module: 33 or 37; OPTM SLOT 0: 41 or 45; OPTM SLOT 1: 49 or 53.</p> <p>portmode quad: Identifies the uplink port as a split 10GbE SFP+ port.</p> <p>Then save the configuration and reload the switch. FTOS# write memory FTOS#reload</p>

Important Points

- You cannot use split ports as stack-link to stack an MXL Switch.
- Split ports cannot be a part of any stacked system.
- The quad port must be in a default configuration before it can be split into 4x10G ports.
- The 40G port is lost in the configuration when the port is split, so be sure the port is also removed from other L2/L3 feature configurations.
- The system must be reloaded after issuing the CLI for the change to take effect.

MTU Size on an Interface

The link MTU is the frame size of a packet. 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 is defined as the entire Ethernet packet (Ethernet header + FCS + payload)

Because different networking vendors define MTU differently, check their documentation when planning MTU sizes across a network.

Table 13-3 lists the range for each transmission media.

Table 13-3. MTU Range

Transmission Media	MTU Range (in bytes)
Ethernet	594-12000 = link MTU 576-11982 = IP MTU

Layer 2 Flow Control Using Ethernet Pause Frames

Ethernet pause frames allow for a temporary stop in data transmission. A situation may arise where a sending device may transmit data faster than a destination device can accept it. The destination sends a pause frame back to the source, stopping the sender's transmission for a period of time.

The globally assigned 48-bit Multicast address 01-80-C2-00-00-01 is used to send and receive pause frames. To allow full duplex flow control, stations implementing the pause operation instruct the MAC to enable reception of frames with a destination address equal to this multicast address.

The pause frame is defined by IEEE 802.3x and uses MAC Control frames to carry the pause commands. Ethernet pause frames are supported on full duplex only. The only configuration applicable to half duplex ports is rx off tx off.

Note that if a port is over-subscribed, Ethernet Pause Frame flow control does not ensure no loss behavior.

The following error message appears when trying to enable flow control when half duplex is already configured:

```
Can't configure flowcontrol when half duplex is configure, config ignored.
```

The following error message appears when trying to enable half duplex and flow control configuration is on:

```
Can't configure half duplex when flowcontrol is on, config ignored.
```

Enable Pause Frames



Note: If rx flow control is disabled, Dell Force10 recommends rebooting the system.

You must enable the Ethernet pause frames flow control on all ports on a chassis. If not, the system may exhibit unpredictable behavior.

The flow-control sender and receiver must be on the same port-pipe. Flow control is not supported across different port-pipes (also refer to [Enabling and Disabling iSCSI Optimization on page 281](#)).

Command Syntax	Command Mode	Purpose
flowcontrol rx [<i>off</i> <i>on</i>] tx [<i>off</i> <i>on</i>] [threshold {<1-2047> <1-2013> <1-2013>}]	INTERFACE	Control how the system responds to and generates 802.3x pause frames on 10 and 40Gig ports. Defaults: rx off
	Parameters:	
	rx on: Enter the keywords rx on to process the received flow control frames on this port.	
	rx off: Enter the keywords rx off to ignore the received flow control frames on this port.	
	tx on: Enter the keywords tx on to send control frames from this port to the connected device when a higher rate of traffic is received.	
	tx off: Enter the keywords tx off so that flow control frames are not sent from this port to the connected device when a higher rate of traffic is received.	



Note: After you disable DCB, if link-level flow control is not automatically enabled on an interface, manually shut down the interface (**shutdown** command) and re-enable it (**no shutdown** command) to enable flow control.

Configure MTU Size on an Interface

If a packet includes a Layer 2 header, the difference in bytes between the link MTU and IP MTU must be enough to include 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-byte IP MTU + 22-byte VLAN Tag = 1422-byte link MTU

The MTU range is 592-12000, with a default of 1554.

Table 13-4 lists the various Layer 2 overheads found in FTOS and the number of bytes.

Table 13-4. 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.

For 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 the 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.

For example, the VLAN contains tagged members with a link MTU of 1522 and an IP MTU of 1500 and untagged members with a link MTU of 1518 and an 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.

Port-Pipes

A high-speed data bus connection used to switch traffic between front-end ports is known as the port pipe. A port pipe is a Dell Force10 term for the hardware path that packets follow through a system. The MXL Switch supports single port pipe only.

Auto-Negotiation on Ethernet Interfaces

Setting Speed and Duplex Mode of Ethernet Interfaces

By default, auto-negotiation of speed and duplex mode is enabled on 100/1000/10000 Base-T Ethernet interfaces. Only 10GbE interfaces do not support auto-negotiation. When using 10GbE interfaces, verify that the settings on the connecting devices are set to no auto-negotiation.

The local interface and the directly connected remote interface must have the same setting. Auto-negotiation is the easiest way to accomplish these settings, as long as the remote interface is capable of auto-negotiation.



Note: As a best practice, Dell Force10 recommends keeping auto-negotiation enabled. Auto-negotiation should only be disabled on switch ports that attach to devices not capable of supporting negotiation or where connectivity issues arise from interoperability issues.

For 100/1000/10000 Ethernet interfaces, 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. In FTOS, the speed 1000 command is an exact equivalent of speed auto 1000 in IOS.

To discover whether the remote and local interface require manual speed synchronization, and to manually synchronize them if necessary, follow these steps (also refer to [Figure 13-26 on page 251](#)).

Step	Task	Command Syntax	Command Mode
1	Determine the local interface status. Refer to Figure 13-25 .	show interfaces [<i>interface</i>] status	EXEC Privilege
2	Determine the remote interface status.	[Use the command on the remote system that is equivalent to the above command.]	EXEC EXEC Privilege
3	Access CONFIGURATION mode.	config	EXEC Privilege
4	Access the port.	interface <i>interface slot/port</i>	CONFIGURATION
5	Set the local port speed.	speed {100 1000 10000 auto}	INTERFACE
6	Optionally, set full- or half-duplex.	duplex {half full}	INTERFACE
7	Disable auto-negotiation on the port. If the speed is set to 1000, you do not need to disable auto-negotiation	no negotiation auto	INTERFACE
8	Verify configuration changes.	show config	INTERFACE



Note: The show interfaces status command ([Figure 13-25](#)) displays link status, but not administrative status. For link and administrative status, use the show ip interface [*interface* | brief] [configuration] command.

Figure 13-25. show interfaces status Command Example

```
FTOS#show interfaces status
Port      Description  Status Speed      Duplex Vlan
Te 0/1    Te 0/1       Down  Auto      Auto  --
Te 0/2    Te 0/2       Down  Auto      Auto  --
Te 0/3    Te 0/3       Down  Auto      Auto  --
Te 0/4    Te 0/4       Down  Auto      Auto  --
Te 0/5    Te 0/5       Down  Auto      Auto  --
Te 0/6    Te 0/6       Down  Auto      Auto  --
Te 0/7    Te 0/7       Down  Auto      Auto  --
Te 0/8    Te 0/8       Down  Auto      Auto  --
Te 0/9    Te 0/9       Down  Auto      Auto  --
Te 0/10   Te 0/10      Down  Auto      Auto  --
Te 0/11   Te 0/11      Down  Auto      Auto  --
Te 0/12   Te 0/12      Down  Auto      Auto  --
Te 0/13   Te 0/13      Down  Auto      Auto  --
[output omitted]
```

In Figure 13-25, several ports display “Auto” in the Speed field, including port 0/1. In Figure 13-26, the speed of port 0/1 is set to 100 Mb and then its auto-negotiation is disabled.

Figure 13-26. Setting Port Speed Example

```
FTOS#configure
FTOS(conf)#interface tengig 0/1
FTOS(Interface 0/1)#speed 100
FTOS(Interface 0/1)#duplex full
FTOS(Interface 0/1)#no negotiation auto
FTOS(Interface 0/1)#show config
!
interface TenGigabitEthernet 0/1
no ip address
speed 100
duplex full
no shutdown
```

Setting Auto-Negotiation Options

The negotiation auto command provides a mode option for configuring an individual port to forced master/forced slave after you enable auto-negotiation.



Caution: Ensure that only one end of the node is configured as forced-master and the other is configured as forced-slave. If both are configured the same (that is, both as forced-master or both as forced-slave), the show interface command flaps between an auto-neg-error and forced-master/slave states.

Table 13-5. Auto-Negotiation, Speed, and Duplex Settings on Different Optics

Command	mode	10GbaseT module	10G SFP+ optics	1G SFP optics	Copper SFP - 1000baseT	Comments
speed 100	interface-config mode	Supported	Not supported (Error message is thrown) (% Error: Speed 100 not supported on this interface, config ignored Te 0/49)	Not supported (Error message is thrown) (% Error: Speed 100 not supported on this interface, config ignored Te 0/49)	% Error: Speed 100 not supported on this interface,	
speed auto	interface-config mode	Supported	Not supported	Not supported	Not supported	Error messages not thrown wherever it says not supported
speed 1000	interface-config mode	Supported	Supported	Supported	Supported	
speed 10000	interface-config mode	Supported	Supported	Not Supported	Not supported	Error messages not thrown wherever it says not supported
negotiation auto	interface-config mode	Supported	Not supported (Should some error message be thrown?)	Not supported	Not supported	Error messages not thrown wherever it says not supported
duplex half	interface-config mode	Supported	CLI not available	CLI not available	Invalid Input error- CLI not available	
duplex full	interface-config mode	Supported	CLI not available	CLI not available	Invalid Input error-CLI not available	

Figure 13-27 shows the auto-negotiation options.

Figure 13-27. Setting Auto-Negotiation Options

```
FTOS(conf)# int tengig 0/0
FTOS(conf-if)#neg auto
FTOS(conf-if-autoneg)# ?

end                Exit from configuration mode
exit              Exit from autoneg configuration mode
mode             Specify autoneg mode
no               Negate a command or set its defaults
show            Show autoneg configuration information
FTOS(conf-if-autoneg)#mode ?
forced-master    Force port to master mode
forced-slave     Force port to slave mode
FTOS(conf-if-autoneg)#
```

Adjust the Keepalive Timer

To change the time interval between keepalive messages on the interfaces, use the `keepalive` command. The interface sends keepalive messages to itself to test network connectivity on the interface.

To change the default time interval between keepalive messages, use the following command:

Command Syntax	Command Mode	Purpose
<code>keepalive [seconds]</code>	INTERFACE	Change the default interval between keepalive messages.

To view the new setting, use the `show config` command in INTERFACE mode.

View Advanced Interface Information

Display Only Configured Interfaces

The following options have been implemented for the `show [ip | running-config] interfaces` command. When you use the configured keyword, only interfaces that have non-default configurations are displayed.

[Figure 13-28](#) shows the possible show commands that have the configured keyword available.

Figure 13-28. show Commands with configured Keyword Examples

```
FTOS#show interfaces configured
FTOS#show interfaces tengigabitEthernet 0 configured
FTOS#show ip interface configured
FTOS#show ip interface tengigabitEthernet 1 configured
FTOS#show interfaces fortygigabitEthernet 0 configured
FTOS#show ip interface fortygigabitEthernet 1 configured
FTOS#show ip interface brief configured
FTOS#show running-config interfaces configured
FTOS#show running-config interface tengigabitEthernet 1 configured
```

In 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 13-29) displays the interface, whether the interface supports IEEE 802.1Q tagging or not, and the VLANs to which the interface belongs.

Figure 13-29. show interfaces switchport Command Example

```
FTOS#show interfaces switchport
Name: TenGigabitEthernet 13/0
802.1QTagged: True
Vlan membership:
Vlan    2

Name: TenGigabitEthernet 13/1
802.1QTagged: True
Vlan membership:
Vlan    2

Name: TenGigabitEthernet 13/2
802.1QTagged: True
Vlan membership:
Vlan    2

Name: TenGigabitEthernet 13/3
802.1QTagged: True
Vlan membership:
Vlan    2
--More--
```


Configure Interface Sampling Size

To configure the number of seconds of traffic statistics to display in the show interfaces output, use the rate-interval command in INTERFACE mode.

You can enter any value between five and 299 seconds (the default). If you enter 1 to 5 seconds, software polling is done at 5 sec interval. If you enter 6 to 10 sec, software polling is done at 10 sec interval. For any other value, software polling is done once every 15 seconds. So, for example, if you enter “19”, you actually get a sample of the past 15 seconds.

All LAG members inherit the rate interval configuration from the LAG.

Figure 13-30 shows how to configure rate interval when changing the default value.

Figure 13-30. Configuring Rate Interval Example

```

FTOS#show interfaces
TenGigabitEthernet 10/0 is down, line protocol is down
Hardware is Dell 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

FTOS(conf)#interface tengigabitethernet 10/0
FTOS(conf-if-te-10/0)#rate-interval 100

FTOS#show interfaces
TenGigabitEthernet 10/0 is down, line protocol is down
Hardware is Dell 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 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

Dynamic Counters

By default, counting for the following four applications is enabled:

- IPFLOW
- IPACL
- L2ACL
- L2FIB

For the remaining applications, FTOS automatically turns on counting when you enable the application and is turned off when you disable the application. Note that if you enable more than four counter-dependent applications 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

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 EXEC Privilege mode:

Command Syntax	Command Mode	Purpose
<code>clear counters [interface] [vrrp [vrid] learning-limit]</code>	EXEC Privilege	<p>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.</p> <ul style="list-style-type: none"> • (OPTIONAL) Enter the following interface keywords and slot/port or number information: • For a Loopback interface, enter the keyword <code>loopback</code> followed by a number from 0 to 16383. • For a Port Channel interface, enter the keyword <code>port-channel</code> followed by a number from 1 to 128. • For a 10-Gigabit Ethernet interface, enter the keyword TenGigabitEthernet followed by the slot/port information. • For a 40-Gigabit Ethernet interface, enter the keyword fortyGigE followed by the slot/port information. • For a VLAN, enter the keyword <code>vlan</code> followed by a number from 1 to 4094 • (OPTIONAL) Enter the keyword <code>vrrp</code> to clear statistics for all VRRP groups configured. Enter a number from 1 to 255 as the <i>vrid</i>. • (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 13-31).

Figure 13-31. Clearing an Interface

```
FTOS#clear counters tengig 0/0
Clear counters on TenGigabitEthernet 0/0 [confirm]
FTOS#
```

IPv4 Routing

The Dell Force10 operating software (FTOS) 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.

- [IP Addresses](#)
- [Directed Broadcast](#)
- [Resolution of Host Names](#)
- [Address Resolution Protocol \(ARP\)](#)
- [Internet Control Message Protocol \(ICMP\)](#)
- [UDP Helper](#)

Table 14-1 lists the defaults for the IP addressing features described in this chapter.

Table 14-1. 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. It also supports classful routing and variable length subnet masks (VLSM). With VLSM, you can configure one network with different masks. Supernetting, which increases the number of subnets, is also supported. Subnetting occurs 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, 10.214.87.131 is represented as:

```
00001010110101100101011110000011
```

For more information about IP addressing, refer to [RFC 791](#), *Internet Protocol*.

Implementation Information

In FTOS, you can configure any IP address as a static route except IP addresses already assigned to interfaces.



Note: FTOS versions 7.7.1.0 and later support 31-bit subnet masks (/31, or 255.255.255.254) as defined by RFC 3021. This feature allows you to save two more IP addresses on point-to-point links than 30-bit masks. FTOS supports RFC 3021 with ARP.

Configuration Task List for IP Addresses

The following list includes the configuration tasks for IP addresses:

- [Assign IP Addresses to an Interface](#) (mandatory)
- [Configure Static Routes](#) (optional)
- [Configure Static Routes for the Management Interface](#) (optional)

For a complete listing of all commands related to IP addressing, refer to *FTOS Command Line Interface Reference Guide*.

Assign IP Addresses to an Interface

Assign primary and secondary IP addresses to physical or logical interfaces (for example, a virtual local area network [VLAN] or port channel) to enable IP communication between the MXL Switch and hosts connected to that interface. In FTOS, you can assign one primary address and up to 255 secondary IP addresses to each interface.

To assign an IP address to an interface, follow these steps, starting in CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	interface <i>interface</i>	CONFIGURATION	<p>Enter the keyword <code>interface</code> followed by the type of interface and slot/port information:</p> <ul style="list-style-type: none"> • For a Loopback interface, enter the keyword <code>loopback</code> followed by a number from 0 to 16383. • For the Management interface, enter the keyword <code>ManagementEthernet</code> followed by the slot/port information. The slot range is 0/0 and the port range is 0/0. • For a port channel interface, enter the keyword <code>port-channel</code> followed by a number from 1 to 128. • For a 10-Gigabit Ethernet interface, enter the keyword <code>TenGigabitEthernet</code> followed by the slot/port information. • For a VLAN interface, enter the keyword <code>vlan</code> followed by a number from 1 to 4094. • For a 40-Gigabit Ethernet interface, enter the keyword <code>fortyGigE</code> followed by the slot/port information.
2	no shutdown	INTERFACE	Enable the interface.

Step	Command Syntax	Command Mode	Purpose
3	<code>ip address <i>ip-address</i> <i>mask</i> [<i>secondary</i>]</code>	INTERFACE	Configure a primary IP address and mask on the interface. <ul style="list-style-type: none"> • <i>ip-address mask</i>: 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.

To view the configuration, use the `show config` command (Figure 13-1) in INTERFACE mode or the `show ip interface` command in EXEC privilege mode (Figure 13-2).

Figure 14-1. show config Command Example

```
FTOS(conf-if-te-0/16)#show conf
!
interface TenGigabitEthernet 0/16
  no ip address
  shutdown
FTOS(conf-if-te-0/16)#
```

Figure 14-2. show ip interface Command Example

```
FTOS#show ip interface tengig 0/16
TenGigabitEthernet 0/16 is down, line protocol is down
Internet address is not set
IP MTU is 1500 bytes
Directed broadcast forwarding is disabled
Proxy ARP is enabled
Split Horizon is enabled
Poison Reverse is disabled
ICMP redirects are not sent
ICMP unreachable are not sent
FTOS#
```

Configure Static Routes

A static route is an IP address that is manually configured and not learned by a routing protocol, such as open shortest path first (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 CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
<pre>ip route <i>ip-address mask</i> [<i>ip-address</i> <i>interface</i> [<i>ip-address</i>]] [<i>distance</i>] [permanent] [tag <i>tag-value</i>]</pre>	CONFIGURATION	<p>Configure a static IP address. Use the following required and optional parameters:</p> <ul style="list-style-type: none"> • <i>ip-address</i>: Enter an address in dotted decimal format (A.B.C.D). • <i>mask</i>: Enter a mask in slash prefix-length format (/X). • <i>interface</i>: Enter an interface type followed by slot/port information. • <i>distance</i> range: 1 to 255 (optional). • <i>permanent</i>: Keep the static route in the routing table (if <i>interface</i> option is used) even if the interface with the route is disabled. (optional) • <i>tag tag-value</i> 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 (Figure 14-3).

Figure 14-3. show ip route static Command Example (partial)

```
FTOS#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--
```

FTOS installs a next hop that is on the directly connected subnet of the current IP address on the interface (for example, if interface tengig 0/0 is on 172.31.5.0 subnet, FTOS installs the static route).

FTOS 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 tengig 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 an interface goes down, FTOS withdraws the route.

- When an interface comes up, FTOS re-installs the route.
- When a recursive resolution is “broken,” FTOS withdraws the route.
- When a 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 CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
management route <i>ip-address mask</i> { <i>forwarding-router-address</i> ManagementEthernet <i>slot/port</i> }	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 EXEC privilege mode (Figure 14-4).

Figure 14-4. show ip management-route Command Example

```
FTOS#show ip management-route all
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
FTOS#
```

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 INTERFACE mode:

Command Syntax	Command Mode	Purpose
ip directed-broadcast	INTERFACE	Enable directed broadcast.

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

Resolution of Host Names

Domain name service (DNS) maps host names to IP addresses. This feature simplifies commands such 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 you enable the feature, the system resolves only host names entered into the host table with the ip host command.

- [Enable Dynamic Resolution of Host Names](#)
- [Specify Local System Domain and a List of Domains](#)
- [DNS with Traceroute](#)

Enable Dynamic Resolution of Host Names

By default, dynamic resolution of host names (DNS) is disabled.

To enable DNS, use the following commands in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ip domain-lookup	CONFIGURATION	Enable dynamic resolution of host names.
ip name-server <i>ip-address</i> [<i>ip-address2</i> ... <i>ip-address6</i>]	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 ([Figure 14-5](#)).

Figure 14-5. show hosts Command Example

```
FTOS>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
FTOS>
```

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 FTOS cannot resolve the domain, it tries the domain name assigned to the local system. If that does not resolve the partial domain, FTOS searches the list of domains configured

To configure a domain name, use the following command in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
<code>ip domain-name <i>name</i></code>	CONFIGURATION	Enter up to 63 characters to configure one domain name.

To configure a list of domain names, use the following command in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
<code>ip domain-list <i>name</i></code>	CONFIGURATION	Enter up to 63 characters to configure names to complete unqualified host names. Configure this command up to six times to specify a list of possible domain names. FTOS 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, use the following commands in CONFIGURATION mode.

Command Syntax	Command Mode	Purpose
<code>ip domain-lookup</code>	CONFIGURATION	Enable dynamic resolution of host names.
<code>ip name-server <i>ip-address</i> [<i>ip-address2</i> ... <i>ip-address6</i>]</code>	CONFIGURATION	Specify up to six name servers. The order you entered the servers determines the order of their use.

Command Syntax	Command Mode	Purpose
tracert [host ip-address]	CONFIGURATION	<p>When you enter the tracert command without specifying an IP address (Extended Tracert), you are prompted for:</p> <ul style="list-style-type: none"> • 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) • port number (default is 33434). <p>To keep the default setting for those parameters, press the ENTER key.</p>

Figure 14-6 shows an example output of DNS using the tracert command.

Figure 14-6. Tracert Command Example

```

FTOS#tracert www.forcel0networks.com

Translating "www.forcel0networks.com"...domain server (10.11.0.1) [OK]
Type Ctrl-C to abort.

-----

Tracing the route to www.forcel0networks.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.forcel0networks.com (10.11.30.126) 005.000 ms 001.000 ms 001.000 ms
 3  fw-sjc-01.forcel0networks.com (10.11.127.254) 000.000 ms 000.000 ms 000.000 ms
 4  www.forcel0networks.com (10.11.84.18) 000.000 ms 000.000 ms 000.000 ms
FTOS#

```

Address Resolution Protocol (ARP)

FTOS uses two forms of address resolution: ARP and proxy ARP.

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 ARP cache. Dynamically learned addresses are removed after a defined period of time.

For more information about ARP, refer to 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 about 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](#) (optional)
- [Enable Proxy ARP](#) (optional)
- [Clear ARP Cache](#) (optional)
- [ARP Learning via Gratuitous ARP](#)
- [ARP Learning via ARP Request](#)
- [Configurable ARP Retries](#)

For a complete listing of all ARP-related commands, refer to the *FTOS Command Line Reference Guide*.

Configure Static ARP Entries

ARP dynamically maps 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 CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
<code>arp ip-address mac-address interface</code>	CONFIGURATION	Configure an IP address and MAC address mapping for an interface. <ul style="list-style-type: none">• <i>ip-address</i>: IP address in dotted decimal format (A.B.C.D).• <i>mac-address</i>: MAC address in nnnn.nnnn.nnnn format• <i>interface</i>: 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 in EXEC privilege mode (Figure 14-7).

Figure 14-7. show arp static Command Example

```
FTOS#show arp
Protocol      Address          Age(min)  Hardware Address  Interface  VLAN      CPU
-----
Internet     10.11.68.14     94       00:01:e9:45:00:03  Ma 0/0    -         CP
Internet     10.11.209.254   0        00:01:e9:45:00:03  Ma 0/0    -         CP
FTOS#
```

Enable Proxy ARP

By default, proxy ARP is enabled. To disable Proxy ARP, use the `no proxy-arp` command in INTERFACE mode.

To re-enable proxy ARP, use the following command in INTERFACE mode:

Command Syntax	Command Mode	Purpose
<code>ip proxy-arp</code>	INTERFACE	Re-enable proxy ARP.

To view if proxy ARP is enabled on the interface, use the `show config` command in INTERFACE mode. If it is not listed in the `show config` command output, it is enabled. Only non-default 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 EXEC Privilege mode:

Command Syntax	Command Mode	Purpose
<code>clear arp-cache [interface ip ip-address] [no-refresh]</code>	EXEC privilege	<p>Clear the ARP caches for all interfaces or for a specific interface by entering the following information:</p> <ul style="list-style-type: none">• For a port channel interface, enter the keyword port-channel followed by a number from 1 to 128.• 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.• For a 40-Gigabit Ethernet interface, enter the keyword fortyGigE followed by the slot/port information.• 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 <i>interface</i> or <i>ip ip-address</i> to specify which dynamic ARP entries you want to delete. <p>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.</p>

ARP Learning via Gratuitous ARP

Gratuitous ARP can mean an ARP request or reply. In the context of ARP learning via gratuitous ARP on FTOS, the gratuitous ARP is a request. A gratuitous ARP request is an ARP request that is not needed according to the ARP specification, but one that hosts may send to. Gratuitous ARP can:

- detect IP address conflicts
- inform switches of their presence on a port so that packets can be forwarded
- update the ARP table of other nodes on the network in case of an address change

In the request, the host uses its own IP address in the Sender Protocol Address and Target Protocol Address fields. When a gratuitous ARP is received, FTOS installs an ARP entry on the CPU.

To enable ARP learning via gratuitous ARP, use the following command in CONFIGURATION mode:

Task	Command Syntax	Command Mode
Enable ARP learning via gratuitous ARP.	<code>arp learn-enable</code>	CONFIGURATION

ARP Learning via ARP Request

In FTOS versions prior to 8.3.1.0, FTOS learns via ARP requests only if the target IP specified in the packet matches the IP address of the receiving router interface. This is the case when a host is attempting to resolve the gateway address.

If the target IP does not match the incoming interface, the packet is dropped. If there is an existing entry for the requesting host, it is updated (Figure 14-8).

Beginning with FTOS version 8.3.1.0, when you enable ARP learning via gratuitous ARP, the system installs a new ARP entry, or updates an existing entry for all received ARP requests (Figure 14-9).

Figure 14-8. Learning via Gratuitous ARP not Enabled


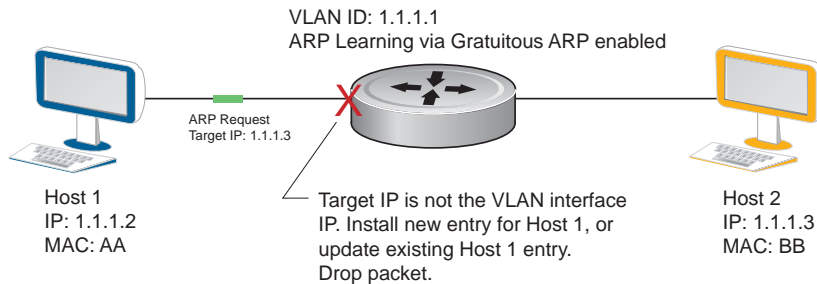


Figure 14-9. Learning via Gratuitous ARP Enabled



Whether you enable or disable ARP learning via gratuitous ARP, the system does not look up the Target IP. It only updates the ARP entry for the Layer 3 interface with the source IP of the request.

Configurable ARP Retries

Beginning with FTOS version 8.3.1.0, the number of ARP retries is configurable.

The default backoff interval remains at 20 seconds. On the MXL switch platform, with FTOS version 8.3.8.0 and later, the time between ARP re-send is configurable. This timer is an exponential backoff timer. Over the specified period, the time between ARP requests increases. This reduces the potential for the system to slow down while waiting for a multitude of ARP responses.

Task	Command Syntax	Command Mode
Set the number of ARP retries.	arp retries <i>number</i> Default: 5 Range: 1-20	CONFIGURATION
Set the an exponential timer for resending unresolved ARPs.	arp backoff-time Default: 30 Range: 1 to 3600	CONFIGURATION
Display all ARP entries learned via gratuitous ARP.	show arp retries	EXEC Privilege

Internet Control Message Protocol (ICMP)

For diagnostics, ICMP provides 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 271](#)

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

Enable ICMP Unreachable Messages

By default, ICMP unreachable messages are disabled. When you enable them, ICMP unreachable messages are created and sent out to all interfaces. To disable ICMP unreachable messages, use the no ip unreachable command.

To reenble the creation of ICMP unreachable messages on the interface, use the following command in 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 INTERFACE mode. If it is not listed in the show config command output, it is enabled. Only non-default information is displayed in the show config command output.

UDP Helper

UDP helper allows you to direct the forwarding IP/UDP broadcast traffic by creating special broadcast addresses and rewriting the destination IP address of packets to match those addresses. Configurations using this feature are described in [Configurations Using UDP Helper](#).

Configuring UDP Helper

Configuring FTOS to direct UDP broadcast is a two-step process:

1. Enable UDP helper and specify the UDP ports for which traffic is forwarded ([Enabling UDP Helper](#)).

Important Points to Remember

- The existing ip directed broadcast command is rendered meaningless if you enable UDP helper on the same interface.
- The broadcast traffic rate must not exceed 200 packets per second when you enable UDP helper.
- You may specify a maximum of 16 UDP ports.
- UDP helper is compatible with IP helper (ip helper-address):
 - UDP broadcast traffic with port number 67 or 68 are unicast to the dynamic host configuration protocol (DHCP) server per the ip helper-address configuration whether or not the UDP port list contains those ports.
 - If the UDP port list contains ports 67 or 68, UDP broadcast traffic is forwarded on those ports.

Enabling UDP Helper

To enable UDP helper, use the ip udp-helper udp-ports command ([Figure 14-10](#)).

Figure 14-10. Enabling UDP Helper

```
FTOS(conf-if-te-1/1)#ip udp-helper udp-port 1000
FTOS(conf-if-te-1/1)#show config
!
interface TenGigabitEthernet 1/1
 ip address 2.1.1.1/24
 ip udp-helper udp-port 1000
 no shutdown
```

To view the interfaces and ports on which UDP helper is enabled, use the show ip udp-helper command from EXEC Privilege mode ([Figure 14-11](#)).

Figure 14-11. Viewing the UDP Broadcast Configuration

```
FTOS#show ip udp-helper
-----
Port          UDP port list
-----
TenGig 1/1    1000
```

Configurations Using UDP Helper

When you enable UDP helper and the destination IP address of an incoming packet is a broadcast address, FTOS suppresses the destination address of the packet. The following sections describe various configurations that employ UDP helper to direct broadcasts.

- [UDP Helper with Broadcast-All Addresses](#)
- [UDP Helper with Subnet Broadcast Addresses](#)
- [UDP Helper with Configured Broadcast Addresses](#)
- [UDP Helper with No Configured Broadcast Addresses](#)

UDP Helper with Broadcast-All Addresses

When the destination IP address of an incoming packet is the IP broadcast address, FTOS rewrites the address to match the configured broadcast address.


In [Figure 14-12](#):

1. Packet 1 is dropped at ingress if no UDP helper address is configured.
2. If you enable UDP helper (using the `ip udp-helper udp-port` command), and the UDP destination port of the packet matches the UDP port configured, the system changes the destination address to the configured broadcast 1.1.255.255 and routes the packet to VLANs 100 and 101. If an IP broadcast address is not configured (using the `ip udp-broadcast-address` command) on VLANs 100 or 101, the packet is forwarded using the original destination IP address 255.255.255.255.

Packet 2, sent from a host on VLAN 101 has a broadcast MAC address and IP address. In this case:

1. It is flooded on VLAN 101 without changing the destination address because the forwarding process is Layer 2.
2. If you enable UDP helper, the system changes the destination IP address to the configured broadcast address 1.1.255.255 and forwards the packet to VLAN 100.
3. Packet 2 is also forwarded to the ingress interface with an unchanged destination address because it does not have broadcast address configured.

Figure 14-12. UDP helper with All Broadcast Addresses



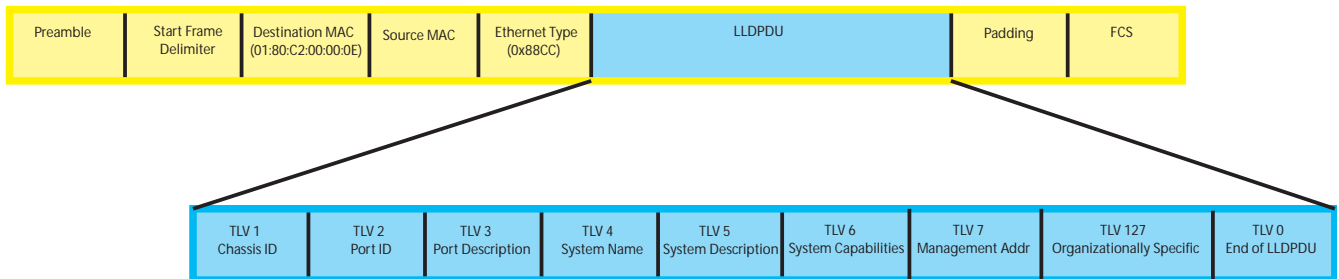
UDP Helper with Subnet Broadcast Addresses

When the destination IP address of an incoming packet matches the subnet broadcast address of any interface, the system changes the address to the configured broadcast address and sends it to a matching interface.

In [Figure 14-13](#), Packet 1 has the destination IP address 1.1.1.255, which matches the subnet broadcast address of VLAN 101. If you configured UDP helper and the packet matches the specified UDP port, the system changes the address to the configured IP broadcast address and floods the packet on VLAN 101.

Packet 2 is sent from host on VLAN 101. It has a broadcast MAC address and a destination IP address of 1.1.1.255. In this case, it is flooded on VLAN 101 in its original condition as the forwarding process is Layer 2.

Figure 14-13. UDP helper with Subnet Broadcast Addresses




UDP Helper with Configured Broadcast Addresses

Incoming packets with a destination IP address matching the configured broadcast address of any interface are forwarded to the matching interfaces.

In [Figure 14-14](#), Packet 1 has a destination IP address that matches the configured broadcast address of VLAN 100 and 101. If you enabled UDP helper and the UDP port number matches, the packet is flooded on both VLANs with an unchanged destination address.

Packet 2 is sent from a host on VLAN 101. It has broadcast MAC address and a destination IP address that matches the configured broadcast address on VLAN 101. In this case, Packet 2 is flooded on VLAN 101 with the destination address unchanged because the forwarding process is Layer 2. If you enabled UDP helper, the packet is flooded on VLAN 100 as well.

Figure 14-14. UDP Helper with Configured Broadcast Addresses



UDP Helper with No Configured Broadcast Addresses

If the incoming packet has a broadcast destination IP address, the unaltered packet is routed to all Layer 3 interfaces. If the incoming packet has a destination IP address that matches the subnet broadcast address of any interface, the unaltered packet is routed to the matching interfaces.

Troubleshooting UDP Helper

To display debugging information, use the `debug ip udp-helper` command (Figure 14-15).

Figure 14-15. Debugging UDP Broadcast

```
FTOS(conf)# debug ip udp-helper
01:20:22: Pkt rcvd on TenGig 5/0 with IP DA (0xffffffff) will be sent on TenGig 5/1 TenGig 5/2 Vlan 3
01:44:54: Pkt rcvd on TenGig 7/0 is handed over for DHCP processing.
```

When using the IP helper and UDP helper on the same interface, use the `debug ip dhcp` command (Figure 14-16).

Figure 14-16. Debugging IP Helper with UDP Helper

```
Packet 0.0.0.0:68 -> 255.255.255.255:67 TTL 128

2005-11-05 11:59:35 %RELAY-I-PACKET, BOOTP REQUEST (Unicast) received at interface
172.21.50.193 BOOTP Request, XID = 0x9265f901, secs = 0 hwaddr = 00:02:2D:8D:46:DC, giaddr =
0.0.0.0, hops = 2

2005-11-05 11:59:35 %RELAY-I-BOOTREQUEST, Forwarded BOOTREQUEST for 00:02:2D:8D:46:DC to
137.138.17.6

2005-11-05 11:59:36 %RELAY-I-PACKET, BOOTP REPLY (Unicast) received at interface
194.12.129.98 BOOTP Reply, XID = 0x9265f901, secs = 0 hwaddr = 00:02:2D:8D:46:DC, giaddr =
172.21.50.193, hops = 2

2005-07-05 11:59:36 %RELAY-I-BOOTREPLY, Forwarded BOOTREPLY for 00:02:2D:8D:46:DC to
128.141.128.90 Packet 0.0.0.0:68 -> 255.255.255.255:67 TTL 128
```

iSCSI Optimization

iSCSI optimization is supported on the MXL 10/40GbE Switch.

This chapter describes how to configure internet small computer system interface (iSCSI) optimization, which enables quality-of-service (QoS) treatment for iSCSI traffic. The topics covered in this chapter include:

- [iSCSI Optimization Overview](#)
- [Default iSCSI Optimization Values](#)
- [iSCSI Optimization Prerequisites](#)
- [Configuring iSCSI Optimization](#)
- [Displaying iSCSI Optimization Information](#)

iSCSI Optimization Overview

iSCSI is a TCP/IP-based protocol for establishing and managing connections between IP-based storage devices and initiators in a storage area network (SAN).

iSCSI optimization provides a means of monitoring iSCSI sessions and applying QoS policies on iSCSI traffic. When enabled, iSCSI optimization allows a switch to monitor (snoop) the establishment and termination of iSCSI connections. The switch uses the snooped information to detect iSCSI sessions and connections established through the switch.

iSCSI optimization allows you to reduce deployment time and management complexity in data centers. In a data center network, Dell EqualLogic and Compellent iSCSI storage arrays are connected to a converged Ethernet network using the data center bridging exchange protocol (DCBX) through stacked and/or non-stacked Ethernet switches.

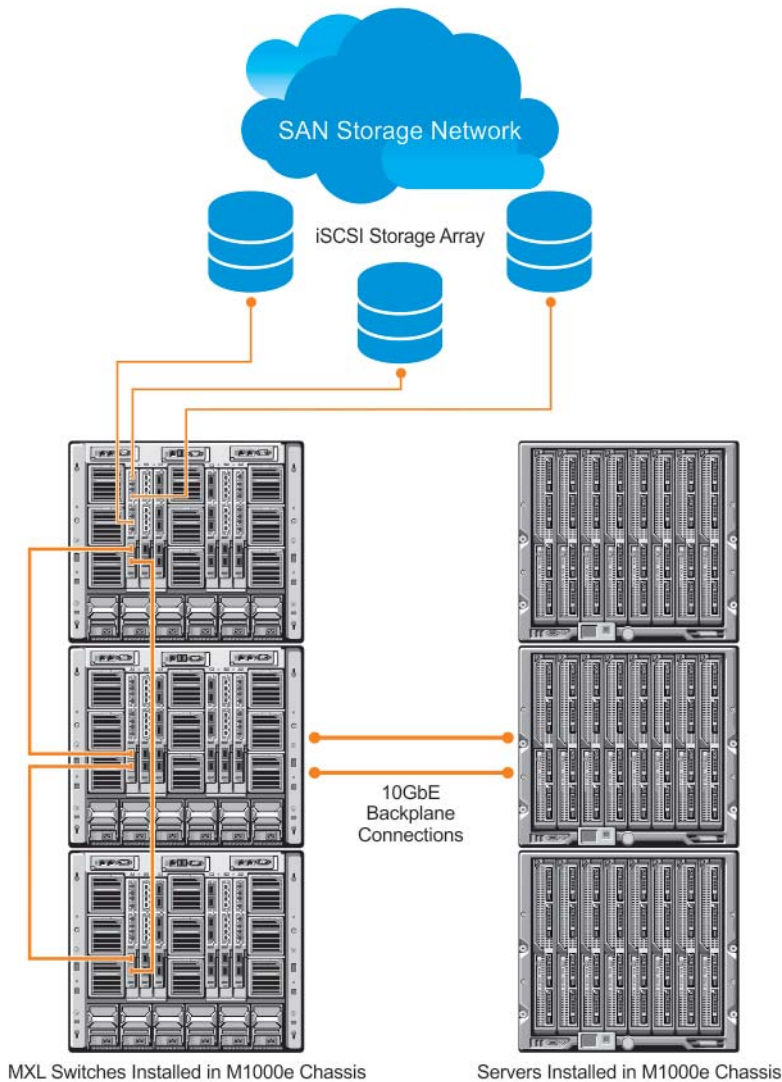
On the MXL Switch, iSCSI optimization functions as follows:

- Auto-detection of EqualLogic storage arrays—The switch detects any active EqualLogic array directly attached to its ports.
- Manual configuration to detect Compellent storage arrays where auto-detection is not supported.
- Automatic configuration of switch ports after detection of storage arrays.
- iSCSI monitoring sessions—The switch monitors and tracks active iSCSI sessions in connections on the switch, including port information and iSCSI session information.

- iSCSI QoS—A user-configured iSCSI class of service (CoS) profile is applied to all iSCSI traffic. Classifier rules are used to direct the iSCSI data traffic to queues that can be given preferential QoS treatment over other data passing through the switch. Preferential treatment helps to avoid session interruptions during times of congestion that would otherwise cause iSCSI packets to be dropped.
- iSCSI DCBX TLVs are supported.

Figure 15-1 shows iSCSI optimization between servers and a storage array in which a stack of three MXL Switches connect installed servers (iSCSI initiators) to a storage array (iSCSI targets) in a SAN network. iSCSI optimization running on the master MXL switch is configured to use dot1p priority-queue assignments to ensure that iSCSI traffic in these sessions receives priority treatment when forwarded on stacked switch hardware.

Figure 15-1. iSCSI Optimization Example



Monitoring iSCSI Traffic Flows

The switch snoops iSCSI session-establishment and termination packets by installing classifier rules that trap iSCSI protocol packets to the CPU for examination. Devices that initiate iSCSI sessions usually use well-known TCP ports 3260 or 860 to contact targets. When you enable iSCSI optimization, by default the switch identifies IP packets to or from these ports as iSCSI traffic.

You can configure the switch to monitor traffic for additional port numbers or a combination of port number and target IP address, and you can remove the well-known port numbers from monitoring.

Application of Quality of Service to iSCSI Traffic Flows

The iSCSI CoS mode is user-configurable and controls whether CoS (dot1p priority) queue assignment and/or packet marking is performed on iSCSI traffic. When you enable iSCSI CoS mode, the CoS policy is applied to iSCSI traffic. When you disable iSCSI CoS mode, iSCSI sessions and connections are still detected and displayed in the status tables, but no CoS policy is applied to iSCSI traffic.

You can configure whether the iSCSI optimization feature uses the VLAN priority or IP DSCP mapping to determine the traffic class queue. By default, iSCSI flows are assigned to dot1p priority 4. Use the CoS dot1p-priority command to map incoming iSCSI traffic on an interface to a dot1p priority-queue other than 4 (refer to [QoS dot1p Traffic Classification and Queue Assignment](#)).

You can configure whether iSCSI frames are re-marked to contain the configured VLAN priority tag or IP DSCP when forwarded through the switch.



Note: On a switch in which a large proportion of traffic is iSCSI, CoS queue assignment may interfere with other network control-plane traffic, such as ARP or LACP. Preferential treatment of iSCSI traffic needs to be balanced against the needs of other critical data in the network.

Information Monitored in iSCSI Traffic Flows

iSCSI optimization examines the following data in packets and uses the data to track the session and create the classifier entries that enable QoS treatment:

- Initiator's IP Address
- Target's IP Address
- ISID (Initiator defined session identifier)
- Initiator's IQN (iSCSI qualified name)
- Target's IQN
- Initiator's TCP Port
- Target's TCP Port

If no iSCSI traffic is detected for a session during a user-configurable aging period, the session data is cleared.

Detection and Autoconfiguration for Dell EqualLogic Arrays

The iSCSI optimization feature includes auto-provisioning support with the ability to detect directly connected Dell EqualLogic storage arrays and automatically reconfigure the switch to enhance storage traffic flows.

The MXL Switch uses the link layer discovery protocol (LLDP) to discover Dell EqualLogic devices on the network. LLDP is enabled by default. For more information about LLDP, refer to [Link Layer Discovery Protocol \(LLDP\)](#).

The following message is displayed the first time a Dell EqualLogic array is detected and describes the configuration changes that are automatically performed:

```
%STKUNIT0-M:CP %IFMGR-5-IFM_ISCSI_AUTO_CONFIG: This switch is being configured for optimal conditions to support iSCSI traffic which will cause some automatic configuration to occur including jumbo frames and flow-control on all ports; no storm control and spanning-tree port fast to be enabled on the port of detection.
```

The following syslog message is generated the first time an EqualLogic array is detected:

```
%STKUNIT0-M:CP %LLDP-5-LLDP_EQL_DETECTED: EqualLogic Storage Array detected on interface Te 1/43
```

- At the first detection of an EqualLogic array, an MTU of 1200 is enabled on all ports and port-channels (if it has not already been enabled).
- Spanning-tree portfast is enabled on the interface identified by LLDP if the port is in L2 mode.
- Unicast storm control is disabled on the interface identified by LLDP.

Detection and Port Configuration for Dell Compellent Arrays

MXL Switches support the `iscsi profile-compellent` command to configure a port connected to a Dell Compellent storage array. The command configures a port for the best iSCSI traffic conditions and must be entered in INTERFACE Configuration mode.

The following message is displayed the first time you use the `iscsi profile-compellent` command to configure a port connected to a Dell Compellent storage array and describes the configuration changes that are automatically performed:

```
%STKUNIT0-M:CP %IFMGR-5-IFM_ISCSI_AUTO_CONFIG: This switch is being configured for optimal conditions to support iSCSI traffic which will cause some automatic configuration to occur including jumbo frames and flow-control on all ports; no storm control and spanning-tree port fast to be enabled on the port of detection.
```

After you execute the `iscsi profile-compellent` command, the following actions occur:

- Jumbo frame size is set to 1200 for all interfaces on all ports and port-channels, if it is not already enabled.
- Spanning-tree portfast is enabled on the interface identified by LLDP if the port is in L2 mode.
- Unicast storm control is disabled on the interface identified by LLDP.

You must enter the `iscsi profile-compellent` command in interface configuration mode; for example:

```
FTOS(conf-if-te-0/50# iscsi profile-compellent
```

Enabling and Disabling iSCSI Optimization



Note: iSCSI optimization is enabled by default.

When you enable iSCSI on the switch, the following actions occur:

- Link-level flow control is globally enabled, if it is not already enabled, and PFC is disabled.
- iSCSI session snooping is enabled.
- iSCSI LLDP monitoring starts to automatically detect EqualLogic arrays.

The following message is displayed when you enable iSCSI on a switch and describes the configuration changes that are automatically performed:

```
%STKUNIT0-M:CP %IFMGR-5-IFM_ISCSI_ENABLE: iSCSI has been enabled causing flow control to be enabled on all interfaces. EQL detection and enabling iscsi profile-compellent on an interface may cause some automatic configurations to occur like jumbo frames on all ports and no storm control and spanning tree port-fast on the port of detection.
```

You can reconfigure any of the auto-provisioned configuration settings that result when you enable iSCSI on a switch.

When you disable the iSCSI feature, iSCSI resources are released and the detection of EqualLogic arrays using LLDP is disabled. Disabling iSCSI does not remove the MTU, flow control, portfast, or storm control configuration applied as a result of enabling iSCSI.



Note: When you enable the iSCSI feature using the `iscsi enable` command, flow control settings are set to **rx on tx off** on all interfaces.

Default iSCSI Optimization Values

Table 15-1 shows the default values for the iSCSI optimization feature.

Table 15-1. iSCSI Optimization: Default Parameters

Parameter	Default Value
iSCSI Optimization global setting	Enabled
iSCSI CoS mode (802.1p priority queue mapping)	Enabled: dot1p priority 4 without remark setting
iSCSI CoS Packet classification	iSCSI packets are classified by VLAN instead of by DSCP values.
VLAN priority tag	iSCSI flows are assigned by default to dot1p priority 4 without remark setting.
DSCP	None: user-configurable.
Remark	Not configured.
iSCSI session aging time	10 minutes
iSCSI optimization target ports	iSCSI well-known ports 3260 and 860 are configured as default (with no IP address or name) but can be removed as any other configured target.

iSCSI Optimization Prerequisites

- iSCSI optimization requires that you enable LLDP on the switch. LLDP is enabled by default (refer to [Link Layer Discovery Protocol \(LLDP\)](#)).
- iSCSI optimization requires two ingress ACL groups to be configured. iSCSI is allocated two ACL groups by default (refer to [CAM Allocation](#)).

Configuring iSCSI Optimization

To configure iSCSI optimization on a switch, follow these steps:

Step	Task	Command	Command Mode
1	Globally enable iSCSI optimization. Default: Enabled.	[no] iscsi enable	CONFIGURATION
2	Configure the iSCSI target ports and optionally the IP addresses on which iSCSI communication will be monitored, where: <ul style="list-style-type: none"> <i>tcp-port-n</i> is the TCP port number or a list of TCP port numbers on which the iSCSI target listens to requests. You can configure up to 16 target TCP ports on the switch in one command or multiple commands. Default: 860, 3260. Separate port numbers with a comma. <i>ip-address</i> specifies the IP address of the iSCSI target. When you enter the no form of the command, and the TCP port to be deleted is one bound to a specific IP address, the IP address value must be included in the command. 	[no] iscsi target port <i>tcp-port-1</i> [<i>tcp-port-2...tcp-port-16</i>] [address <i>ip-address</i>]	CONFIGURATION
3	Set the QoS policy that will be applied to iSCSI flows, where: <ul style="list-style-type: none"> <i>enable</i> enables the application of preferential QoS treatment to iSCSI traffic so that iSCSI packets are scheduled in the switch with a dot1p priority 4 regardless of the VLAN priority tag in the packet. Default: iSCSI packets are handled with dot1p priority 4 without remark. <i>disable</i> disables the application of preferential QoS treatment to iSCSI frames. <i>dot1p vlan-priority-value</i> specifies the VLAN priority tag assigned to incoming packets in an iSCSI session. Range: 0 to 7. Default: The dot1p value in ingress iSCSI frames is not changed and is used in iSCSI TLV advertisements if the iscsi priority-bits command is not entered (Step 5). <i>dscp dscp-value</i> specifies the DSCP value assigned to incoming packets in an iSCSI session. Range: 0 to 63. Default: The DSCP value in ingress packets is not changed. <i>remark</i> marks incoming iSCSI packets with the configured dot1p or DSCP value when they egress the switch. Default: The dot1p and DSCP values in egress packets are not changed. 	[no] iscsi cos {enable disable dot1p <i>vlan-priority-value</i> [remark] dscp <i>dscp-value</i> [remark]}	CONFIGURATION
4	Set the aging time for iSCSI sessions. Valid values: 5 to 43,200 minutes. Default: 10 minutes.	[no] iscsi aging time <i>time</i>	CONFIGURATION

Step	Task	Command	Command Mode
5	(Optional) Configures DCBX to send iSCSI TLV advertisements. You can configure iSCSI TLVs to be sent either globally or on a specified interface. The interface configuration takes priority over global configuration. Default: Enabled.	[no] advertise dcbx-app-tlv iscsi	CONFIGURATION or INTERFACE
6	(Optional) Configures the priority bitmap to be advertised in iSCSI application TLVs. Default: 4 (0x10 in the bitmap).	[no] iscsi priority-bits	CONFIGURATION
7	(Optional) Enter interface configuration mode to configure the auto-detection of Compellent disk arrays.	interface <i>port-type slot/port</i>	CONFIGURATION
8	(Optional) Configures the autodetection of Compellent arrays on a port. Default: Compellent disk arrays are not detected.	[no] iscsi profile-compellent	INTERFACE

Displaying iSCSI Optimization Information

Use the show commands in [Table 15-2](#) to display information on iSCSI optimization.

Table 15-2. Displaying iSCSI Optimization Information

Command	Output
show iscsi (Figure 15-2)	Displays the currently configured iSCSI settings.
show iscsi sessions (Figure 15-3)	Displays information on active iSCSI sessions on the switch.
show iscsi sessions detailed [session <i>isid</i>] (Figure 15-4)	Displays detailed information on active iSCSI sessions on the switch. To display detailed information on specified iSCSI session, enter the session's iSCSI ID.
show run iscsi	Displays all globally-configured non-default iSCSI settings in the current FTOS session.

Figure 15-2. show iscsi Command Example

```
FTOS# show iscsi

iSCSI is enabled
iSCSI COS : dot1p is 4 no-remark
Session aging time: 10
Maximum number of connections is 256
-----
iSCSI Targets and TCP Ports:
-----
TCP Port      Target IP Address
3260
860
```

Figure 15-3. show iscsi sessions Command Example

```
FTOS# show iscsi sessions
Session 0:
-----
Target: iqn.2001-05.com.equallogic:0-8a0906-0e70c2002-10a0018426a48c94-iom010
Initiator: iqn.1991-05.com.microsoft:win-x918v27yajg
ISID: 400001370000

Session 1:
-----
Target: iqn.2001-05.com.equallogic:0-8a0906-0f60c2002-0360018428d48c94-iom011
Initiator: iqn.1991-05.com.microsoft:win-x918v27yajg
ISID: 400001370000.
```

Figure 15-4. show iscsi sessions detailed Command Example

```
FTOS# show iscsi sessions detailed
Session 0      :
-----
Target:iqn.2010-11.com.ixia:ixload:iscsi-TG1
Initiator:iqn.2010-11.com.ixia.ixload:initiator-iscsi-2c
Up Time:00:00:01:28(DD:HH:MM:SS)
Time for aging out:00:00:09:34(DD:HH:MM:SS)
ISID:806978696102
Initiator      Initiator      Target      Target      Connection
IP Address     TCP Port      IP Address  TCPPort     ID
10.10.0.44     33345        10.10.0.101 3260        0
Session 1      :
-----
Target:iqn.2010-11.com.ixia:ixload:iscsi-TG1
Initiator:iqn.2010-11.com.ixia.ixload:initiator-iscsi-35
Up Time:00:00:01:22(DD:HH:MM:SS)
Time for aging out:00:00:09:31(DD:HH:MM:SS)
ISID:806978696102
Initiator      Initiator      Target      Target      Connection
IP Address     TCP Port      IP Address  TCPPort     ID
10.10.0.53     33432        10.10.0.101 3260        0
```


Link Aggregation Control Protocol (LACP)

The major sections in this chapter include:

- [Introduction to Dynamic LAGs and LACP](#)
- [LACP Configuration Tasks](#)
- [Shared LAG State Tracking](#)
- [LACP Basic Configuration Example](#)

Introduction to Dynamic LAGs and LACP

A link aggregation group (LAG), referred to as a port channel by the Dell Force10 operating software (FTOS), provides both load-sharing and port redundancy across stack units. You can enable LAGs as static or dynamic. The benefits and constraints are basically the same, as described in [Port Channel Interfaces in Interfaces](#).

The unique benefit of a dynamic LAG is that its ports can toggle between participating in the LAG or acting as dedicated ports, whereas ports in a static LAG must be specifically removed from the LAG in order to act alone.

FTOS uses LACP to create dynamic LAGs. LACP provides a standardized means of exchanging information between two systems (also called partner systems) and automatically establishes the LAG between the systems. LACP permits the exchange of messages on a link to allow their LACP instances to:

- Reach agreement on the identity of the LAG to which the link belongs.
- Move the link to that LAG.
- 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 protocol data unit (PDUs) across LAN Ethernet links. The protocol packets are only exchanged between ports that are configured as LACP capable.

Important Points to Remember

- LACP allows you to add members to a port channel (LAG) as long as it has no static members. Conversely, if the LAG already contains a statically defined member (channel-member command), the port-channel mode command is not permitted.
- A static LAG cannot be created if a dynamic LAG using the selected number already exists.
- No dual membership in static and dynamic LAGs:
 - If a physical interface is a part of a static LAG, the port-channel-protocol lacp command is rejected on that interface.
 - If a physical interface is a part of a dynamic LAG, it cannot be added as a member of a static LAG. The channel-member tengigabitethernet *x/y* command is rejected in the static LAG interface for that physical interface.
- You can create a dynamic LAG with any type of configuration.
- There is a difference between the shutdown command and the no interface port-channel command:
 - The shutdown command on LAG “xyz” disables the LAG and retains the user commands. However, the system does not allow the channel number “xyz” to be statically created.
 - The no interface port-channel *channel-number* command deletes the specified LAG, including a dynamically created LAG. This command causes all LACP-specific commands on the member interfaces to be removed. The interfaces are restored to a state that is ready to be configured.

Note: There is no configuration on the interface because that condition is required for an interface to be part of a LAG.
- You can configure link dampening on individual members of a LAG. For more information, refer to [MTU Size on an Interface](#).

LACP Modes

FTOS provides the following three modes for configuration of LACP:

- **Off**—In this state, an interface is not capable of being part of a dynamic LAG. LACP does not run on any port that is configured to be in this state.
- **Active**—In this state, the interface is said to be in the “active negotiating state.” LACP runs on any link that is configured to be in this state. A port in Active state also automatically initiates negotiations with other ports by initiating LACP packets.
- **Passive**—In this state, the interface is not in an active negotiating state, but LACP runs on the link. A port in Passive state also responds to negotiation requests (from ports in Active state). Ports in Passive state respond to LACP packets.

FTOS supports LAGs in the following cases:

- A port in Active state can set up a port channel (LAG) with another port in Active state.
- A port in Active state can set up a LAG with another port in Passive state.
- A port in Passive state cannot set up a LAG with another port in Passive state.

LACP Configuration Commands

If you configure aggregated ports with compatible LACP modes (Off, Active, Passive), LACP can automatically link them, as defined in IEEE 802.3, Section 43. The following commands configure LACP:

Command Syntax	Command Mode	Purpose
[no] lacp system-priority <i>priority-value</i>	CONFIGURATION	Configure the system priority. Range: 1– 65535 (the higher the number, the lower the priority) Default: 32768
[no] port-channel-protocol lacp	INTERFACE	Enable or disable LACP on any LAN port: <ul style="list-style-type: none">• Default is LACP disabled• This command creates a new context.
[no] port-channel <i>number</i> mode [active passive off]	LACP	Configure LACP mode. <ul style="list-style-type: none">• Default is LACP active• <i>number</i> cannot statically contain any links
[no] lacp port-priority <i>priority-value</i>	LACP	Configure port priority. <ul style="list-style-type: none">• Ranges: 1 – 65535 (the higher the number, the lower the priority)• Default: 32768

LACP Configuration Tasks

The tasks covered in this section are:

- [Create a LAG](#)
- [Configure the LAG Interfaces as Dynamic](#)
- [Set the LACP Long Timeout](#)
- [Monitor and Debugging LACP](#)
- [Configure Shared LAG State Tracking](#)

Create a LAG

To create a dynamic port channel (LAG), define the LAG and then the LAG interfaces. Use the interface port-channel and switchport commands ([Figure 16-1](#)), which uses the example of LAG 32:

Figure 16-1. Placing a LAG into the Default VLAN

```
FTOS(config)#interface port-channel 32
FTOS(config-if-po-32)#no shutdown
FTOS(config-if-po-32)#switchport
```

The LAG is in the default VLAN. To place the LAG into a non-default VLAN, use the tagged command on the LAG (Figure 16-2):

Figure 16-2. Placing a LAG into a Non-default VLAN

```
FTOS(conf)#interface vlan 10
FTOS(conf-if-vl-10)#tagged port-channel 32
```

Configure the LAG Interfaces as Dynamic

After creating a LAG, to configure the dynamic LAG interfaces, use the port-channel-protocol lacp command. Figure 16-3 shows ports 3/15, 3/16, 4/15, and 4/16 added to LAG 32 in LACP mode.

Figure 16-3. Creating a Dynamic LAG Example

```
FTOS(conf)#interface TenGigabitEthernet 3/15
FTOS(conf-if-te-3/15)#no shutdown
FTOS(conf-if-te-3/15)#port-channel-protocol lacp
FTOS(conf-if-te-3/15-lacp)#port-channel 32 mode active
...
FTOS(conf)#interface TenGigabitEthernet 3/16
FTOS(conf-if-te-3/16)#no shutdown
FTOS(conf-if-te-3/16)#port-channel-protocol lacp
FTOS(conf-if-te-3/16-lacp)#port-channel 32 mode active
...
FTOS(conf)#interface TenGigabitEthernet 4/15
FTOS(conf-if-te-4/15)#no shutdown
FTOS(conf-if-te-4/15)#port-channel-protocol lacp
FTOS(conf-if-te-4/15-lacp)#port-channel 32 mode active
...
FTOS(conf)#interface TenGigabitEthernet 4/16
FTOS(conf-if-te-4/16)#no shutdown
FTOS(conf-if-te-4/16)#port-channel-protocol lacp
FTOS(conf-if-te-4/16-lacp)#port-channel 32 mode active
```

The port-channel 32 mode active command in Figure 16-3 may be successfully issued as long as there is no existing static channel-member configuration in LAG 32.

Set the LACP Long Timeout

PDUs are exchanged between port channel (LAG) interfaces to maintain LACP sessions. PDUs are transmitted at either a slow or fast transmission rate, depending on the LACP timeout value. The timeout value is the amount of time that a LAG interface waits for a PDU from the remote system before bringing the LACP session down. The default timeout value is 1 second but you can configure it to be 30 seconds. Invoking the longer timeout might prevent the LAG from flapping if the remote system is up but temporarily unable to transmit PDUs due to a system interruption.



Note: The 30-second timeout is available for dynamic LAG interfaces only. You can use the lacp long-timeout command for static LAGs, but it has no effect.

To configure the LACP long timeout, follow the step below.

Step	Task	Command Syntax	Command Mode
1	Set the LACP timeout value to 30 seconds.	lACP long-timeout	CONFIG-INT-PO

Figure 16-4 shows the **no shutdown** command.

Figure 16-4. Invoking the LACP Long Timeout

```

FTOS(conf)# interface port-channel 32
FTOS(conf-if-po-32)#no shutdown
FTOS(conf-if-po-32)#switchport
FTOS(conf-if-po-32)#lACP long-timeout
FTOS(conf-if-po-32)#end
FTOS# 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 TenGig 10/6 is enabled, LACP is enabled and mode is lACP
Actor Admin: State ADEHJLMP Key 1 Priority 128

```



Note: To view PDU exchanges and the timeout value, use the `debug lACP` command. For more information, refer to [Monitor and Debugging LACP](#).

Monitor and Debugging LACP

The system log (syslog) records faulty LACP actions.

To debug LACP, use the following command:


Command Syntax	Command Mode	Purpose
[no] debug lACP [config events pdu [in out [interface [in out]]]]	EXEC	Debug LACP, including configuration and events.

Shared LAG State Tracking

Shared LAG state tracking provides the flexibility to bring down a port channel (LAG) based on the operational state of another LAG. At any time, only two LAGs can be a part of a group such that the fate (status) of one LAG depends on the other LAG.

In [Figure 16-5](#), line-rate traffic from R1 destined for R4 follows the lowest-cost route via R2. Traffic is equally distributed between LAGs 1 and 2. If LAG 1 fails, all traffic from R1 to R4 flows across LAG 2 only. This condition over-subscribes the link and packets are dropped.

Figure 16-5. LAGs using ECMP without Shared LAG State Tracking



To avoid packet loss, traffic must be re-directed through the next lowest-cost link (R3 to R4). FTOS has the ability to bring LAG 2 down in the event that LAG 1 fails, so that traffic can be re-directed. This is shared LAG state tracking. To achieve this functionality, you must group LAG 1 and LAG 2 into a single entity, called a failover group.

Configure Shared LAG State Tracking

To configure shared LAG state tracking, you must first configure a failover group. Follow these steps:

Step	Task	Command	Command Mode
1	Enter port-channel failover group mode.	port-channel failover-group	CONFIGURATION
2	Create a failover group and specify the two port-channels that will be members of the group.	group <i>number</i> port-channel <i>number</i> port-channel <i>number</i>	CONFIG-PO-FAILOVER-GRP

In Figure 16-6, LAGs 1 and 2 have been placed into to the same failover group.

Figure 16-6. Configuring Shared LAG State Tracking

```
FTOS#config
FTOS(conf)#port-channel failover-group
FTOS(conf-po-failover-grp)#group 1 port-channel 1 port-channel 2
```


To view the failover group configuration, use the show running-configuration po-failover-group command (Figure 16-7).

Figure 16-7. Viewing Shared LAG State Tracking in the Running-configuration

```
FTOS#show running-config po-failover-group
!
port-channel failover-group
group 1 port-channel 1 port-channel 2
```

In Figure 16-8, LAGs 1 and 2 are members of a failover group. LAG 1 fails and LAG 2 is brought down upon the failure. This effect is logged by Message 1, in which a console message declares both LAGs down at the same time.

Figure 16-8. Shared LAG State Tracking



Message 1 Shared LAG State Tracking Console Message

```
May 16 06:19:37: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Po 1
May 16 06:19:37: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Po 2
```

To view the status of a failover group member, use the `show interface port-channel` command (Figure 16-9).

Figure 16-9. Viewing Status of a Failover Group Member

```
FTOS#show interface Port-channel 2
Port-channel 2 is up, line protocol is down (Failover-group 1 is down)
Hardware address is 00:01:e8:05:e8:4c, Current address is 00:01:e8:05:e8:4c
Interface index is 1107755010
Minimum number of links to bring Port-channel up is 1
Port-channel is part of failover-group 1
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 1000 Mbit
Members in this channel: TenGig 1/17(U)
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:01:28
Queueing strategy: fifo
```



Note: The set of console messages shown in [Message 1](#) appear only if you configure shared LAG state tracking on that router (you can configure the feature on one or both sides of a link). For example, in [Figure 16-8](#), if shared LAG state tracking is configured on R2 only, no messages appears on R4 regarding the state of LAGs in a failover group.

Important Points about Shared LAG State Tracking

- This feature is available for static and dynamic LAGs.
- Only a LAG can be a member of a failover group.
- You can configure shared LAG state tracking on one side of a link or on both sides.
- If a LAG that is part of a failover group is deleted, the failover group is deleted.
- If a LAG moves to the down state due to this feature, its members may still be in the up state.

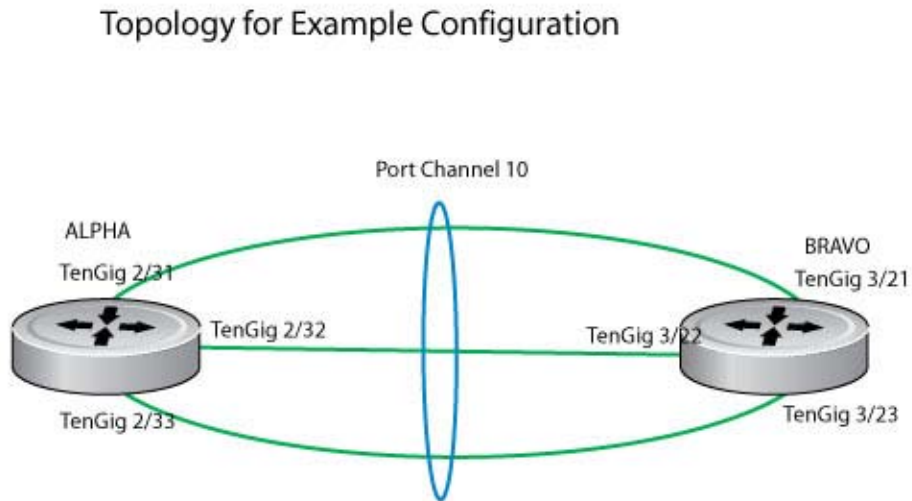
LACP Basic Configuration Example

The screenshots in this section are based on the example topology shown in [Figure 16-10](#). Two routers are named ALPHA and BRAVO, and their hostname prompts reflect those names.

The sections are:

- [Configuring a LAG on ALPHA](#)
- [Summary of the Configuration on ALPHA](#)
- [Summary of the Configuration on BRAVO](#)

Figure 16-10. LACP Sample Topology



Configuring a LAG on ALPHA

Figure 16-11 shows creating a LAG (ALPHA).

Figure 16-11. Creating a LAG on ALPHA

```
Alpha(conf)#interface port-channel 10
Alpha(conf-if-po-10)#no ip address
Alpha(conf-if-po-10)#switchport
Alpha(conf-if-po-10)#no shutdown
Alpha(conf-if-po-10)#show config
!
interface Port-channel 10
  no ip address
  switchport
  no shutdown
!
Alpha(conf-if-po-10)#
```

Figure 16-12 shows the LAG port configuration (ALPHA).

Figure 16-12. Inspecting a LAG Port Configuration on ALPHA

```
Alpha#sh int tengig 0/16
TenGigabitEthernet 0/16 is up, line protocol is down
Hardware is DellForce10Eth, address is 00:1e:c9:bb:02:c2
  Current address is 00:1e:c9:bb:02:c2
Server Port AdminState is Down
Pluggable media not present
Interface index is 38080769
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :tenG145001ec9bb02c2
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 2w0d0h
Queueing strategy: fifo
Input Statistics:
  0 packets,0 bytes
  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:
  0 packets,0 bytes,0 underruns
  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 Unicasts
  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:2w0d0h
```

Shows the status of this physical interface

Figure 16-13 shows inspecting the LAG 10 configuration (ALPHA).

Figure 16-13. Inspecting Configuration of LAG 10 on ALPHA

```
ALPHA#show int port 10
Port-channel 10 is up,line protocol is up
Created by LACP protocol
Hardware address is 00:1e:c9:f1:00:cd, Current address is 00:1e:c9:f1:00:cd
Interface index is 1107755009
Minimum number of links to bring Port-channel up is 1
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :lag1001ec9f100cd
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 30000 Mbit
Members in this channel: Te 3/21(U) Te 3/22(U) Te 3/23(U)
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 6d18h52m
Queueing strategy: fifo
Input Statistics:
 464817 packets, 57464390 bytes
 31962 64-byte pkts, 122 over 64-byte pkts, 432733 over 127-byte pkts
 0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
 464817 Multicasts, 0 Broadcasts
 0 runts, 0 giants, 0 throttles
 0 CRC, 0 overrun, 0 discarded
Output Statistics:
 579948 packets, 72708857 bytes, 0 underruns
 57 64-byte pkts, 118899 over 64-byte pkts, 459925 over 127-byte pkts
 1067 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
 578881 Multicasts, 1067 Broadcasts, 0 Unicasts
 0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
 Input 00.00 Mbits/sec,    2 packets/sec, 0.00% of line-rate
 Output 00.00 Mbits/sec,    2 packets/sec, 0.00% of line-rate
Time since last interface status change: 18:47:40
```

Indicates the MAC address assigned to the LAG.
This does NOT match any of the
physical interface MAC addresses.

Confirms the number of links to bring up
the LAG and that this is a switch
port instead of a router port.

Confirms the total bandwidth for this
LAG and which interfaces are active.

To Verify LAG 10 Status on ALPHA, use the show lacp command (Figure 16-13).

Figure 16-14. show lacp Command Example

```
Alpha#sho lacp 10
Port-channel 10 admin up, oper up, mode lacp
Actor System ID: Priority 32768, Address 0001.e806.953e
Partner System ID: Priority 32768, Address 0001.e809.c24a
Actor Admin Key 10, Oper Key 10, Partner Oper Key 10
LACP LAG 10 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 TenGi 2/31 is enabled, LACP is enabled and mode is lacp
Actor Admin: State ACEHJLMP Key 10 Priority 32768
Oper: State ACEGIKNP Key 10 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 10 Priority 32768

Port TenGi 2/32 is enabled, LACP is enabled and mode is lacp
Actor Admin: State ACEHJLMP Key 10 Priority 32768
Oper: State ACEGIKNP Key 10 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 10 Priority 32768

Port TenGi 2/33 is enabled, LACP is enabled and mode is lacp
Actor Admin: State ACEHJLMP Key 10 Priority 32768
Oper: State ACEGIKNP Key 10 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 10 Priority 32768
Alpha#
```

Shows LAG status

Interfaces participating in the LAG are included here.

Summary of the Configuration on ALPHA

Figure 16-15 shows the summary of the configuration (ALPHA)

Figure 16-15. Summary of the Configuration on ALPHA

```
Alpha(conf-if-po-10)#int tengig 2/31
Alpha(conf-if-te-2/31)#no ip address
Alpha(conf-if-te-2/31)#no switchport
Alpha(conf-if-te-2/31)#shutdown
Alpha(conf-if-te-2/31)#port-channel-protocol lacp
Alpha(conf-if-te-2/31-lacp)#port-channel 10 mode active
Alpha(conf-if-te-2/31-lacp)#no shut
Alpha(conf-if-te-2/31)#show config

!
interface TenGigabitEthernet 2/31
  no ip address
!
  port-channel-protocol LACP
  port-channel 10 mode active
  no shutdown
!
Alpha(conf-if-te-2/31)#

interface Port-channel 10
  no ip address
  switchport
  no shutdown

interface TenGigabitEthernet 2/31
  no ip address
  no switchport
  switchport
  port-channel-protocol LACP
  port-channel 10 mode active
  no shutdown
```

Summary of the Configuration on BRAVO

Figure 16-16 shows the summary of the configuration (BRAVO).

Figure 16-16. Summary of the Configuration on BRAVO

```
Bravo(conf-if-te-3/21)#int port-channel 10
Bravo(conf-if-po-10)#no ip add
Bravo(conf-if-po-10)#switch
Bravo(conf-if-po-10)#no shut
Bravo(conf-if-po-10)#show config
!
interface Port-channel 10
  no ip address
  switchport
  no shutdown
!
Bravo(conf-if-po-10)#exit

Bravo(conf)#int tengig 3/21
Bravo(conf)#no ip address
Bravo(conf)#no switchport
Bravo(conf)#shutdown
Bravo(conf-if-te-3/21)#port-channel-protocol lacp
Bravo(conf-if-te-3/21-lacp)#port-channel 10 mode active
Bravo(conf-if-te-3/21-lacp)#no shut
Bravo(conf-if-te-3/21)#end

!
interface TenGigabitEthernet 3/21
  no ip address
!
  port-channel-protocol LACP
  port-channel 10 mode active
  no shutdown
Bravo(conf-if-te-3/21)#end

int port-channel 10
no ip address
switchport
no shutdown
show config

int tengig 3/21
no ip address
no switchport
shutdown
port-channel-protocol lacp
port-channel 10 mode active
no shut
show config
end
```

To inspect a LAG port on BRAVO, use the show interface command (Figure 16-17).

Figure 16-17. Inspect the LAG Port on BRAVO

```
Bravo#show interfaces tengigabitethernet 3/21
TenGigabitEthernet 3/21 is up, line protocol is up
Port is part of Port-channel 10
Hardware is DellForce10Eth, address is 00:1e:c9:f1:00:cd
Current address is 00:1e:c9:f1:00:cd
Port is present
Pluggable media present, SFP+ type is 10GBASE-CU2M
Medium is MultiRate

Interface index is 113840385
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :tenG434001ec9f100cd

MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 10000 Mbit
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 6d19h0m
Queueing strategy: fifo
Input Statistics:
 250266 packets, 30006383 bytes
 31962 64-byte pkts, 122 over 64-byte pkts, 218182 over 127-byte pkts
 0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
 250266 Multicasts, 0 Broadcasts
 0 runts, 0 giants, 0 throttles
 0 CRC, 0 overrun, 0 discarded
Output Statistics:
 291403 packets, 36650619 bytes, 0 underruns
 57 64-byte pkts, 59540 over 64-byte pkts, 230739 over 127-byte pkts
 1067 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
 290336 Multicasts, 1067 Broadcasts, 0 Unicasts
 0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
 Input 00.00 Mbits/sec, 1 packets/sec, 0.00% of line-rate
 Output 00.00 Mbits/sec, 1 packets/sec, 0.00% of line-rate
Time since last interface status change: 19:01:37
```

Shows the status of this interface.
Also shows it is part of LAG 10.

Shows that this is a Layer 2 port.

Shows the speed of this physical interface.
Also shows it is the Master of the TenGigE link.

To inspect the LAG, use the show interfaces port-channel command (Figure 16-18).

Figure 16-18. show interfaces port-channel Command Example to inspect LAG 10

```
ALPHA#show int port 10
Port-channel 10 is up, line protocol is up
Created by LACP protocol
Hardware address is 00:1e:c9:f1:00:cd, Current address is 00:1e:c9:f1:00:cd
Interface index is 1107755009
Minimum number of links to bring Port-channel up is 1
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :lag1001ec9f100cd
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 30000 Mbit
Members in this channel: Te 3/21(U) Te 3/22(U) Te 3/23(U)
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 6d18h52m
Queueing strategy: fifo
Input Statistics:
 464817 packets, 57464390 bytes
 31962 64-byte pkts, 122 over 64-byte pkts, 432733 over 127-byte pkts
 0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
464817 Multicasts, 0 Broadcasts
 0 runts, 0 giants, 0 throttles
 0 CRC, 0 overrun, 0 discarded
Output Statistics:
579948 packets, 72708857 bytes, 0 underruns
 57 64-byte pkts, 118899 over 64-byte pkts, 459925 over 127-byte pkts
1067 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
578881 Multicasts, 1067 Broadcasts, 0 Unicasts
 0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
 Input 00.00 Mbits/sec,    2 packets/sec, 0.00% of line-rate
 Output 00.00 Mbits/sec,    2 packets/sec, 0.00% of line-rate
Time since last interface status change: 18:47:40
```

Indicates the MAC address assigned to the LAG. This does NOT match any of the physical interface MAC addresses.

Confirms the number of links to bring up the LAG and that this is a switch port instead of a router port.

Confirms the total bandwidth for this LAG and which interfaces are active.

To inspect the LAG status, use the show lacp command (Figure 16-19).

Figure 16-19. show lacp Command Example to Inspect LAG status

```
FTOS#show lacp 10
Port-channel 10 admin up, oper up, mode lacp
Actor System ID: Priority 32768, Address 0001.e809.c24a
Partner System ID: Priority 32768, Address 0001.e806.953e
Actor Admin Key 10, Oper Key 10, Partner Oper Key 10
LACP LAG 10 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 TenGig 3/21 is enabled, LACP is enabled and mode is lacp
Actor Admin: State ACEHJLMP Key 10 Priority 32768
Oper: State ACEGIKNP Key 10 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 10 Priority 32768

Port TenGig 3/22 is enabled, LACP is enabled and mode is lacp
Actor Admin: State ACEHJLMP Key 10 Priority 32768
Oper: State ACEGIKNP Key 10 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 10 Priority 32768

Port TenGig 3/23 is enabled, LACP is enabled and mode is lacp
Actor Admin: State ACEHJLMP Key 10 Priority 32768
Oper: State ACEGIKNP Key 10 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 10 Priority 32768
FTOS#
```

Shows LAG status

Interfaces participating in the LAG are included here.

Layer 2

This chapter describes the following Layer 2 features:

- [Managing the MAC Address Table](#)
- [MAC Learning Limit](#)
- [Network Interface Controller \(NIC\) Teaming](#)

Managing the MAC Address Table

The Dell Force10 operating system (FTOS) provides the following management activities for the MAC address table:

- [Clear the MAC Address Table](#)
- [Set the Aging Time for Dynamic Entries](#)
- [Configure a Static MAC Address](#)
- [Display the MAC Address Table](#)

Clear the MAC Address Table

To clear the MAC address table of dynamic entries, use the following command:

Task	Command Syntax	Command Mode
Clear a MAC address table of dynamic entries. <ul style="list-style-type: none"> • address deletes the specified entry • all deletes all dynamic entries • interface deletes all entries for the specified interface • vlan deletes all entries for the specified VLAN 	<pre>clear mac-address-table dynamic {address all <i>interface</i> vlan}</pre>	EXEC Privilege

Set the Aging Time for Dynamic Entries

Learned MAC addresses are entered in the table as dynamic entries, which means that they are subject to aging. For any dynamic entry, if no packet arrives on the switch with the MAC address as the source or destination address within the timer period, the address is removed from the table. The default aging time is 1800 seconds.

To set the aging time for dynamic entries, use the following commands:

Task	Command Syntax	Command Mode
Disable MAC address aging for all dynamic entries.	<code>mac-address-table aging-time 0</code>	CONFIGURATION
Specify an aging time.	<code>mac-address-table aging-time <i>seconds</i></code> Range: 10-1000000	CONFIGURATION



FTOS Behavior: The time elapsed before the configured MAC aging time expires is not precisely as configured. For example, the VLAN configuration `mac-address-table aging-time 1`, does not remove dynamic entries from the CAM after precisely 1 second. The actual minimum aging time for entries is approximately 5 seconds because this is the default MAC address table scanning interval. Therefore, MAC aging configurations of less than 5 seconds, as in this example, might be ineffective. Configuring `mac-address-table station-move time-interval 500`, solves this limitation. Reducing the scanning interval to the minimum (500 milliseconds), increases the detection speed, which results in FTOS clearing entries closer to the actual desired aging time.

Configure a Static MAC Address

A static entry is one that is not subject to aging. Static entries must be entered manually. To configure a static MAC address, use the following command:

Task	Command Syntax	Command Mode
Create a static MAC address entry in the MAC address table.	<code>mac-address-table static</code>	CONFIGURATION

Display the MAC Address Table

To display the contents of the MAC address table, use the following command:

Task	Command Syntax	Command Mode
Display the contents of the MAC address table. <ul style="list-style-type: none"> • <code>address</code> displays the specified entry. • <code>aging-time</code> displays the configured aging-time. • <code>count</code> displays the number of dynamic and static entries for all VLANs, and the total number of entries. • <code>dynamic</code> displays only dynamic entries • <code>interface</code> displays only entries for the specified interface. • <code>static</code> displays only static entries. • <code>vlan</code> displays only entries for the specified VLAN. 	<code>show mac-address-table [address aging-time [vlan <i>vlan-id</i>] count dynamic interface static vlan]</code>	EXEC Privilege

MAC Learning Limit

This section describes the following:

- [MAC Learning Limit Dynamic](#)
- [MAC Learning Limit Station-Move](#)
- [Learning Limit Violation Actions](#)
- [Station Move Violation Actions](#)
- [Recovering from Learning Limit and Station Move Violations](#)

The MAC address learning limit is a method of port security on Layer 2 port-channel and physical interfaces, and virtual local area networks (VLANs). It allows you to set an upper limit on the number of MAC addresses that are learned on an interface/VLAN. After the limit is reached, the system drops all traffic from a device with an unlearned MAC address.



FTOS Behavior: When configuring the MAC learning limit on a port, the configuration is accepted (becomes part of the running-config and show mac learning-limit interface) before the system verifies that sufficient content addressable memory (CAM) space exists. If the CAM check fails, a message is displayed:

```
%E90MH:5 %ACL_AGENT-2-ACL_AGENT_LIST_ERROR: Unable to apply access-list  
Mac-Limit on TenGigabitEthernet 5/84
```

In this case, the configuration is still present in the running-config and the show output. Remove the configuration before re-applying a MAC learning limit with a lower value. Also, ensure that you can view the syslog message on your session.

To set a MAC learning limit on an interface, use the following command:

Task	Command Syntax	Command Mode
Specify the number of MAC addresses that the system can learn off a Layer 2 interface.	mac learning-limit <i>address_limit</i>	INTERFACE

Three options are available with the mac learning-limit command: dynamic, no-station-move, and station-move.



Note: A simple network management protocol (SNMP) trap is available for mac learning-limit station-move. No other SNMP traps are available for the MAC learning limit, including limit violations.

MAC Learning Limit Dynamic

The MAC address table is stored on the Layer 2 forwarding information base (FIB) region of the CAM. The Layer 2 FIB region allocates space for static MAC address entries and dynamic MAC address entries. When you enable MAC learning limit, entries created on this port are static by default. When you configure the dynamic option, learned MAC addresses are stored in the dynamic region and are subject to aging. Entries created before this option is set are not affected.



FTOS Behavior: If you do not configure the dynamic option, the MXL Switch does not detect station moves in which a MAC address learnt off of a MAC-limited port is learnt on another port on same stack unit or different stack.

MAC Learning Limit Station-Move

The station-move option allows a MAC address already in the table to be learned off of another interface. For example, if you disconnect a network device from one interface and reconnect it to another interface, the MAC address is learned on the new interface. When the system detects this “station move,” the system clears the entry learned on the original interface and installs a new entry on the new interface.

Learning Limit Violation Actions

To configure the system to take an action when the MAC learning limit is reached on an interface and a new address is received, use one of the following options with the `mac learning-limit` command:

Task	Command Syntax	Command Mode
Generate a system log message when the MAC learning limit is exceeded.	<code>mac learn-limit-violation log</code>	INTERFACE
Shut down the interface and generate a system log message when the MAC learning limit is exceeded.	<code>mac learn-limit-violation shutdown</code>	INTERFACE

Station Move Violation Actions

`no-station-move` is the default behavior. To configure the system to take an action if a station move occurs, use one of the following options with the `mac learning-limit` command:.

Task	Command Syntax	Command Mode
Generate a system log message indicating a station move.	<code>mac station-move-violation log</code>	INTERFACE
Shut down the first port to learn the MAC address.	<code>mac station-move-violation shutdown-original</code>	INTERFACE
Shut down the second port to learn the MAC address.	<code>mac station-move-violation shutdown-offending</code>	INTERFACE

Task	Command Syntax	Command Mode
Shut down both the first and second port to learn the MAC address.	mac station-move-violation shutdown-both	INTERFACE


To display a list of interfaces configured with MAC learning limit or station move violation actions, use the following command:

Task	Command Syntax	Command Mode
Display a list of all of the interfaces configured with MAC learning limit or station move violation.	show mac learning-limit violate-action	CONFIGURATION

Recovering from Learning Limit and Station Move Violations

After a learning-limit or station-move violation shuts down an interface, you must manually reset it. To do this, use the following commands:


Task	Command Syntax	Command Mode
Reset interfaces in ERR_Disabled state caused by a learning limit violation or station move violation.	mac learning-limit reset	CONFIGURATION
Reset interfaces in ERR_Disabled state caused by a learning limit violation.	mac learning-limit reset learn-limit-violation [<i>interface</i> all]	CONFIGURATION
Reset interfaces in ERR_Disabled state caused by a station move violation.	mac learning-limit reset station-move-violation [<i>interface</i> all]	CONFIGURATION

 **Note:** Alternatively, you can reset the interface by shutting it down using the shutdown command and then reenabling it using the command the no shutdown command.

Network Interface Controller (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 17-1 shows a topology where two NICs have been teamed together. In this case, if the primary NIC fails, traffic switches to the secondary NIC, because they are represented by the same set of addresses.

Figure 17-1. Redundant NICs with NIC Teaming

When you use NIC teaming, consider that the server MAC address is originally learned on Port 0/1 of the switch (Figure 17-2). 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 `mac-address-table station-move refresh-arp` command on the Dell Force10 switch at the time that NIC teaming is being configured on the server.



Note: If you do not configure this command, traffic continues to be forwarded to the failed NIC until the ARP entry on the switch times out.

Figure 17-2. Configuring mac-address-table station-move refresh-arp Command

MAC Move Optimization

Station-move detection takes 5000ms because this is the interval at which the detection algorithm runs.

threshold is the number of times a station move must be detected in a single interval in order to trigger a system log message. For example, if you configure `mac-address-table station-move threshold 2 time-interval 5000`, and 4 station moves occur in 5000ms, two log messages are generated.

Link Layer Discovery Protocol (LLDP)

This chapter contains the following sections:

- [Overview](#)
- [TIA-1057 \(LLDP-MED\) Overview](#)
- [Configuring LLDP](#)

Overview

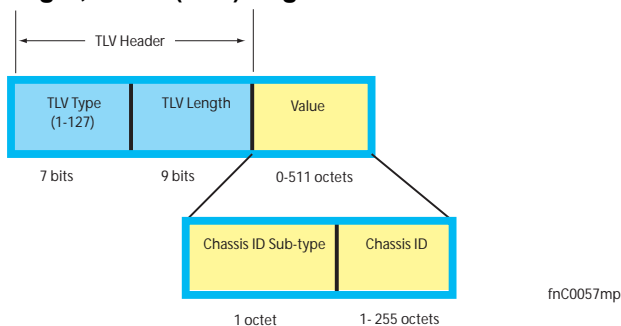
Link layer discovery protocol (LLDP)—defined by IEEE 802.1AB—is a protocol that enables a local area network (LAN) device to advertise its configuration and receive configuration information from adjacent LLDP-enabled LAN infrastructure devices. The collected information is stored in a management information base (MIB) on each device, and is accessible via a simple network management protocol (SNMP).

Protocol Data Units

Configuration information is exchanged in the form of type, length, value (TLV) segments. [Figure 18-1](#) shows the Chassis ID TLV.

- **Type**—the kind of information included in the TLV
- **Length**—the value, in octets, of the TLV after the Length field
- **Value**—the configuration information that the agent is advertising

Figure 18-1. Type, Length, Value (TLV) Segment



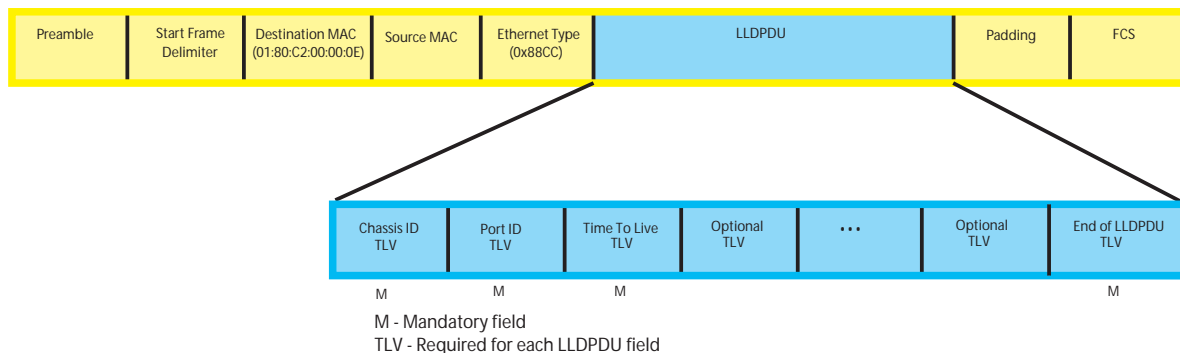
TLVs are encapsulated in a frame called an LLDP data unit (LLDPDU) (Figure 18-2), which is transmitted from one LLDP-enabled device to its LLDP-enabled neighbors. LLDP is a one-way protocol. LLDP-enabled devices (LLDP agents) can transmit and/or receive advertisements, but they cannot solicit and do not respond to advertisements.

There are five types of TLVs (Table 18-1). All types are mandatory in the construction of an LLDPDU except Optional TLVs. You can configure the inclusion of individual Optional TLVs.

Table 18-1. Type, Length, Value (TLV) Types

Type	TLV	Description
0	End of LLDPDU	Marks the end of an LLDPDU.
1	Chassis ID	An administratively assigned name that identifies the LLDP agent.
2	Port ID	An administratively assigned name that identifies a port through which TLVs are sent and received.
3	Time to Live	A value that tells the receiving agent how long the information contained in the TLV Value field is valid.
—	Optional	Includes sub-types of TLVs that advertise specific configuration information. These sub-types are Management TLVs, IEEE 802.1, IEEE 802.3, and TIA-1057 Organizationally Specific TLVs.

Figure 18-2. LLDPDU Frame



fnC0047mp

Optional TLVs

The Dell Force10 operating software (FTOS) supports the following optional TLVs:

- Management TLVs
- IEEE 802.1 and 802.3 Organizationally Specific TLVs
- TIA-1057 Organizationally Specific TLVs

Management TLVs


A Management TLV is an Optional TLVs sub-type. This kind of TLV contains essential management information about the sender. The five types are described in [Table 18-2](#).

Organizationally Specific TLVs

Organizationally specific TLVs can be defined by a professional organization or a vendor. They have two mandatory fields ([Figure 18-3](#)) in addition to the basic TLV fields ([Figure 18-1](#)):

- Organizationally Unique Identifier (OUI)—a unique number assigned by the IEEE to an organization or vendor.
- OUI Sub-type—These sub-types indicate the kind of information in the following data field. The sub-types are determined by the owner of the OUI.

Figure 18-3. Organizationally Specific TLV



IEEE Organizationally Specific TLVs

Eight TLV types have been defined by the IEEE 802.1 and 802.3 working groups ([Table 18-2](#)) as a basic part of LLDP; the IEEE OUI is 00-80-C2. You can configure the Dell Force10 system to advertise any or all of these TLVs.

Table 18-2. Optional TLV Types

Type	TLV	Description
Optional TLVs		
4	Port description	A user-defined alphanumeric string that describes the port. FTOS does not currently support this TLV.
5	System name	A user-defined alphanumeric string that identifies the system.
6	System description	A user-defined alphanumeric string that describes the system.
7	System capabilities	Identifies the chassis as one or more of the following: repeater, bridge, WLAN Access Point, Router, Telephone, DOCSIS cable device, end station only, or other.
8	Management address	Indicates the network address of the management interface. FTOS does not currently support this TLV.
IEEE 802.1 Organizationally Specific TLVs		
127	Port-VLAN ID	On Dell Force10 systems, indicates the untagged VLAN to which a port belongs.

Table 18-2. Optional TLV Types

Type	TLV	Description
127	Port and Protocol VLAN ID	On Dell Force10 systems, indicates the tagged VLAN to which a port belongs (and the untagged VLAN to which a port belongs if the port is in hybrid mode).
127	VLAN Name	Indicates the user-defined alphanumeric string that identifies the VLAN.
127	Protocol Identity	Indicates the protocols that the port can process. FTOS does not currently support this TLV.
IEEE 802.3 Organizationally Specific TLVs		
127	MAC/PHY Configuration/Status	Indicates the capability and current setting of the duplex status and bit rate, and whether the current settings are the result of auto-negotiation. This TLV is not available in the FTOS implementation of LLDP, but is available and mandatory (non-configurable) in the LLDP-MED implementation.
127	Power via MDI	Dell Force10 supports the LLDP-MED protocol, which recommends that Power via MDI TLV is not implemented, and therefore Dell Force10 implements Extended Power via MDI TLV only.
127	Link Aggregation	Indicates whether the link is capable of being aggregated, whether it is currently in a LAG, and the port identification of the LAG. FTOS does not currently support this TLV.
127	Maximum Frame Size	Indicates the maximum frame size capability of the MAC and PHY.

TIA-1057 (LLDP-MED) Overview

Link layer discovery protocol—media endpoint discovery (LLDP-MED)—as defined by ANSI/TIA-1057— provides additional organizationally specific TLVs so that endpoint devices and network connectivity devices can advertise their characteristics and configuration information; the OUI for the Telecommunications Industry Association (TIA) is 00-12-BB.

- **LLDP-MED Endpoint Device**—any device that is on an IEEE 802 LAN network edge can communicate using IP and uses the LLDP-MED framework.
- **LLDP-MED Network Connectivity Device**—any device that provides access to an IEEE 802 LAN to an LLDP-MED endpoint device and supports IEEE 802.1AB (LLDP) and TIA-1057 (LLDP-MED). The Dell Force10 system is an LLDP-MED network connectivity device.

With regard to connected endpoint devices, LLDP-MED provides network connectivity devices with the ability to:

- manage inventory
- manage Power over Ethernet (PoE)
- identify physical location
- identify network policy

LLDP-MED is designed for, but not limited to, voice over IP (VoIP) endpoints.

TIA Organizationally Specific TLVs

The Dell Force10 system is an LLDP-MED Network Connectivity Device (Device Type 4). Network connectivity devices are responsible for:

- transmitting an LLDP-MED capabilities TLV to endpoint devices
- storing the information that endpoint devices advertise

Table 18-3 list the five types of TIA-1057 Organizationally Specific TLVs.

Table 18-3. TIA-1057 (LLDP-MED) Organizationally Specific TLVs

Type	Sub-type	TLV	Description
127	1	LLDP-MED Capabilities	Indicates: <ul style="list-style-type: none"> • whether the transmitting device supports LLDP-MED • what LLDP-MED TLVs it supports • LLDP device class
127	2	Network Policy	Indicates the application type, VLAN ID, Layer 2 Priority, and DSCP value
127	3	Location Identification	Indicates the physical location of the device expressed in one of three possible formats: <ul style="list-style-type: none"> • Coordinate Based LCI • Civic Address LCI • Emergency Call Services ELIN
127	4	Extended Power via MDI	Indicates power requirements, priority, and power status
Inventory Management TLVs			Implementation of this set of TLVs is optional in LLDP-MED devices. None or all TLVs must be supported. FTOS does not currently support these TLVs.
127	5	Inventory - Hardware Revision	Indicates the hardware revision of the LLDP-MED device.
127	6	Inventory - Firmware Revision	Indicates the firmware revision of the LLDP-MED device.
127	7	Inventory - Software Revision	Indicates the software revision of the LLDP-MED device.
127	8	Inventory - Serial Number	Indicates the device serial number of the LLDP-MED device.
127	9	Inventory - Manufacturer Name	Indicates the manufacturer of the LLDP-MED device.
127	10	Inventory - Model Name	Indicates the model of the LLDP-MED device.
127	11	Inventory - Asset ID	Indicates a user specified device number to manage inventory.
127	12-255	Reserved	—

LLDP-MED Capabilities TLV

The LLDP-MED Capabilities TLV communicates the types of TLVs that the endpoint device and the network connectivity device support. LLDP-MED network connectivity devices must transmit the Network Policies TLV.

- The value of the LLDP-MED Capabilities field in the TLV is a 2 octet bitmap (Figure 18-4), each bit represents an LLDP-MED capability (Table 18-4).
- The possible values of the LLDP-MED Device Type is listed in Table 18-5. The Dell Force10 system is a Network Connectivity device, which is Type 4.

When you enable LLDP-MED in FTOS (using the advertise med command), the system begins transmitting this TLV.

Figure 18-4. LLDP-MED Capabilities TLV




Table 18-4. FTOS LLDP-MED Capabilities

Bit Position	TLV	FTOS Support
0	LLDP-MED Capabilities	Yes
1	Network Policy	Yes
2	Location Identification	Yes
3	Extended Power via MDI-PSE	Yes
4	Extended Power via MDI-PD	No
5	Inventory	No
6-15	reserved	No

Table 18-5. LLDP-MED Device Types

Value	Device Type
0	Type Not Defined
1	Endpoint Class 1
2	Endpoint Class 2
3	Endpoint Class 3
4	Network Connectivity
5-255	Reserved

LLDP-MED Network Policies TLV

A network policy in the context of LLDP-MED is a device's virtual local area network (VLAN) configuration and associated Layer 2 and Layer 3 configurations, specifically:

- VLAN ID
- VLAN tagged or untagged status
- Layer 2 priority
- DSCP value


The application type is represented by an integer (the Type integer in [Table 18-6](#)), which indicates a device function for which a unique network policy is defined. An individual LLDP-MED Network Policy TLV is generated for each application type that you specify with the FTOS command line interface (CLI) ([Advertising TLVs](#)).



Note: With regard to [Table 18-6](#), signaling is a series of control packets that are exchanged between an endpoint device and a network connectivity device to establish and maintain a connection. These signal packets might require a different network policy than the media packets for which a connection is made. In this case, configure the signaling application.

Table 18-6. Network Policy Applications


Type	Application	Description
0	Reserved	—
1	Voice	Specify this application type for dedicated IP telephony handsets and other appliances supporting interactive voice services.
2	Voice Signaling	Specify this application type only if voice control packets use a separate network policy than voice data.
3	Guest Voice	Specify this application type for a separate limited voice service for guest users with their own IP telephony handsets and other appliances supporting interactive voice services.
4	Guest Voice Signaling	Specify this application type only if guest voice control packets use a separate network policy than voice data.
5	Softphone Voice	Softphone is a computer program that enables IP telephony on a computer, rather than using a phone. Specify this application type for this type of endpoint device.
6	Video Conferencing	Specify this application type for dedicated video conferencing and other similar appliances supporting real-time interactive video.
7	Streaming Video	Specify this application type for broadcast or multicast based video content distribution and other similar applications supporting streaming video services. This does not include video applications relying on TCP with buffering.
8	Video Signaling	Specify this application type only if video control packets use a separate network policy than video data.
9-255	Reserved	—

Figure 18-5. LLDP-MED Policies TLV

Extended Power via MDI TLV

The Extended Power via MDI TLV enables advanced power over ethernet (PoE) management between LLDP-MED endpoints and network connectivity devices (Figure 18-6). Advertise the Extended Power via MDI on all ports that are connected to an 802.3af powered, LLDP-MED endpoint device.

- **Power Type**—there are two possible power types: power sourcing entity (PSE) or power device (PD). The Dell Force10 system is a PSE, which corresponds to a value of 0, based on the TIA-1057 specification.
- **Power Source**—there are two possible power sources: Primary and Backup. The Dell Force10 system is a Primary Power Source, which corresponds to a value of 1, based on the TIA-1057 specification.
- **Power Priority**—there are three possible priorities: Low, High, and Critical. On Dell Force10 systems, the default power priority is High, which corresponds to a value of 2 based on the TIA-1057 specification. You can configure a different power priority through the CLI. Dell Force10 also honors the power priority value sent by the powered device. However, the CLI configuration takes precedence.
- **Power Value**—Dell Force10 advertises the maximum amount of power that can be supplied on the port. By default it is 15.4W, which corresponds to a Power Value of 130, based on the TIA-1057 specification. You can advertise a different Power Value using the max-milliwatts option with the power inline auto | static command. Dell Force10 also honors the power value (power requirement) sent by the powered device when the port is configured for power inline auto.

Figure 18-6. Extended Power via MDI TLV

Configuring LLDP

Configuring LLDP is a two-step process:

1. [Enabling LLDP](#)
2. [Advertising TLVs](#)

Related Configuration Tasks

- [Viewing the LLDP Configuration](#)
- [Viewing Information Advertised by Adjacent LLDP Agents](#)
- [Configuring LLDPDU Intervals](#)
- [Configuring Transmit and Receive Mode](#)
- [Configuring a Time to Live](#)
- [Debugging LLDP](#)

Important Points to Remember

- LLDP is disabled by default.
- Dell Force10 systems support up to eight neighbors per interface.
- Dell Force10 systems support a maximum of 8000 total neighbors per system. If the number of interfaces multiplied by eight exceeds the maximum, the system does not configure more than 8000.
- INTERFACE level configurations override all CONFIGURATION level configurations.
- LLDP is not hitless.

LLDP Compatibility

- Spanning tree and Force10 Ring Protocol “blocked” ports allow LLDPDUs.

CONFIGURATION versus INTERFACE Configurations

All LLDP configuration commands are available in PROTOCOL LLDP mode, which is a sub-mode of CONFIGURATION mode and INTERFACE mode ([Figure 18-7](#)).

- Configurations made at the CONFIGURATION level are global, that is, they affect all interfaces on the system.
- Configurations made at the INTERFACE level affect only the specific interface. They override CONFIGURATION level configurations.

Figure 18-7. Configuration and Interface mode LLDP Commands

```

R1(conf)#protocol lldp
R1(conf-lldp)#?
advertise          Advertise TLVs
dcbx               Configure Dcbx Parameters
disable            Disable LLDP protocol globally
end                Exit from configuration mode
exit               Exit from LLDP configuration mode
fcoe               Configure priority bits for FCoE traffic
hello              LLDP hello configuration
iscsi              Configure priority bits for ISCSI traffic
mode               LLDP mode configuration (default = rx and tx)
multiplier         LLDP multiplier configuration
no                 Negate a command or set its defaults
show               Show LLDP configuration
R1(conf-lldp)#exit
R1(conf)#interface tengigabitethernet 1/31
R1(conf-if-te-1/31)#protocol lldp
R1(conf-if-te-1/31-lldp)#?
advertise          Advertise TLVs
dcbx               Configure Dcbx Parameters
disable            Disable LLDP protocol on this interface
end                Exit from configuration mode
exit               Exit from LLDP configuration mode
hello              LLDP hello configuration
mode               LLDP mode configuration (default = rx and tx)
multiplier         LLDP multiplier configuration
no                 Negate a command or set its defaults
show               Show LLDP configuration
no                 Negate a command or set its defaults
show               Show LLDP configuration
R1(conf-if-te-1/31-lldp)#

```

Enabling LLDP

LLDP is disabled by default. You can enable and disable LLDP globally or per interface. If you enable LLDP globally, all up interfaces send periodic LLDPDUs. To enable LLDP, follow these steps:

Step	Task	Command	Command Mode
1	Enter Protocol LLDP mode.	protocol lldp	CONFIGURATION or INTERFACE
2	Enable LLDP.	no disable	PROTOCOL LLDP

Disabling and Undoing LLDP

- Disable LLDP globally or for an interface using the disable command.
- Undo an LLDP configuration by preceding the relevant command with the keyword no command.

Advertising TLVs

You can configure the system to advertise TLVs out of all interfaces or out of specific interfaces.


- If you configure the system globally, all interfaces send LLDPDUs with the specified TLVs.
- If you configure an interface, only the interface sends LLDPDUs with the specified TLVs.

If you configure LLDP both globally and at interface level, the interface-level configuration overrides the global configuration. To advertise TLVs, follow these steps:

Step	Task	Command	Command Mode
1	Enter LLDP mode.	<code>protocol lldp</code>	CONFIGURATION or INTERFACE
2	Advertise one or more TLVs. Include the keyword for each TLV you want to advertise. <ul style="list-style-type: none">• For management TLVs: <code>system-capabilities</code>, <code>system-description</code>• For 802.1 TLVs: <code>port-protocol-vlan-id</code>, <code>port-vlan-id</code>, <code>vlan-name</code>• For 802.3 TLVs: <code>max-frame-size</code>• For TIA-1057 TLVs:<ul style="list-style-type: none">• <code>guest-voice</code>• <code>guest-voice-signaling</code>• <code>location-identification</code>• <code>power-via-mdi</code>• <code>softphone-voice</code>• <code>streaming-video</code>• <code>video-conferencing</code>• <code>video-signaling</code>• <code>voice</code>• <code>voice-signaling</code>	<code>advertise {management-tlv dot1-tlv dot3-tlv med dcbx-appln-tlv dcbx-tlv interface-port-desc}</code>	PROTOCOL LLDP

In [Figure 18-8](#), LLDP is enabled globally. R1 and R2 are transmitting periodic LLDPDUs that contain management, 802.1, and 802.3 TLVs.

Figure 18-8. Configuring LLDP



Viewing the LLDP Configuration

To display the LLDP configuration, use the `show config` command in either CONFIGURATION or INTERFACE mode ([Figure 18-9](#)) and ([Figure 18-10](#)).

Figure 18-9. Viewing LLDP Global Configurations

```
R1(conf)#protocol lldp
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  hello 10
  no disable
R1(conf-lldp)#
```

Figure 18-10. Viewing LLDP Interface Configurations

```
R1(conf-lldp)#exit
R1(conf)#interface tengigabitethernet 1/31
R1(conf-if-te-1/31)#show config
!
interface TenGigabitEthernet 1/31
  no ip address
!
  no shutdown
R1(conf-if-te-1/31)#protocol lldp
R1(conf-if-te-1/31-lldp)#show config
!
  protocol lldp
R1(conf-if-te-1/31-lldp)#
```

Viewing Information Advertised by Adjacent LLDP Agents

To display brief information about adjacent devices, use the `show lldp neighbors` command (Figure 18-11). To display all of the information that neighbors are advertising, use the `show lldp neighbors detail` command (Figure 18-12).

Figure 18-11. Viewing Brief Information Advertised by Adjacent LLDP Agents

```
R1(conf-if-te-1/31-lldp)#end
R1(conf-if-te-1/31)#do show lldp neighbors
Loc PortID    Rem Host Name      Rem Port Id        Rem Chassis Id
-----
Te 0/2       -                  00:00:c9:b1:3b:82  00:00:c9:b1:3b:82
Te 0/3       -                  00:00:c9:ad:f6:12  00:00:c9:ad:f6:12
```

Figure 18-12. Viewing All Information Advertised by Adjacent LLDP Agent

```
FTOS#show lldp neighbors detail
=====
Local Interface Te 0/2 has 1 neighbor
Total Frames Out: 16843
Total Frames In: 17464
Total Neighbor information Age outs: 0
Total Multiple Neighbors Detected: 0
Total Frames Discarded: 0
Total In Error Frames: 0
Total Unrecognized TLVs: 0
Total TLVs Discarded: 0
Next packet will be sent after 16 seconds
The neighbors are given below:
-----

Remote Chassis ID Subtype: Mac address (4)
Remote Chassis ID: 00:00:c9:b1:3b:82
Remote Port Subtype: Mac address (3)
Remote Port ID: 00:00:c9:b1:3b:82
Local Port ID: TenGigabitEthernet 0/2
Locally assigned remote Neighbor Index: 7
Remote TTL: 120
Information valid for next 105 seconds
Time since last information change of this neighbor: 1d21h56m
Remote System Desc: Emulex OneConnect 10Gb Multi function Adapter
Existing System Capabilities: Station only
Enabled System Capabilities: Station only
-----

=====
Local Interface Te 0/3 has 1 neighbor
Total Frames Out: 39165
Total Frames In: 40650
Total Neighbor information Age outs: 0
Total Multiple Neighbors Detected: 0
Total Frames Discarded: 0
Total In Error Frames: 0
Total Unrecognized TLVs: 0
Total TLVs Discarded: 0
Next packet will be sent after 4 seconds
The neighbors are given below:
-----

Remote Chassis ID Subtype: Mac address (4)
Remote Chassis ID: 00:00:c9:ad:f6:12
Remote Port Subtype: Mac address (3)
Remote Port ID: 00:00:c9:ad:f6:12
Local Port ID: TenGigabitEthernet 0/3
```


Configuring LLDPDU Intervals

LLDPDUs are transmitted periodically; the default interval is 30 seconds. To configure a non-default transmit interval—at CONFIGURATION level or INTERFACE level—use the hello command (Figure 18-13).

Figure 18-13. Configuring LLDPDU Transmit and Receive Mode

```
R1(conf)#protocol lldp
R1(conf-lldp)#show config
!
protocol lldp
 advertise dot1-tlv port-protocol-vlan-id port-vlan-id
 advertise dot3-tlv max-frame-size
 advertise management-tlv system-capabilities system-description
 no disable
R1(conf-lldp)#mode ?
rx                Rx only
tx                Tx only
R1(conf-lldp)#mode tx
R1(conf-lldp)#show config
!
protocol lldp
 advertise dot1-tlv port-protocol-vlan-id port-vlan-id
 advertise dot3-tlv max-frame-size
 advertise management-tlv system-capabilities system-description
 mode tx
 no disable
R1(conf-lldp)#no mode
R1(conf-lldp)#show config
!
protocol lldp
 advertise dot1-tlv port-protocol-vlan-id port-vlan-id
 advertise dot3-tlv max-frame-size
 advertise management-tlv system-capabilities system-description
 no disable
R1(conf-lldp)#
```

Configuring Transmit and Receive Mode

After you enable LLDP, Dell Force10 systems transmit *and* receive LLDPDUs by default. You can configure the system—at CONFIGURATION level or INTERFACE level—to transmit only by executing the mode tx command, or receive only by executing the mode rx command. To return to the default setting, use the no mode command (Figure 18-14).

Figure 18-14. Configuring LLDPDU Transmit and Receive Mode

```
R1(conf)#protocol lldp
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#mode ?
rx                Rx only
tx                Tx only
R1(conf-lldp)#mode tx
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  mode tx
  no disable
R1(conf-lldp)#no mode
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#
```

Configuring a Time to Live

The information received from a neighbor expires after a specific amount of time (measured in seconds) called a Time to Live (TTL). The TTL is the product of the LLDPDU transmit interval (hello) and an integer called a *multiplier*. The default multiplier is 4, which results in a default TTL of 120 seconds. To adjust the TTL value—at CONFIGURATION level or INTERFACE level—use the multiplier command. To return to the default multiplier value, use the no multiplier command (Figure 18-15).

Figure 18-15. Configuring LLDPDU Time to Live

```
R1(conf-lldp)#show config
!
protocol lldp
 advertise dot1-tlv port-protocol-vlan-id port-vlan-id
 advertise dot3-tlv max-frame-size
 advertise management-tlv system-capabilities system-description
 no disable
R1(conf-lldp)#multiplier ?
<2-10>                Multiplier (default=4)
R1(conf-lldp)#multiplier 5
R1(conf-lldp)#show config
!
protocol lldp
 advertise dot1-tlv port-protocol-vlan-id port-vlan-id
 advertise dot3-tlv max-frame-size
 advertise management-tlv system-capabilities system-description
 multiplier 5
 no disable
R1(conf-lldp)#no multiplier
R1(conf-lldp)#show config
!
protocol lldp
 advertise dot1-tlv port-protocol-vlan-id port-vlan-id
 advertise dot3-tlv max-frame-size
 advertise management-tlv system-capabilities system-description
 no disable
R1(conf-lldp)#
```

Debugging LLDP

The debug lldp command allows you to view the TLVs that your system is sending and receiving.

- Use the debug lldp brief command to view a readable version of the TLVs.
- Use the debug lldp detail command to view a readable version of the TLVs plus a hexadecimal version of the entire LLDPDU.

Figure 18-16. debug lldp detail—LLDPDU Packet Dissection

```

FTOS# debug lldp interface tengigabitethernet 1/2 packet detail tx
FTOS#1w1d19h :Transmit timer blew off for local interface TenGig 1/2
1w1d19h :Forming LLDP pkt to send out of interface TenGig 1/2
1w1d19h :TLV:Chassis ID,Len: 7, Subtype:Mac address (4),Value:00:01:e8:0d:b6:d6
1w1d19h :TLV:Port ID,Len: 20, Subtype:Interface name (5),Value:TenGigabitEthernet 1/2
1w1d19h :TLV:TTL,Len: 2,Value: 120
1w1d19h :TLV:SYS_DESC,Len: 207,Value:Force10 Networks Real Time Operating System Software.Force10
Operating System Version: 1.0.Force10 Application Software Version: E_MAIN4.7.5.276.Copyright (c)1999-Build
Time: Fri Oct 26 12:22:22 PDT 2007
1w1d19h :TLV:SYSTEM_CAPAB,Len: 4,Value:Existing:Repeater Bridge Router,Enabled:Repeater Bridge Router
1w1d19h :TLV:ENDOFDPDU,Len: 0
1w1d19h :Sending LLDP pkt out of TenGig 1/2 of length 270
1w1d19h :Packet dump:
1w1d19h : 01 80 c2 00 00 0e 00 01 e8 0d b7 3b 81 00 00 00
1w1d19h : 88 cc 02 07 04 00 01 e8 0d b6 d6 04 14 05 47 69
1w1d19h : 67 61 62 69 74 45 74 68 65 72 6e 65 74 20 31 2f
1w1d19h : 32 06 02 00 78 0c cf 46 6f 72 63 65 31 30 20 4e
1w1d19h : 65 74 77 6f 72 6b 73 20 52 65 61 6c 20 54 69 6d
1w1d19h : 65 20 4f 70 65 72 61 74 69 6e 67 20 53 79 73 74
1w1d19h : 65 6d 20 53 6f 66 74 77 61 72 65 2e 20 46 6f 72
1w1d19h : 63 65 31 30 20 4f 70 65 72 61 74 69 6e 67 20 53
1w1d19h : 79 73 74 65 6d 20 56 65 72 73 69 6f 6e 3a 20 31
1w1d19h : 2e 30 2e 20 46 6f 72 63 65 31 30 20 41 70 70 6c
1w1d19h : 69 63 61 74 69 6f 6e 20 53 6f 66 74 77 61 72 65
1w1d19h : 20 56 65 72 73 69 6f 6e 3a 20 45 5f 4d 41 49 4e
1w1d19h : 34 2e 37 2e 35 2e 32 37 36 2e 20 43 6f 70 79 72
1w1d19h : 69 67 68 74 20 28 63 29 20 31 39 39 39 2d 42 75
1w1d19h : 69 6c 64 20 54 69 6d 65 3a 20 46 72 69 20 4f 63
1w1d19h : 74 20 32 36 20 31 32 3a 32 32 3a 32 32 20 50 44
1w1d19h : 54 20 32 30 30 37 0e 04 00 16 00 16 00 00
1w1d19h :LLDP frame sent out successfully of TenGig 1/2
1w1d19h :Started Transmit timer for Loc interface TenGig 1/2 for time 30 sec

```

Relevant Management Objects

FTOS supports all IEEE 802.1AB MIB objects.

- [Table 18-7](#) lists the objects associated with received and transmitted TLVs.
- [Table 18-8](#) lists the objects associated with the LLDP configuration on the local agent.
- [Table 18-9](#) lists the objects associated with IEEE 802.1AB Organizationally Specific TLVs.
- [Table 18-10](#) lists the objects associated with received and transmitted LLDP-MED TLVs.

Table 18-7. LLDP Configuration MIB Objects

MIB Object Category	LLDP Variable	LLDP MIB Object	Description
LLDP Configuration	adminStatus	lldpPortConfigAdminStatus	Whether the local LLDP agent is enabled for transmit, receive, or both
	msgTxHold	lldpMessageTxHoldMultiplier	Multiplier value
	msgTxInterval	lldpMessageTxInterval	Transmit Interval value
	rxInfoTTL	lldpRxInfoTTL	Time to Live for received TLVs
	txInfoTTL	lldpTxInfoTTL	Time to Live for transmitted TLVs
Basic TLV Selection	mibBasicTLVsTxEnable	lldpPortConfigTLVsTxEnable	Indicates which management TLVs are enabled for system ports
	mibMgmtAddrInstanceTxEnable	lldpManAddrPortsTxEnable	The management addresses defined for the system and the ports through which they are enabled for transmission
LLDP Statistics	statsAgeoutsTotal	lldpStatsRxPortAgeoutsTotal	Total number of times that a neighbors information is deleted on the local system due to an rxInfoTTL timer expiration
	statsFramesDiscardedTotal	lldpStatsRxPortFramesDiscardedTotal	Total number of LLDP frames received then discarded
	statsFramesInErrorsTotal	lldpStatsRxPortFramesErrors	Total number of LLDP frames received on a port with errors
	statsFramesInTotal	lldpStatsRxPortFramesTotal	Total number of LLDP frames received through the port
	statsFramesOutTotal	lldpStatsTxPortFramesTotal	Total number of LLDP frames transmitted through the port
	statsTLVsDiscardedTotal	lldpStatsRxPortTLVsDiscardedTotal	Total number of TLVs received then discarded
	statsTLVsUnrecognizedTotal	lldpStatsRxPortTLVsUnrecognizedTotal	Total number of all TLVs the local agent does not recognize

Table 18-8. LLDP System MIB Objects

TLV Type	TLV Name	TLV Variable	System	LLDP MIB Object
1	Chassis ID	chassis ID subtype	Local	lldpLocChassisIdSubtype
			Remote	lldpRemChassisIdSubtype
		chassid ID	Local	lldpLocChassisId
			Remote	lldpRemChassisId

Table 18-8. LLDP System MIB Objects

TLV Type	TLV Name	TLV Variable	System	LLDP MIB Object
2	Port ID	port subtype	Local	lldpLocPortIdSubtype
			Remote	lldpRemPortIdSubtype
		port ID	Local	lldpLocPortId
			Remote	lldpRemPortId
4	Port Description	port description	Local	lldpLocPortDesc
			Remote	lldpRemPortDesc
5	System Name	system name	Local	lldpLocSysName
			Remote	lldpRemSysName
6	System Description	system description	Local	lldpLocSysDesc
			Remote	lldpRemSysDesc
7	System Capabilities	system capabilities	Local	lldpLocSysCapSupported
			Remote	lldpRemSysCapSupported
8	Management Address	enabled capabilities	Local	lldpLocSysCapEnabled
			Remote	lldpRemSysCapEnabled
		management address length	Local	lldpLocManAddrLen
			Remote	lldpRemManAddrLen
		management address subtype	Local	lldpLocManAddrSubtype
			Remote	lldpRemManAddrSubtype
		management address	Local	lldpLocManAddr
			Remote	lldpRemManAddr
		interface numbering subtype	Local	lldpLocManAddrIfSubtype
			Remote	lldpRemManAddrIfSubtype
		interface number	Local	lldpLocManAddrIfId
			Remote	lldpRemManAddrIfId
		OID	Local	lldpLocManAddrOID
			Remote	lldpRemManAddrOID

Table 18-9. LLDP 802.1 Organizationally Specific TLV MIB Objects

TLV Type	TLV Name	TLV Variable	System	LLDP MIB Object
127	Port-VLAN ID	PVID	Local	lldpXdot1LocPortVlanId
			Remote	lldpXdot1RemPortVlanId

Table 18-9. LLDP 802.1 Organizationally Specific TLV MIB Objects

TLV Type	TLV Name	TLV Variable	System	LLDP MIB Object
127	Port and Protocol VLAN ID	port and protocol VLAN supported	Local	lldpXdot1LocProtoVlanSupported
			Remote	lldpXdot1RemProtoVlanSupported
		port and protocol VLAN enabled	Local	lldpXdot1LocProtoVlanEnabled
			Remote	lldpXdot1RemProtoVlanEnabled
		PPVID	Local	lldpXdot1LocProtoVlanId
			Remote	lldpXdot1RemProtoVlanId
127	VLAN Name	VID	Local	lldpXdot1LocVlanId
			Remote	lldpXdot1RemVlanId
		VLAN name length	Local	lldpXdot1LocVlanName
			Remote	lldpXdot1RemVlanName
		VLAN name	Local	lldpXdot1LocVlanName
			Remote	lldpXdot1RemVlanName

Table 18-10. LLDP-MED System MIB Objects

TLV Sub-Type	TLV Name	TLV Variable	System	LLDP-MED MIB Object
1	LLDP-MED Capabilities	LLDP-MED Capabilities	Local	lldpXMedPortCapSupported lldpXMedPortConfigTLVsTx Enable
			Remote	lldpXMedRemCapSupported, lldpXMedRemConfigTLVsTx Enable
		LLDP-MED Class Type	Local	lldpXMedLocDeviceClass
			Remote	lldpXMedRemDeviceClass

Table 18-10. LLDP-MED System MIB Objects

TLV Sub-Type	TLV Name	TLV Variable	System	LLDP-MED MIB Object
2	Network Policy	Application Type	Local	lldpXMedLocMediaPolicyAppType
			Remote	lldpXMedRemMediaPolicyAppType
		Unknown Policy Flag	Local	lldpXMedLocMediaPolicyUnknown
			Remote	lldpXMedLocMediaPolicyUnknown
		Tagged Flag	Local	lldpXMedLocMediaPolicyTagged
			Remote	lldpXMedLocMediaPolicyTagged
		VLAN ID	Local	lldpXMedLocMediaPolicyVlanID
			Remote	lldpXMedRemMediaPolicyVlanID
		L2 Priority	Local	lldpXMedLocMediaPolicyPriority
			Remote	lldpXMedRemMediaPolicyPriority
		DSCP Value	Local	lldpXMedLocMediaPolicyDscp
			Remote	lldpXMedRemMediaPolicyDscp
3	Location Identifier	Location Data Format	Local	lldpXMedLocLocationSubtype
			Remote	lldpXMedRemLocationSubtype
		Location ID Data	Local	lldpXMedLocLocationInfo
			Remote	lldpXMedRemLocationInfo

Table 18-10. LLDP-MED System MIB Objects

TLV Sub-Type	TLV Name	TLV Variable	System	LLDP-MED MIB Object
4	Extended Power via MDI	Power Device Type	Local	lldpXMedLocXPoEDeviceType
			Remote	lldpXMedRemXPoEDeviceType
		Power Source	Local	lldpXMedLocXPoEPSEPowerSource, lldpXMedLocXPoEPDPowerSource
			Remote	lldpXMedRemXPoEPSEPowerSource, lldpXMedRemXPoEPDPowerSource
		Power Priority	Local	lldpXMedLocXPoEPDPowerPriority, lldpXMedLocXPoEPSEPortPDPriority
			Remote	lldpXMedRemXPoEPSEPowerPriority, lldpXMedRemXPoEPDPowerPriority
		Power Value	Local	lldpXMedLocXPoEPSEPortPowerAv, lldpXMedLocXPoEPDPowerReq
			Remote	lldpXMedRemXPoEPSEPowerAv, lldpXMedRemXPoEPDPowerReq

Multiple Spanning Tree Protocol (MSTP)


Overview

Multiple spanning tree protocol (MSTP)—specified in IEEE 802.1Q-2003—is an rapid spanning tree protocol (RSTP)-based spanning tree variation that improves on PVST+. MSTP allows multiple spanning tree instances and allows you to map many virtual local area networks (VLANs) to one spanning tree instance to reduce the total number of required instances.

In contrast, per-VLAN spanning tree plus (PVST+) allows a spanning tree instance for each VLAN. This 1:1 approach is not suitable if you have many VLANs, because each spanning tree instance costs bandwidth and processing resources.

In [Figure 19-1](#), three VLANs are mapped to two multiple spanning tree instances (MSTI). VLAN 100 traffic takes a different path than VLAN 200 and 300 traffic. The behavior in [Figure 19-1](#) demonstrates how you can use MSTP to achieve load balancing.

Figure 19-1. MSTP with Three VLANs Mapped to Two Spanning Tree Instances



The Dell Force10 operating software (FTOS) supports three other variations of Spanning Tree (Table 19-1).

Table 19-1. FTOS Supported Spanning Tree Protocols

Dell Force10 Term	IEEE Specification
Spanning Tree Protocol	802.1d
Rapid Spanning Tree Protocol	802.1w
Multiple Spanning Tree Protocol	802.1s
Per-VLAN Spanning Tree Plus	Third Party

Implementation Information

- The FTOS MSTP implementation is based on IEEE 802.1Q-2003 and interoperates only with bridges that also use this standard implementation.
- MSTP is compatible with STP and RSTP.
- FTOS supports only one MSTP region.
- When you enable MSTP, all ports in Layer 2 mode participate in MSTP.
- On the MXL Switch, you can configure 64 MSTIs including the default instance 0 (CIST).

Configure Multiple Spanning Tree Protocol

Configuring Multiple Spanning Tree is a four-step process:

1. Configure interfaces for Layer 2.
2. Place the interfaces in VLANs.
3. Enable MSTP
4. Create MSTP instances and map VLANs to them.

Related Configuration Tasks

- [Create Multiple Spanning Tree Instances](#)
- [Create Multiple Spanning Tree Instances](#)
- [Influence MSTP Root Selection](#)
- [Interoperate with Non-FTOS Bridges](#)
- [Modify Global Parameters](#)
- [Enable BPDU Filtering globally](#)
- [Enable BPDU Filtering globally](#)
- [Configure an EdgePort](#)
- [Flush MAC Addresses after a Topology Change](#)
- [Debugging and Verifying an MSTP Configuration](#)

- Preventing Network Disruptions with BPDU Guard
- SNMP Traps for Root Elections and Topology Changes

Enable Multiple Spanning Tree Globally

MSTP is not enabled by default. To enable MSTP, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Enter PROTOCOL MSTP mode.	protocol spanning-tree mstp	CONFIGURATION
2	Enable MSTP.	no disable	PROTOCOL MSTP

To verify that MSTP is enabled, use the show config command from PROTOCOL MSTP mode (Figure 19-2).

Figure 19-2. Verifying MSTP is Enabled

```
FTOS(conf)#protocol spanning-tree mstp
FTOS(conf-mstp)#no disable
FTOS(conf-mstp)#show config
!
protocol spanning-tree mstp
no disable
FTOS#
```

When you enable MSTP, all physical, VLAN, and port-channel interfaces that are enabled and in Layer 2 mode are automatically part of the MSTI 0.

- Within an MSTI, only one path from any bridge to any other bridge is enabled.
- Bridges block a redundant path by disabling one of the link ports.

Create Multiple Spanning Tree Instances

A single MSTI provides no more benefit than RSTP. To take full advantage of MSTP, you must create multiple MSTIs and map VLANs to them.

To create an MSTI, use the msti command from PROTOCOL MSTP mode. Specify the keyword vlan followed by the VLANs that you want to participate in the MSTI (Figure 19-3).

Figure 19-3. Mapping VLANs to MSTI Instances

```

FTOS(conf)#protocol spanning-tree mstp
FTOS(conf-mstp)#msti 1 vlan 100
FTOS(conf-mstp)#msti 2 vlan 200-300
FTOS(conf-mstp)#show config
!
protocol spanning-tree mstp
no disable
MSTI 1 VLAN 100
MSTI 2 VLAN 200-300

```

All bridges in the MSTP region must have the same VLAN-to-instance mapping. To view which instance a VLAN is mapped, use the `show spanning-tree mst vlan <VLAN-ID>` command from EXEC Privilege mode.

To view the forwarding/discarding state of the ports participating in an MSTI, use the `show spanning-tree msti` command from EXEC Privilege mode (Figure 19-4).

Figure 19-4. Viewing MSTP Port States

```

FTOS#show spanning-tree msti 1
MSTI 1 VLANs mapped 100
Root Identifier has priority 32768, Address 0001.e806.953e
Root Bridge hello time 2, max age 20, forward delay 15, max hops 19
Bridge Identifier has priority 32768, Address 0001.e80d.b6d6
Configured hello time 2, max age 20, forward delay 15, max hops 20
Current root has priority 32768, Address 0001.e806.953e
Number of topology changes 2, last change occurred 1d2h ago on TenGig 1/21
Port 374 (TenGigabitEthernet 1/21) is root Forwarding
Port path cost 2000, Port priority 128, Port Identifier 128.374
Designated root has priority 32768, address 0001.e806.953e
Designated bridge has priority 32768, address 0001.e806.953e
Designated port id is 128.374, designated path cost 2000
Number of transitions to forwarding state 1
BPDU (MRecords): sent 93671, received 46843
The port is not in the Edge port mode, bpdu filter is disabled
Port 384 (TenGigabitEthernet 1/31) is alternate Discarding
Port path cost 2000, Port priority 128, Port Identifier 128.384
Designated root has priority 32768, address 0001.e806.953e
Designated bridge has priority 32768, address 0001.e809.c24a
Designated port id is 128.384, designated path cost 2000
Number of transitions to forwarding state 1
BPDU (MRecords): sent 39291, received 7547
The port is not in the Edge port mode, bpdu filter is disabled

```

Influence MSTP Root Selection

MSTP determines the root bridge, but you can assign one bridge a lower priority to increase the probability of it becoming the root bridge.

To change the bridge priority, use the following command:

Task	Command Syntax	Command Mode
Assign a number as the bridge priority. A lower number increases the probability that the bridge becomes the root bridge. Range: 0 to 61440, in increments of 4096 Default: 32768	<code>msti <i>instance</i> bridge-priority <i>priority</i></code>	PROTOCOL MSTP

The simple configuration (Figure 19-1) by default yields the same forwarding path for both MSTIs. Figure 19-5 shows how R3 is assigned bridge priority 0 for MSTI 2, which elects a different root bridge than MSTI 2.

To view the bridge priority, use the show config command from PROTOCOL MSTP mode (Figure 19-5).

Figure 19-5. Changing the Bridge Priority

```
FTOS(conf-mstp)#msti 2 bridge-priority 0

FTOS(conf-mstp)#show config
!
protocol spanning-tree mstp
 MSTI 2 bridge-priority 0
FTOS(conf-mstp)#
```

Interoperate with Non-FTOS Bridges

FTOS supports only one MSTP region. A region is a combination of three unique qualities:

- **Name** is a mnemonic string you assign to the region. The default region name on FTOS is null.
- **Revision** is a two-byte number. The default revision number on FTOS is 0.
- **VLAN-to-instance mapping** is the placement of a VLAN in an MSTI.

For a bridge to be in the same MSTP region as another, all three of these qualities must match exactly. The default values for the name and revision matches on all Dell Force10 FTOS equipment. If you have non-FTOS equipment that participates in MSTP, ensure these values to match on all the equipment.



Note: Some non-FTOS equipment may implement a non-null default region name. SFTOS, for example, uses the Bridge ID, while others may use a MAC address.

To change the region name or revision, use the following commands:

Task	Command Syntax	Command Mode
Change the region name.	name <i>name</i>	PROTOCOL MSTP
Change the region revision number. <ul style="list-style-type: none"> • Range: 0 to 65535 • Default: 0 	revision <i>number</i>	PROTOCOL MSTP

To view the current region name and revision, use the show spanning-tree mst configuration command from EXEC Privilege mode (Figure 19-6).

Figure 19-6. Viewing the MSTP Region Name and Revision

```
FTOS(conf-mstp)#name my-mstp-region
FTOS(conf-mstp)#exit
FTOS(conf)#do show spanning-tree mst config
MST region name: my-mstp-region
Revision: 0
MSTI    VID
  1      100
  2      200-300
```

Modify Global Parameters

The root bridge sets the values for forward-delay, hello-time, max-age, and max-hops and overwrites the values set on other MSTP bridges.

- **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 MSTP 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 MST topology.
- **Max-hops** is the maximum number of hops a BPDU can travel before a receiving switch discards it.



Note: Dell Force10 recommends that only experienced network administrators change MSTP parameters. Poorly planned modification of MSTP parameters can negatively impact network performance.

To change MSTP parameters, use the following commands on the root bridge:

Task	Command Syntax	Command Mode
Change the forward-delay parameter. <ul style="list-style-type: none"> • Range: 4 to 30 • Default: 15 seconds 	forward-delay <i>seconds</i>	PROTOCOL MSTP

Task	Command Syntax	Command Mode
Change the hello-time parameter. Note: With large configurations (especially those with more ports) Dell Force10 recommends that you increase the hello-time. Range: 1 to 10 Default: 2 seconds	hello-time <i>seconds</i>	PROTOCOL MSTP
Change the max-age parameter. Range: 6 to 40 Default: 20 seconds	max-age <i>seconds</i>	PROTOCOL MSTP
Change the max-hops parameter. Range: 1 to 40 Default: 20	max-hops <i>number</i>	PROTOCOL MSTP

To view the current values for MSTP parameters, use the show running-config spanning-tree mstp command from EXEC privilege mode.

Figure 19-7. Viewing the Current Values for MSTP Parameters


```

FTOS(conf-mstp)#forward-delay 16
FTOS(conf-mstp)#exit
FTOS(conf)#do show running-config spanning-tree mstp
!
protocol spanning-tree mstp
no disable
name my-mstp-region
MSTI 1 VLAN 100
MSTI 2 VLAN 200-300
forward-delay 16
MSTI 2 bridge-priority 4096
FTOS(conf)#

```

Enable BPDU Filtering globally

The enabling of BPDU Filtering stops transmitting of BPDUs on the operational port fast enabled ports by default. When BPDUs are received, the spanning tree is automatically prepared. By default global bpdu filtering is disabled.

Figure 19-8. BPDU Filtering enabled globally

Task	Command Syntax	Command Mode
Enable BPDU Filter globally to filter transmission of BPDU port fast enabled interfaces.	edge-port bpdu filter default	PROTOCOL MSTP

Modify Interface Parameters

You can adjust two interface parameters to increase or decrease the probability that a port becomes a forwarding port:

- **Port cost** is a value that is based on the interface type. The greater the port cost, the less likely the port is selected to be a forwarding port.
- **Port priority** influences the likelihood that a port is selected to be a forwarding port in case that several ports have the same port cost.

Table 19-2 lists the default values for port cost by interface.

Table 19-2. MSTP Default Port Cost Values

Port Cost	Default Value
1000-Mb/s Ethernet interfaces	20000
40-Gigabit Ethernet interfaces	1400
10-Gigabit Ethernet interfaces	2000
Port Channel with one 10-Gigabit Ethernet interface	2000
Port Channel with one 40-Gigabit Ethernet interface	1400

Table 19-2. MSTP Default Port Cost Values

Port Cost	Default Value
Port Channel with two 10-Gigabit Ethernet interfaces	1800
Port Channel with two 40-Gigabit Ethernet interfaces	600

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 2000000 Default: refer to Table 19-2 .	<code>spanning-tree msti <i>number</i> cost <i>cost</i></code>	INTERFACE
Change the port priority of an interface. Range: 0 to 240, in increments of 16 Default: 128	<code>spanning-tree msti <i>number</i> priority <i>priority</i></code>	INTERFACE

To view the current values for these interface parameters, use the `show config` command from INTERFACE mode ([Figure 19-9](#)).

Configure an EdgePort

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. The `bpduguard shutdown-on-violation` option causes the interface hardware to be shutdown when it receives a BPDU. When you only implement `bpduguard`, although the interface is placed in an error disabled state when receiving the BPDU, the physical interface remains up and spanning-tree drops packets in the hardware after a BPDU violation. BPDUs are dropped in the software after receiving the BPDU violation. 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 you enable it on an interface connected to a network.

The enabling of BPDU Filtering stops sending and receiving of BPDUs on the port fast enabled ports. When both BPDU guard and BPDU filter are enabled on the port, BPDU filter takes the higher precedence. By default, BPDU filtering on an interface is disabled.

To enable EdgePort on an interface, use the following command:

Task	Command Syntax	Command Mode
Enable EdgePort on an interface.	<code>spanning-tree mstp edge-port [bpduguard [shutdown-on-violation] bpdupfilter]</code>	INTERFACE

To verify that EdgePort is enabled on a port, use the show config command from INTERFACE mode (Figure 19-9).



FTOS Behavior: Regarding bpduguard shutdown-on-violation behavior:

- 1 If the interface to be shutdown is a port channel, all the member ports are disabled in the hardware.
- 2 When a physical port is added to a port channel already in error disable state, the new member port is also disabled in the hardware.
- 3 When a physical port is removed from a port channel in error disable state, the error disabled state is cleared on this physical port (the physical port is enabled in the hardware).
- 4 You can clear the error disabled state with any of the following methods:
 - Perform a shutdown command on the interface.
 - Disable the shutdown-on-violation command on the interface (use the no spanning-tree mstp edge-port [bpduguard | [shutdown-on-violation]]) command).
 - Disable spanning tree on the interface (use the no spanning-tree command in INTERFACE mode).
- 5 Disabling global spanning tree (use the no spanning-tree command in CONFIGURATION mode).

Figure 19-9. Configuring EdgePort

```
FTOS(conf-if-te-3/41)#spanning-tree mstp edge-port
FTOS(conf-if-te-3/41)#show config
!
interface TenGigabitEthernet 3/41
no ip address
switchport
spanning-tree mstp edge-port
spanning-tree MSTI 1 priority 144
no shutdown
FTOS(conf-if-te-3/41)#
```

Flush MAC Addresses after a Topology Change

FTOS has an optimized MAC address flush mechanism for RSTP, MSTP, and PVST+ that flushes addresses only when necessary, which allows for faster convergence during topology changes. However, to activate the flushing mechanism defined by 802.1Q-2003, use the tc-flush-standard command, which flushes MAC addresses after every topology change notification. To view the enable status of this feature, use the show running-config spanning-tree mstp command from EXEC Privilege mode.

MSTP Sample Configurations

The running-configurations in Figure 19-11, Figure 19-12, and Figure 19-13 support the topology shown in Figure 19-10. The configurations are from FTOS systems. An MXL Switch system using FTOS, configured as shown in Figure 19-14, could be substituted for an FTOS router in this sample following topology and MSTP would function as designed.

Figure 19-10. MSTP with Three VLANs Mapped to Two Spanning Tree Instances

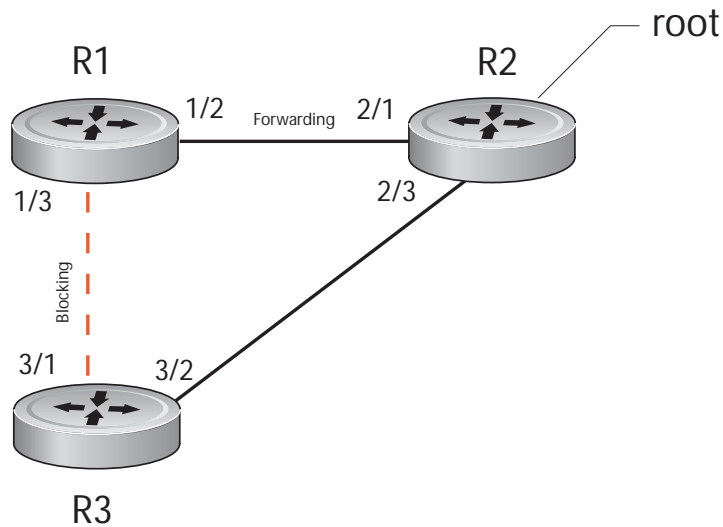


Figure 19-11. Router 1 Running-configuration

```

protocol spanning-tree mstp
no disable
name Tahiti
revision 123
MSTI 1 VLAN 100
MSTI 2 VLAN 200,300
!
interface TenGigabitEthernet 1/21
no ip address
switchport
no shutdown
!
interface TenGigabitEthernet 1/31
no ip address
switchport
no shutdown
!
interface Vlan 100
no ip address
tagged TenGigabitEthernet 1/21,31
no shutdown
!
interface Vlan 200
no ip address
tagged TenGigabitEthernet 1/21,31
no shutdown
!
interface Vlan 300
no ip address
tagged TenGigabitEthernet 1/21,31
no shutdown

```

Annotations for the configuration:

- Enable MSTP globally
Set Region Name and Revision
Map MSTP Instances to VLANs
- Assign Layer-2 interfaces to MSTP topology
- Create VLANs mapped to MSTP Instances
Tag interfaces to VLANs

Figure 19-12. Router 2 Running-configuration

```
protocol spanning-tree mstp
no disable
name Tahiti
revision 123
MSTI 1 VLAN 100
MSTI 2 VLAN 200,300
!
interface TenGigabitEthernet 2/11
no ip address
switchport
no shutdown
!
interface TenGigabitEthernet 2/31
no ip address
switchport
no shutdown
!
interface Vlan 100
no ip address
tagged TenGigabitEthernet 2/11,31
no shutdown
!
interface Vlan 200
no ip address
tagged TenGigabitEthernet 2/11,31
no shutdown
!
interface Vlan 300
no ip address
tagged TenGigabitEthernet 2/11,31
no shutdown
```

Enable MSTP globally
Set Region Name and Revision
Map MSTP Instances to VLANs

Assign Layer-2 interfaces
to MSTP topology

Create VLANs mapped to MSTP Instances
Tag interfaces to VLANs

Figure 19-13. Router 3 Running-configuration

```
protocol spanning-tree mstp
no disable
name Tahiti
revision 123
MSTI 1 VLAN 100
MSTI 2 VLAN 200,300
!
interface TenGigabitEthernet 3/11
no ip address
switchport
no shutdown
!
interface TenGigabitEthernet 3/21
no ip address
switchport
no shutdown
!
interface Vlan 100
no ip address
tagged TenGigabitEthernet 3/11,21
no shutdown
!
interface Vlan 200
no ip address
tagged TenGigabitEthernet 3/11,21
no shutdown
!
interface Vlan 300
no ip address
tagged TenGigabitEthernet 3/11,21
no shutdown
```

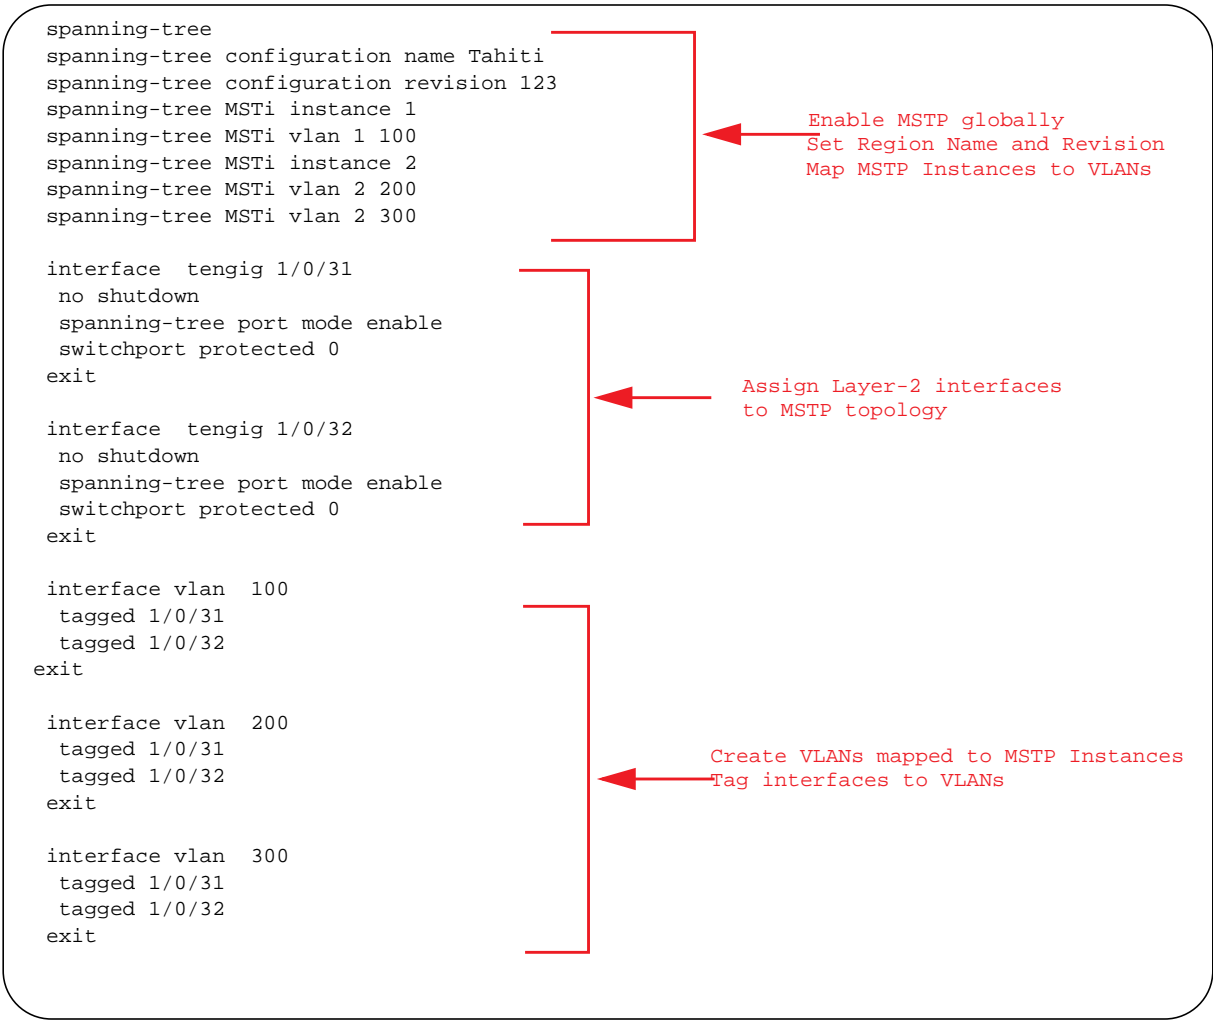
The diagram shows three red arrows pointing from explanatory text to specific configuration blocks in the code. The first arrow points to the global MSTP configuration (lines 1-5). The second arrow points to the Layer-2 interface configuration (lines 10-15). The third arrow points to the VLAN configuration (lines 20-30).

Enable MSTP globally
Set Region Name and Revision
Map MSTP Instances to VLANs

Assign Layer-2 interfaces
to MSTP topology

Create VLANs mapped to MSTP Instances
Tag interfaces to VLANs

Figure 19-14. FTOS Example Running-Configuration



Debugging and Verifying an MSTP Configuration

To display BPDUs, use the `debug spanning-tree mstp bpd` command from EXEC Privilege mode (Figure 19-15). To display MSTP-triggered topology change messages, use the `debug spanning-tree mstp events` command.

Figure 19-15. Displaying BPDUs and Events

```
FTOS#debug spanning-tree mstp bpd
1w1d17h : MSTP: Sending BPDU on TenGig 1/31 :
ProtId: 0, Ver: 3, Bpdu Type: MSTP, Flags 0x68
CIST Root Bridge Id: 32768:0001.e806.953e, Ext Path Cost: 20000
Regional Bridge Id: 32768:0001.e809.c24a, CIST Port Id: 128:384
Msg Age: 2, Max Age: 20, Hello: 2, Fwd Delay: 15, Ver1 Len: 0, Ver3 Len: 96
Name: my-mstp-region, Rev: 0, Int Root Path Cost: 20000
Rem Hops: 19, Bridge Id: 32768:0001.e80d.b6d6
E1200#1w1d17h : INST 1: Flags: 0x28, Reg Root: 32768:0001.e809.c24a, Int Root Co
    Brg/Port Prio: 32768/128, Rem Hops: 19
INST 2: Flags: 0x68, Reg Root: 4096:0001.e809.c24a, Int Root Cost: 20000
    Brg/Port Prio: 32768/128, Rem Hops: 19
[output omitted]
FTOS#debug spanning-tree mstp events
1w1d17h : MSTP: TC flag set in the incoming BPDU on port TenGig 1/31 for instance 0

1w1d17h : MSTP: TC flag set in the incoming BPDU on port TenGig 1/31 for instance 0

1w1d17h : MSTP: TC flag set in the incoming BPDU on port TenGig 1/31 for instance 0
```

Examine your individual routers to ensure all the necessary parameters match.

1. Region Name
2. Region Version
3. VLAN to Instance mapping

The `show spanning-tree mst` commands show various portions of the MSTP configuration. To view the overall MSTP configuration on the router, use the `show running-configuration spanning-tree mstp` command in EXEC Privilege mode (Figure 19-16).

To monitor and verify that the MSTP configuration is connected and communicating as desired, use the `debug spanning-tree mstp bpd` command (Figure 19-17).

Key items to look for in the debug report:

- MSTP flags indicate communication received from the same region.
 - In Figure 19-17, the output shows that the MSTP routers are located in the same region.
 - Does the debug log indicate that packets are coming from a “Different Region” (Figure 19-18)? If so, one of the key parameters is not matching.
- MSTP Region Name and Revision
 - The configured name and revisions *must* be identical among all the routers.
 - Is the Region name blank? That may mean that a name was configured on one router and but was not configured or was configured differently on another router (spelling and capitalization counts).

- MSTP Instances.
 - Use the show commands to verify the VLAN to MSTP instance mapping.
 - Are there “extra” MSTP Instances in the Sending or Received logs? That may mean that an additional MSTP instance was configured on one router but not the others.

Figure 19-16. Sample Output for show running-configuration spanning-tree mstp command

```
FTOS#show run spanning-tree mstp
!
protocol spanning-tree mstp
  name Tahiti
  revision 123
  MSTI 1 VLAN 100
  MSTI 2 VLAN 200,300
no disable
```

Figure 19-17. Displaying BPDUs and Events - Debug Log of Successful MSTP Configuration

```
FTOS#debug spanning-tree mstp bpdu
MSTP debug bpdu is ON
FTOS#
4w0d4h : MSTP: Sending BPDU on Tengig 2/21 :
ProtId: 0, Ver: 3, Bpdu Type: MSTP, Flags 0x6e
CIST Root Bridge Id: 32768:0001.e806.953e, Ext Path Cost: 0
Regional Bridge Id: 32768:0001.e806.953e, CIST Port Id: 128:470
Msg Age: 0, Max Age: 20, Hello: 2, Fwd Delay: 15, Ver1 Len: 0, Ver3 Len: 96
Name: Tahiti, Rev: 123, Int Root Path Cost: 0
Rem Hops: 20, Bridge Id: 32768:0001.e806.953e
4w0d4h : INST 1: Flags: 0x6e, Reg Root: 32768:0001.e806.953e, Int Root Cost: 0
      Brg/Port Prio: 32768/128, Rem Hops: 20
INST 2: Flags: 0x6e, Reg Root: 32768:0001.e806.953e, Int Root Cost: 0
      Brg/Port Prio: 32768/128, Rem Hops: 20

4w0d4h : MSTP: Received BPDU on Tengig 2/21 :
ProtId: 0, Ver: 3, Bpdu Type: MSTP, Flags 0x78 Same Region ← Indicates MSTP
CIST Root Bridge Id: 32768:0001.e806.953e, Ext Path Cost: 0 routers are in the
Regional Bridge Id: 32768:0001.e806.953e, CIST Port Id: 128:470 (single) region
Msg Age: 0, Max Age: 20, Hello: 2, Fwd Delay: 15, Ver1 Len: 0, Ver3 Len: 96
Name: Tahiti, Rev: 123, Int Root Path Cost: 0
Rem Hops: 19, Bridge Id: 32768:0001.e8d5.cbbd
4w0d4h : INST 1: Flags: 0x78, Reg Root: 32768:0001.e806.953e, Int Root Cost: 0
      Brg/Port Prio: 32768/128, Rem Hops: 19
INST 2: Flags: 0x78, Reg Root: 32768:0001.e806.953e, Int Root Cost: 0
      Brg/Port Prio: 32768/128, Rem Hops: 19

MSTP Instance
MSTP Region name
and revision
```

Figure 19-18. Displaying BPDUs and Events - Debug Log of Unsuccessful MSTP Configuration

```
4w0d4h : MSTP: Received BPDU on TenGig 2/21 :  
ProtId: 0, Ver: 3, Bpdu Type: MSTP, Flags 0x78  
CIST Root Bridge Id: 32768:0001.e806.953e, Ext Path Cost: 0  
Regional Bridge Id: 32768:0001.e806.953e, CIST Port Id: 128:470  
Msg Age: 0, Max Age: 20, Hello: 2, Fwd Delay: 15, Ver1 Len: 0, Ver  
Name: Tahiti, Rev: 123, Int Root Path Cost: 0  
Rem Hops: 20, Bridge Id: 32768:0001.e8d5.cbbd  
4w0d4h : INST 1: Flags: 0x70, Reg Root: 32768:0001.e8d5.cbbd, Int  
      Brg/Port Prio: 32768/128, Rem Hops: 20  
INST 2: Flags: 0x70, Reg Root: 32768:0001.e8d5.cbbd, Int Root Cost  
      Brg/Port Prio: 32768/128, Rem Hops: 20
```

Different Region

Indicates MSTP
routers are in
different regions and
are not communicating
with each other

Open Shortest Path First (OSPFv2)

This chapter includes the following topics:

- [Overview](#)
- [Implementing OSPF with FTOS](#)
 - [Fast Convergence \(OSPFv2, IPv4 only\)](#)
 - [Multi-Process OSPF \(OSPFv2, IPv4 only\)](#)
 - [RFC-2328 Compliant OSPF Flooding](#)
 - [OSPF ACK Packing](#)
 - [OSPF Adjacency with Cisco Routers](#)
- [Configuration Information](#)
 - [Configuration Task List for OSPFv2 \(OSPF for IPv4\)](#)
 - [Troubleshooting OSPFv2](#)
- [Sample Configurations for OSPFv2](#)

OSPF protocol standards are listed in the [Chapter 40, Standards Compliance](#) chapter.

Overview

Open shortest path first (OSPF) routing is a link-state routing protocol that calls for the sending of link-state advertisements (LSAs) to all other routers within the same autonomous system (AS) areas. Information on attached interfaces, metrics used, and other variables is included in OSPF LSAs. As OSPF routers accumulate link-state information, they use the shortest path first algorithm (SPF algorithm) to calculate the shortest path to each node.

OSPF routers initially exchange HELLO messages to set up adjacencies with neighbor routers. Use the HELLO process to establish adjacencies between routers of the AS. It is not required that every router within the autonomous system areas establish adjacencies. If two routers on the same subnet agree to become neighbors through the HELLO process, they begin to exchange network topology information in the form of LSAs.

In OSPFv2, neighbors on broadcast and NBMA links are identified by their interface addresses, while neighbors on other types of links are identified by RID.

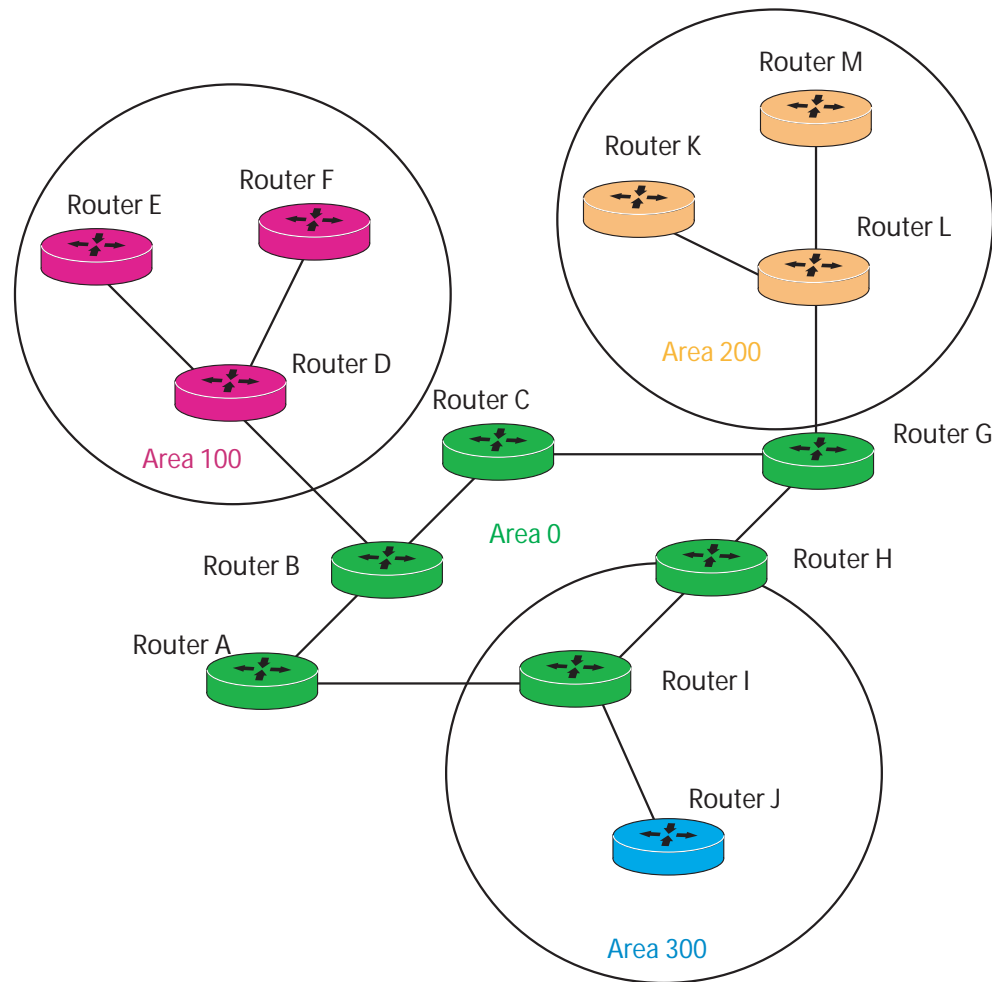
Autonomous System (AS) Areas

OSPF operates in a type of hierarchy. The largest entity within the hierarchy is the AS, which is a collection of networks under a common administration that share a common routing strategy (Figure 20-1). OSPF is an intra-AS (interior gateway) routing protocol, although it is capable of receiving routes from and sending routes to other ASs.

You can divide an AS into a number of areas, which are groups of contiguous networks and attached hosts. Routers with multiple interfaces can participate in multiple areas. These routers, area border routers (ABRs), maintain separate databases for each area. Areas are a logical grouping of OSPF routers identified by an integer or dotted-decimal number.

Areas allow you to further organize your routers within the AS. One or more areas are required within the AS. Areas are valuable in that they allow sub-networks to “hide” within the AS, thus minimizing the size of the routing tables on all routers. An area within the AS may not see the details of another area's topology. AS areas are known by their area number or the router's IP address.

Figure 20-1. Autonomous System Areas



Area Types

The **Backbone** of the network is Area 0. It is also called Area 0.0.0.0 and is the core of any AS. All other areas must connect to Area 0. Areas can be defined in such a way that the backbone is not contiguous. In this case, backbone connectivity must be restored through virtual links. Virtual links are configured between any backbone routers that share a link to a non-backbone area and function as if they were direct links.

An OSPF backbone is responsible for distributing routing information between areas. It consists of all area border routers, networks not wholly contained in any area, and their attached routers.

The backbone is the only area with a default area number. All other areas can have their Area ID assigned in the configuration.

Figure 20-1 shows that Routers A, B, C, G, H, and I are the backbone routers.

A **Stub Area (SA)** does not receive external route information, except for the default route. These areas do receive information from inter-area (IA) routes. Note that you must configure all routers within an assigned stub area as stubby so that they do not generate LSAs that do not apply. For example, a Type 5 LSA is intended for external areas and the stubby area routers may not generate external LSAs. Stubby areas cannot be traversed by a virtual link.

A **not-so-stubby area (NSSA)** can import AS external route information and send it to the backbone. It cannot receive external AS information from the backbone or other areas. It can be traversed by a virtual link.

Totally stubby areas are referred to as “no summary areas” in FTOS.

Networks and Neighbors

As a link-state protocol, OSPF sends routing information to other OSPF routers concerning the state of the links between them. The state (up or down) of those links is important.

Routers that share a link become neighbors on that segment. OSPF uses the hello protocol as a neighbor discovery and keep alive mechanism. After two routers are neighbors, they may proceed to exchange and synchronize their databases, which creates an adjacency.

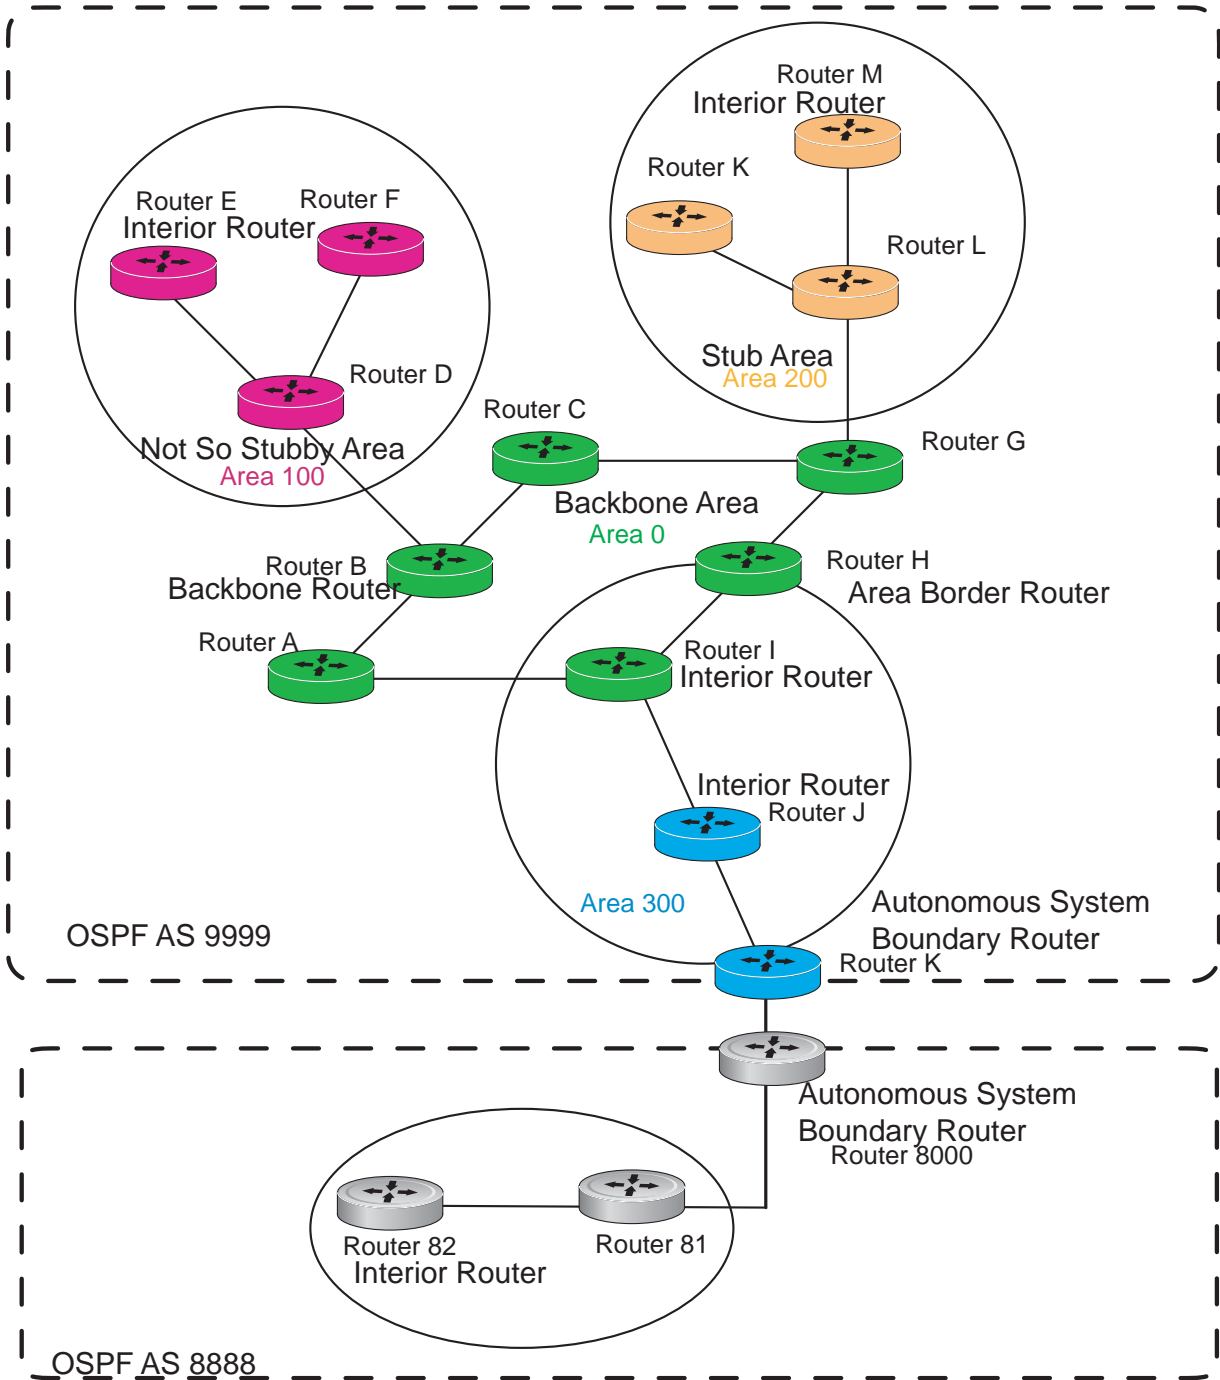
Router Types

Router types are attributes of the OSPF process. A given physical router may be a part of one or more OSPF processes. For example, a router connected to more than one area, receiving routing from a border gateway protocol (BGP) process connected to another AS acts as both an area border router and an AS router.

Each router has a unique ID, written in decimal format (A.B.C.D). The router ID does not have to be associated with a valid IP address. However, Dell Force10 recommends that the router ID and the router's IP address reflect each other, to make troubleshooting easier.

Figure 20-2 shows some examples of the different router designations.

Figure 20-2. OSPF Routing Examples



Backbone Router (BR)

A backbone router (BR) is part of the OSPF backbone, Area 0. This includes all ABRs. It can also include any routers that connect only to the backbone and another ABR, but are only part of Area 0, such as Router I in [Figure 20-2](#).

Area Border Router (ABR)

Within an AS, an area border router (ABR) connects one or more areas to the backbone. The ABR keeps a copy of the link-state database for every area it connects to, so it may keep multiple copies of the link state database. An ABR takes information it has learned on one of its attached areas and can summarize it before sending it out on other areas it is connected to. An ABR can connect to many areas in an AS, and is considered a member of each area it connects to.

Autonomous System Border Router (ASBR)

The autonomous system border area router (ASBR) connects to more than one AS and exchanges information with the routers in other ASs. In general, the ASBR connects to a non-interior gate protocol (IGP) such as BGP or uses static routes.

Internal Router (IR)

The internal router (IR) has adjacencies with ONLY routers in the same area, as Router E, M, and I are in [Figure 20-2](#).

Designated and Backup Designated Routers

OSPF elects a designated router (DR) and a backup designated router (BDR). Among other things, the designated router is responsible for generating LSAs for the entire multi-access network. Designated routers allow a reduction in network traffic and in the size of the topological database.

- The DR maintains a complete topology table of the network and sends the updates to the other routers using multicast. All routers in an area form a slave/master relationship with the DR. Every time a router sends an update, it sends it to the DR and BDR. The DR sends the update out to all other routers in the area.
- The BDR is the router that takes over if the DR fails.

Each router exchanges information with the DR and BDR. The DR and BDR relay the information to the other routers. On broadcast network segments, the number of OSPF packets is further reduced by the DR and BDR sending such OSPF updates to a multicast IP address that all OSPF routers on the network segment are listening on.

These router designations are not the same as the router IDs described earlier. The DR and BDR are configurable in FTOS. If no DR or BDR is defined in FTOS, the system assigns them. OSPF looks at the priority of the routers on the segment to determine which routers are the DR and BDR. The router with the highest priority is elected the DR. If there is a tie, the router with the higher router ID takes precedence. After the DR is elected, the BDR is elected the same way. A router with a router priority set to zero cannot become the DR or BDR.

Link-State Advertisements (LSAs)

A LSA communicates the router's local routing topology to all other local routers in the same area.

The LSA types supported by Dell Force10 are defined as follows:

- Type 1 - Router LSA
 - The router lists links to other routers or networks in the same area. Type 1 LSAs are flooded across their own area only. The Link-State ID of the Type 1 LSA is the originating router ID.
- Type 2 - Network LSA
 - The DR in an area lists which routers are joined together within the area. Type 2 LSAs are flooded across their own area only. The Link-State ID of the Type 2 LSA is the IP interface address of the DR.
- Type 3 - Summary LSA (OSPFv2)
 - An ABR takes information it has learned on one of its attached areas and can summarize it before sending it out on other areas it is connected to. The Link-State ID of the Type 3 LSA is the destination network number.
- Type 4 - AS Border Router Summary LSA (OSPFv2)
 - In some cases, Type 5 External LSAs are flooded to areas where the detailed next-hop information may not be available. An ABR flood the information for the router (for example, the Autonomous System Border Router [ASBR]) where the Type 5 advertisement originated. The Link-State ID for Type 4 LSAs is the router ID of the described ASBR.
- Type 5 - External LSA
 - These LSAs contain information imported into OSPF from other routing processes. They are flooded to all areas, except stub areas. The Link-State ID of the Type 5 LSA is the external network number.
- Type 7
 - Routers in a NSSA do not receive external LSAs from ABRs, but are allowed to send external routing information for redistribution. They use Type 7 LSAs to tell the ABRs about these external routes, which the ABR then translates to Type 5 external LSAs and floods as normal to the rest of the OSPF network.
- Type 9 - Link Local LSA (OSPFv2)
 - For OSPFv2, this is a link-local "opaque" LSA as defined by RFC2370.

For all LSA types, there are 20-byte LSA headers. One of the fields of the LSA header is the Link-State ID.

Each router link is defined as one of four types: type 1, 2, 3, or 4. The LSA includes a link ID field that identifies, by the network number and mask, the object this link connects to.

Depending on the type, the link ID has different meanings.

- 1: point-to-point connection to another router neighboring router
- 2: connection to a transit network IP address of Designated Router
- 3: connection to a stub network IP network/subnet number
- 4: virtual link neighboring router ID

LSA Throttling

LSA throttling provides configurable interval timers to improve OSPF convergence times. The default OSPF static timers (5 seconds for transmission, 1 second for acceptance) ensure sufficient time for sending and resending LSAs and for system acceptance of arriving LSAs. However, some networks may require reduced intervals for LSA transmission and acceptance. The throttling timers allow for this improved convergence times.

Configure the LSA throttling timers in milliseconds, with the interval time increasing exponentially until a maximum time has been reached. If the maximum time is reached, the system continues to transmit at the max-interval until twice the max-interval time has passed. At that point, the system reverts to the start-interval timer and the cycle begins again.

When you configure the LSA throttle timers, syslog messages appear, indicating the interval times ([Message 1](#)) and ([Message 2](#)).

Message 1 SYSLOG message for LSA transmit timer (45000 msec in this example)

```
Mar 15 09:46:00: %STKUNIT0-M:CP %OSPF-4-LSA_BACKOFF: OSPF Process 10,Router lsa id 2.2.2.2 router-id 2.2.2.2 is backed off to transmit after 45000ms
```

Message 2 SYSLOG message for LSA arrival timer (1000 msec in this example)


```
Mar 15 09:46:06: %STKUNIT0-M:CP %OSPF-4-LSA_BACKOFF: OSPF Process 10,Router lsa id 3.3.3.3 rtrid 3.3.3.3 received before 1000ms time
```

Router Priority and Cost

Router priority and cost is the method the system uses to “rate” the routers ([Figure 20-3](#)). For example, if not assigned, the system selects the router with the highest priority as the DR. The second highest priority is the BDR.

Priority is a numbered rating 0 to 255. The higher the number, the higher the priority.

Cost is a numbered rating 1 to 65535. The higher the number, the greater the cost. The cost assigned reflects the cost should the router fail. When a router fails and the cost is assessed, a new priority number results.

Figure 20-3. Priority and Costs Example

Router 1 selected by the system as DR.
Router 2 selected by the system as BDR.

If R1 fails, the system subtracts 21 from R1's priority number. R1's new priority is 179.

R2 as both the selected BDR and the now-highest priority, becomes the DR.

If R3 fails, the system subtracts 50 from its priority. R2's new priority is 130.

R4 is now the highest priority and becomes the DR.

Implementing OSPF with FTOS

FTOS supports up to 10,000 OSPF routes. Within that 10,000, you can designate up to 8,000 routes as external and up to 2,000 designated as inter/intra area routes.

FTOS version 7.8.1.0 and later supports multiple OSPF processes (OSPF MP). The MXL Switch supports up to 16 processes simultaneously.

FTOS supports SA, Totally Stub (No Summary), and NSSAs and supports the following LSAs, as described earlier in this document.

- 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)

Fast Convergence (OSPFv2, IPv4 only)

Fast convergence allows you to define the speeds at which LSAs are originated and accepted, and reduce OSPFv2 end-to-end convergence time. FTOS allows you to accept and originate LSAs as soon as they are available to speed up route information propagation.



Note: The faster the convergence, the more frequent the route calculations and updates. This impacts CPU utilization and may impact adjacency stability in larger topologies.

Multi-Process OSPF (OSPFv2, IPv4 only)

Multi-process OSPF is supported on OSPFv2 with IPv4 only.

Multi-process OSPF allows multiple OSPFv2 processes on a single router. Multiple OSPFv2 processes allow for isolating routing domains, supporting multiple route policies and priorities in different domains, and creating smaller domains for easier management.

- The MXL Switch supports up to 16 OSPFv2 processes.

Each OSPFv2 process has a unique process ID and must have an associated router ID. There must be an equal number of interfaces in Layer-3 mode for the number of processes created. For example, if five OSPFv2 processes are created on a system, there must be at least five interfaces assigned in Layer-3 mode.

Each OSPFv2 process is independent. If one process loses adjacency, the other processes continue to function.

Processing SNMP and Sending SNMP Traps

Though there are may be several OSPFv2 processes, only one process can process SNMP requests and send SNMP traps. The `mib-binding` command identifies one of the OSPFv2 processes as the process responsible for SNMP management. If the `mib-binding` command is not specified, the first OSPFv2 process created manages the SNMP processes and traps.

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 you require the RFC 2328 flooding behavior, enable it by using the `flood-2328` command in ROUTER OSPF mode. When you enable RFC 2328 flooding, this command configures FTOS to flood LSAs on all interfaces.

To confirm RFC 2328 flooding behavior, use `debug ip ospf packet` command and look for output similar to the following (Figure 20-4).

Figure 20-4. Enabling RFC-2328 Compliant OSPF Flooding

```
00:10:41 : OSPF(1000:00):
Rcv. v:2 t:5(LSAck) l:64 Acks 2 rid:2.2.2.2 ← Printed only for ACK packets
    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:TenGig 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 FTOS version 7.5.1.0, to confirm that RFC-2328 compliant OSPF flooding is enabled, use the `show ip ospf` command (Figure 20-5).

Figure 20-5. Enabling RFC-2328 Compliant OSPF Flooding

```
FTOS#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 SPF's 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 is non-configurable.

OSPF Adjacency with Cisco Routers

To establish an OSPF adjacency between Dell Force10 and Cisco routers, the hello interval and dead interval must be the same on both routers. In FTOS, the OSPF dead interval value is, by default, set to 40 seconds, and is independent of the OSPF hello interval. Configuring a hello interval does not change the dead interval in FTOS. In contrast, the OSPF dead interval on a Cisco router is, by default, four times as long as the hello interval. Changing the hello interval on the Cisco router automatically changes the dead interval as well.

To ensure equal intervals between the routers, manually set the dead interval of the Dell Force10 router to match the Cisco configuration. Use the `ip ospf dead-interval <x>` command in INTERFACE mode (Figure 20-6) and (Figure 20-7).

Figure 20-6. Command Example for ip ospf intervals

```
FTOS(conf)#int tengig- 2/2
FTOS(conf-if-te-2/2)#ip ospf hello-interval 20
FTOS(conf-if-te-2/2)#ip ospf dead-interval 80
FTOS(conf-if-te-2/2)#
```

← Dead Interval Set at 4x Hello Interval

Figure 20-7. OSPF Configuration with intervals set

```
FTOS (conf-if-te-2/2)#ip ospf dead-interval 20
FTOS (conf-if-te-2/2)#do show ip os int tengig 1/3
TenGigabitEthernet 2/2 is up, line protocol is up
Internet Address 20.0.0.1/24, Area 0
Process ID 10, Router ID 1.1.1.2, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 1.1.1.2, Interface address 30.0.0.1
Backup Designated Router (ID) 1.1.1.1, Interface address 30.0.0.2
Timer intervals configured, Hello 20, Dead 80, Wait 20, Retransmit 5
Hello due in 00:00:04
Neighbor Count is 1, Adjacent neighbor count is 1
Adjacent with neighbor 1.1.1.1 (Backup Designated Router)
FTOS (conf-if-te-2/2)#
```

← Dead Interval Set at 4x Hello Interval

For more information regarding this functionality or for assistance, go to www.force10networks.com/support.

Configuration Information

The interfaces must be in Layer-3 mode (assigned an IP address) and enabled so that they can send and receive traffic. The OSPF process must know about these interfaces. To make the OSPF process aware of these interfaces, they must be assigned to OSPF areas.

You must configure OSPF GLOBALLY on the system in CONFIGURATION mode.

To assign OSPF features and functions to each router, use the CONFIG-INTERFACE commands for each interface.



Note: By default, OSPF is disabled.

Configuration Task List for OSPFv2 (OSPF for IPv4)

Configuration takes three steps:

1. Configure a physical interface. Assign an IP address, physical or loopback, to the interface to enable Layer 3 routing.
2. Enable OSPF globally. Assign network area and neighbors.
3. Add interfaces or configure other attributes.

The following configuration steps include two mandatory steps and several optional ones:

- [Enable OSPFv2 \(mandatory\)](#)
- [Enable Multi-Process OSPF](#)
- [Assign an OSPFv2 area \(mandatory\)](#)
- [Enable OSPFv2 on Interfaces](#)
- [Configure Stub Areas](#)
- [Configure LSA Throttling Timers](#)
- [Enable Passive Interfaces](#)
- [Enable Fast-Convergence](#)
- [Change OSPFv2 Parameters on Interfaces](#)
- [Enable OSPFv2 Authentication](#)
- [Redistribute Routes](#)
- [Troubleshooting OSPFv2](#)

For a complete listing of all commands related to OSPFv2, refer to the OSPF section in the *FTOS Command Line Interface* document.

Enable OSPFv2

Assign an IP address to an interface (physical or loopback) to enable Layer 3 routing. By default, OSPF, like all routing protocols, is disabled.

Before enabling OSPFv2 globally, you *must* configure at least one interface for Layer 3.

If implementing multi-process OSPF, you must create an equal number of Layer 3-enabled interfaces and OSPF Process IDs. For example, if you create 4 OSPFv2 process IDs, you must have four interfaces with Layer 3 enabled.

To enable OSPFv2 routing, follow these steps.

Step	Command Syntax	Command Mode	Usage
1	<code>ip address <i>ip-address mask</i></code> If using a loopback interface, refer to Loopback Interfaces on page 230 .	CONFIG-INTERFACE	Assign an IP address to an interface. Format: A.B.C.D/M
2	<code>no shutdown</code>	CONFIG-INTERFACE	Enable the interface.

To enable the OSPF process, return to CONFIGURATION mode. The OSPF process ID is the identifying number assigned to the OSPF process and the Router ID is the IP address associated with the OSPF process.

Command Syntax	Command Mode	Usage
<code>router ospf <i>process-id</i></code>	CONFIGURATION	Enable the OSPFv2 process globally. Range: 1-65535

If you try to enter an OSPF process ID, or if you try to enable more OSPF processes than available Layer 3 interfaces, prior to assigning an IP address to an interface and setting the `no shutdown` command, you will see the following message.

Message 3

```
FTOS(conf)#router ospf 1
% Error: No router ID available.
```

In CONFIGURATION ROUTER OSPF mode, assign the router ID. The router ID is not required to be the router's IP address. Dell Force10 recommends using the IP address as the router ID for easier management and troubleshooting.

Command Syntax	Command Mode	Usage
<code>router-id <i>ip address</i></code>	CONFIG-ROUTER-OSPF-id	Assign the Router ID for the OSPFv2 process. IP Address: A.B.C.D

To disable OSPF, use the `no router ospf process-id` command in CONFIGURATION mode.

To reset the OSPFv2 process, use the `clear ip ospf process-id` command in EXEC Privilege mode.

To view the current OSPFv2 status, use the `show ip ospf process-id` command in EXEC mode ([Figure 18-8](#)).

Figure 20-8. show ip ospf process id Command Example

```

FTOS#show ip ospf 55555
Routing Process ospf 55555 with ID 10.10.10.10
Supports only single TOS (TOS0) routes
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Number of area in this router is 0, normal 0 stub 0 nssa 0
FTOS#

```

Enable Multi-Process OSPF

Multi-process OSPF allows multiple OSPFv2 processes on a single router. The following list shows the number of processes supported on each platform type.

- The MXL Switch supports up to 16 OSPFv2 processes.

When configuring a single OSPF process, follow the same steps described above. Repeat them as often as necessary for the desired number of processes. After you create the process, all other configurations apply as usual. To enable multi-process OSPF, follow these steps:

Step	Command Syntax	Command Mode	Usage
1	ip address <i>ip-address mask</i>	CONFIG-INTERFACE	Assign an IP address to an interface. Format: A.B.C.D/M If using a loopback interface, refer to Loopback Interfaces .
2	no shutdown	CONFIG-INTERFACE	Enable the interface.

To enable the OSPF process, return to CONFIGURATION mode. The OSPF process ID is the identifying number assigned to the OSPF process. The Router ID is the IP address associated with the OSPF process.

Command Syntax	Command Mode	Usage
router ospf <i>process-id</i>	CONFIGURATION	Enable the OSPFv2 process globally. Range: 1-65535

If you try to enable more OSPF processes than available Layer 3 interfaces, you will see the following message.

Message 4

```

FTOS(conf)#router ospf 1
% Error: No router ID available.

```

In CONFIGURATION ROUTER OSPF mode, assign the Router ID. The Router ID is not required to be the router's IP address. For easier management and troubleshooting, Dell Force10 recommends using the IP address as the Router ID.

Command Syntax	Command Mode	Usage
<code>router-id ip address</code>	CONFIG-ROUTER-OSPF-id	Assign the Router ID for the OSPFv2 process. IP Address: A.B.C.D

To disable OSPF, use the `no router ospf process-id` command in CONFIGURATION mode.

To reset the OSPFv2 process, use the `clear ip ospf process-id` command syntax in EXEC Privilege mode.

Assign an OSPFv2 area

After you enable OSPFv2, assign the interface to an OSPF area. To set up OSPF areas and enable OSPFv2 on an interface, use the `network` command.

You must have at least one AS area: Area 0. This is the backbone area. If your OSPF network contains more than one area, you must also configure a backbone area (Area ID 0.0.0.0). Any area besides Area 0 can have any number ID assigned to it.

The OSPFv2 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 10.0.0.0 /8, you cannot assign the network address 10.1.0.0 /16 because 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 OSPFv2.

To set up each neighbor and OSPF area, use the following command in CONFIGURATION ROUTER OSPF mode. The you can assign the area by a number or with an IP interface address.

Command Syntax	Command Mode	Usage
<code>network ip-address mask area area-id</code>	CONFIG-ROUTER-OSPF-id	Enable OSPFv2 on an interface and assign an network address range to a specific OSPF area. IP Address Format: A.B.C.D/M Area ID Range: 0-65535 or A.B.C.D/M

Enable OSPFv2 on Interfaces

Each interface must have OSPFv2 enabled. It must be configured for Layer 3 protocol and not be shutdown. OSPFv2 can also be assigned to a loopback interface as a virtual interface.

OSPF functions and features, such as MD5 Authentication, Grace Period, Authentication Wait Time, etc., are assigned on a per interface basis.



Note: If using features like MD5 Authentication, ensure all the neighboring routers are also configured for MD5.

Figure 20-9 shows an example of assigning an IP address to an interface and then assigning an OSPFv2 area that includes that Layer-3 interface's IP address.

Figure 20-9. Configuring an OSPF Area Example

```
FTOS#(conf)#int tengig 4/44
FTOS(conf-if-te-4/44)#ip address 10.10.10.10/24
FTOS(conf-if-te-4/44)#no shutdown
FTOS(conf-if-te-4/44)#ex
FTOS(conf)#router ospf 1
FTOS(conf-router_ospf-1)#network 1.2.3.4/24 area 0
FTOS(conf-router_ospf-1)#network 10.10.10.10/24 area 1
FTOS(conf-router_ospf-1)#network 20.20.20.20/24 area 2
FTOS(conf-router_ospf-1)#
FTOS#
```

Assign Layer-3 interface with IP Address and no shutdown

Assign interface's IP Address to an Area

Dell Force10 recommends that the OSPFv2 Router ID be the interface IP addresses for easier management and troubleshooting.

To view the configuration, use the show config command in CONFIGURATION 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. To view the interfaces currently active and the areas assigned to the interfaces, use the show ip ospf interface command (Figure 20-10).

Figure 20-10. show ip ospf process-id interface Command Example

```
FTOS>show ip ospf 1 interface

TenGigabitEthernet 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

TenGigabitEthernet 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)
FTOS>
```

Loopback interfaces also assist in the OSPF process. OSPF picks the highest interface address as the router-id and a loopback interface address has a higher precedence than other interface addresses.

Figure 20-11 shows the show ip ospf process-id interface command with a loopback interface.

Figure 20-11. show ip ospf process-id interface Command Example

```

FTOS#show ip ospf 1 int

TenGigabitEthernet 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.
FTOS#

```

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; the 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 OSPFv2 network, never configure the backbone area as a stub area.

To configure a stub area, follow these steps, starting in EXEC Privilege mode.

Step	Command Syntax	Command Mode	Usage
1	show ip ospf <i>process-id</i> database database-summary	EXEC Privilege	Review all areas after they were configured to determine which areas are NOT receiving type 5 LSAs.
2	configure	EXEC Privilege	Enter CONFIGURATION mode.
3	router ospf <i>process-id</i>	CONFIGURATION	Enter ROUTER OSPF mode. Process ID is the ID assigned when configuring OSPFv2 globally.
4	area <i>area-id</i> stub [no-summary]	CONFIG-ROUTER-O SPF-id	Configure the area as a stub area. To prevent transmission to the area of summary ASBR LSAs., use the keywords no-summary. Area ID is the number or IP address assigned when creating the Area.

To view which LSAs are transmitted, use the `show ip ospf database process-id database-summary` command syntax in EXEC Privilege mode (Figure 20-12).

Figure 20-12. show ip ospf process-id database database-summary Command Example

```
FTOS#show ip ospf 34 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
FTOS#
```

To view information on areas, use the `show ip ospf process-id` command in EXEC Privilege mode.

Configure LSA Throttling Timers

Configured LSA timers replace the standard transmit and acceptance times for LSAs. Configure LSA throttling timers in milliseconds, with the interval time increasing exponentially until a maximum time has been reached. If the maximum time is reached, the system continues to transmit at the max-interval. If the system is stable for twice the maximum interval time, the system reverts to the start-interval timer and the cycle begins again. To configure the LSA throttling timers, use the commands below.

Command Syntax	Command Mode	Usage
<code>timers throttle lsa all</code> {start-interval hold-interval max-interval}	CONFIG-ROUTER-OSPF-id	Specify the interval times for all LSA transmissions <ul style="list-style-type: none"> start-interval: Set the minimum interval between initial sending and resending the same LSA. Range: 0-600,000 milliseconds hold-interval: Set the next interval to send the same LSA. This is the time between sending the same LSA after the start-interval has been attempted. Range: 1-600,000 milliseconds max-interval: Set the maximum amount of time the system waits before sending the LSA. Range: 1-600,000 milliseconds
<code>timers throttle lsa arrival</code> <i>arrival-time</i>	CONFIG-ROUTER-OSPF-id	Specify the interval for LSA acceptance. <ul style="list-style-type: none"> arrival-time: Set the interval between receiving the same LSA repeatedly, to allow sufficient time for the system to accept the LSA. Range: 0-600,000 milliseconds

Enable Passive Interfaces

A passive interface is one that does not send or receive routing information. Enabling passive interface suppresses routing updates on an interface. Although the passive interface will neither send nor receive routing updates, the network on that interface is still included in OSPF updates sent using other interfaces.

To suppress the interface's participation on an OSPF interface, use the following command in ROUTER OSPF mode. This command stops the router from sending updates on that interface.

Command Syntax	Command Mode	Usage
passive-interface {default interface}	CONFIG-ROUTER-OSPF-id	<p>Specify whether all or some of the interfaces will be passive.</p> <p>Default enabled passive interfaces on ALL interfaces in the OSPF process.</p> <p>Entering the physical interface type, slot, and number enable passive interface on only the identified interface.</p> <ul style="list-style-type: none"> • For a port channel, enter the keyword <code>port-channel</code> followed by a number from 1 to 128 • For a 10-Gigabit Ethernet interface, enter the keyword <code>TenGigabitEthernet</code> followed by the slot/port information (for example, <code>passive-interface ten 2/3</code>). • For a VLAN, enter the keyword <code>vlan</code> followed by a number from 1 to 4094 (for example, <code>passive-interface vlan 2222</code>). • For a 40-Gigabit Ethernet interface, enter the keyword <code>fortyGigE</code> followed by the slot/port information. <p>The default keyword sets all interfaces on this OSPF process as passive. To remove the passive interface from select interfaces, use the <code>no passive-interface interface</code> command while passive interface default is configured.</p>

To enable both receiving and sending routing updates, use the `no passive-interface interface` command.

When you configure a passive interface, the `show ip ospf process-id interface` command adds the words "passive interface" to indicate that hello packets are not transmitted on that interface (Figure 20-13).

Figure 20-13. show ip ospf process-id interface Command Example

```

FTOS#show ip ospf 34 int

TenGigabitEthernet 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

TenGigabitEthernet 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) ← Interface is not running the
    Neighbor Count is 0, Adjacent neighbor count is 0           OSPF protocol.

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.
FTOS#

```

Enable Fast-Convergence

The fast-convergence CLI sets the minimum origination and arrival LSA parameters to zero (0), allowing rapid route calculation. When you disable fast-convergence, origination and arrival LSA parameters are set to 5 seconds and 1 second, respectively.

Setting the convergence parameter (1-4) indicates the actual convergence level. Each convergence setting adjusts the LSA parameters to zero, but the fast-convergence parameter setting allows for even finer tuning of the convergence speed. The higher the number, the faster the convergence.

To enable or disable fast-convergence, use the following command in ROUTER OSPF mode.


Command Syntax	Command Mode	Usage
<code>fast-convergence {number}</code>	CONFIG-ROUTER-OSPF-id	Enable OSPF fast-convergence and specify the convergence level. Parameter: 1-4 The higher the number, the faster the convergence. When disabled, the parameter is set at 0 (Figure 20-15).
		Note: A higher convergence level can result in occasional loss of OSPF adjacency. Generally, convergence level 1 meets most convergence requirements. Higher convergence levels should only be selected following consultation with Dell Force10 technical support.

Figure 20-14 shows the convergence settings when you enable fast-convergence. Figure 20-15 shows settings when you disable fast-convergence. To view these settings, use the show ip ospf command.

Figure 20-14. show ip ospf process-id (Fast-Convergence Enabled) Command Example

```
FTOS(conf-router_ospf-1)#fast-converge 2
FTOS(conf-router_ospf-1)#exit
FTOS(conf)#exit
FTOS#show ip ospf 1
Routing Process ospf 1 with ID 192.168.67.2
Supports only single TOS (TOS0) routes
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Convergence Level 2
Min LSA origination 0 msec, Min LSA arrival 0 msec
Min LSA hold time 0 msec, Max LSA wait time 0 msec
Number of area in this router is 0, normal 0 stub 0 nssa 0
FTOS#
```

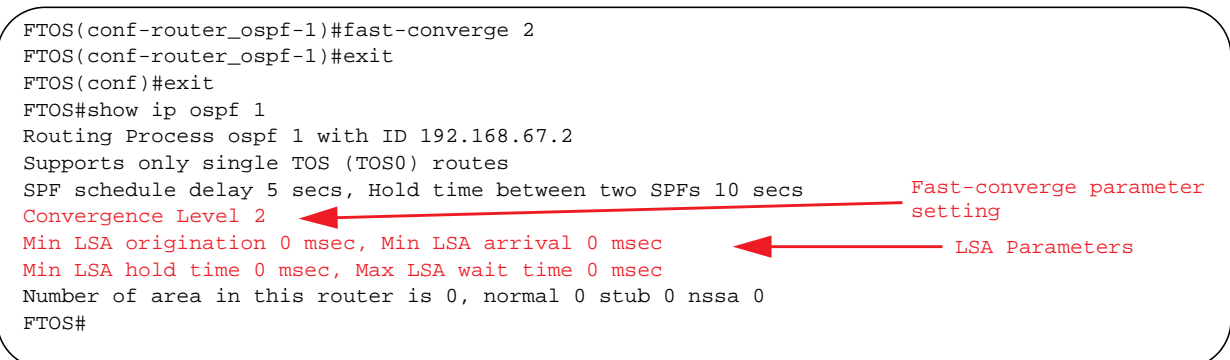
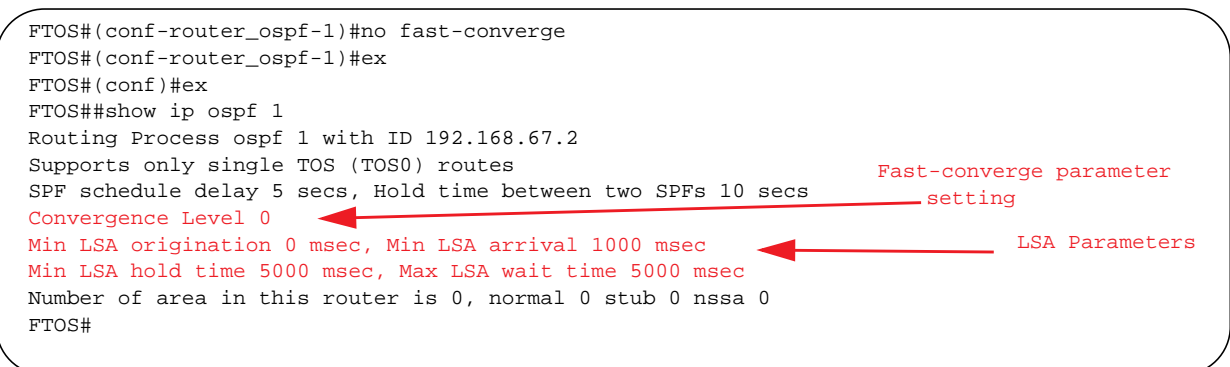


Figure 20-15. show ip ospf process-id (Fast-Convergence Disabled) Command Example

```
FTOS#(conf-router_ospf-1)#no fast-converge
FTOS#(conf-router_ospf-1)#ex
FTOS#(conf)#ex
FTOS##show ip ospf 1
Routing Process ospf 1 with ID 192.168.67.2
Supports only single TOS (TOS0) routes
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Convergence Level 0
Min LSA origination 0 msec, Min LSA arrival 1000 msec
Min LSA hold time 5000 msec, Max LSA wait time 5000 msec
Number of area in this router is 0, normal 0 stub 0 nssa 0
FTOS#
```



Change OSPFv2 Parameters on Interfaces

In FTOS, you can modify the OSPF settings on the interfaces. Some interface parameter values must be consistent across all interfaces to avoid routing errors. For example, to prevent misconfiguration of OSPF neighbors, you must set the same time interval for the hello packets on all routers in the OSPF network.

To change OSPFv2 parameters on the interfaces, use any or all of the following commands in CONFIGURATION INTERFACE mode.

Command Syntax	Command Mode	Usage
<code>ip ospf cost</code>	CONFIG-INTERFACE	Change the cost associated with OSPF traffic on the interface. Cost: 1 to 65535 (default depends on the interface speed).
<code>ip ospf dead-interval seconds</code>	CONFIG-INTERFACE	Change the time interval the router waits before declaring a neighbor dead. Configure Seconds range: 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.
<code>ip ospf hello-interval seconds</code>	CONFIG-INTERFACE	Change the time interval between the hello-packet transmission. Seconds range: from 1 to 65535 (default is 10 seconds). The hello interval must be the same on all routers in the OSPF network.
<code>ip ospf message-digest-key keyid md5 key</code>	CONFIG-INTERFACE	Use the MD5 algorithm to produce a message digest or key, which is sent instead of the key. Keyid range: 1 to 255 Key: a character string Be sure to write down or otherwise record the Key. You cannot learn the key once it is configured. You must be careful when changing this key.
<code>ip ospf priority number</code>	CONFIG-INTERFACE	Change the priority of the interface, which is used to determine the designated router for the OSPF broadcast network. Number range: 0 to 255 (the default is 1).
<code>ip ospf retransmit-interval seconds</code>	CONFIG-INTERFACE	Change the retransmission interval between the LSAs. Seconds range: from 1 to 65535 (default is 5 seconds). The retransmit interval must be the same on all routers in the OSPF network.
<code>ip ospf transmit-delay seconds</code>	CONFIG-INTERFACE	Change the wait period between the link state update packets sent out the interface. Seconds range: from 1 to 65535 (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 CONFIGURATION INTERFACE mode (Figure 20-16). To view the interface status in the OSPF process, use the show ip ospf interface command in EXEC mode.

Figure 20-16. Changing the OSPF Cost Value on an Interface

```

FTOS(conf-if)#ip ospf cost 45
FTOS(conf-if)#show config
!
interface TenGigabitEthernet 0/0
ip address 10.1.2.100 255.255.255.0
no shutdown
ip ospf cost 45
FTOS(conf-if)#end
FTOS#show ip ospf 34 interface

TenGigabitEthernet 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
FTOS#

```

The change is made on the interface and it is reflected in the OSPF configuration

Enable OSPFv2 Authentication

To enable or change various OSPF authentication parameters, use the following commands in CONFIGURATION INTERFACE mode.:

Command Syntax	Command Mode	Usage
ip ospf authentication-key <i>key</i>	CONFIG-INTERFACE	Set the clear text authentication scheme on the interface. Configure a <i>key</i> that is a text string no longer than eight characters. All neighboring routers must share the same password to exchange OSPF information.
ip ospf auth-change-wait-time <i>seconds</i>	CONFIG-INTERFACE	Set the authentication change wait time in <i>seconds</i> 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.

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 CONFIGURATION PREFIX LIST mode prior to assigning it to the OSPF process. To do this, use the following commands.

Command Syntax	Command Mode	Usage
<code>ip prefix-list <i>prefix-name</i></code>	CONFIGURATION	Create a prefix list and assign it a unique name. You are in PREFIX LIST mode.
<code>seq <i>sequence-number</i> {deny permit} <i>ip-prefix</i> [ge <i>min-prefix-length</i>] [le <i>max-prefix-length</i>]</code>	CONFIG- PREFIX LIST	Create a prefix list with a sequence number and a deny or permit action. The optional parameters are: ge <i>min-prefix-length</i> : is the minimum prefix length to be matched (0 to 32). le <i>max-prefix-length</i> : is the maximum prefix length to be matched (0 to 32).

For configuration information on prefix lists, refer to the *IP Access Control Lists, Prefix Lists, and Route-maps* chapter in the *FTOS Configuration Guide*.

To apply prefix lists to incoming or outgoing OSPF routes, use the following commands in CONFIGURATION-ROUTER OSPF mode.

Command Syntax	Command Mode	Usage
<code>distribute-list <i>prefix-list-name</i> in [<i>interface</i>]</code>	CONFIG-ROUTER-OSPF-id	Apply a configured prefix list to incoming OSPF routes.
<code>distribute-list <i>prefix-list-name</i> out [connected ospf <process-id> rip static]</code>	CONFIG-ROUTER-OSPF-id	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, you can include router information protocol (RIP), static, or directly connected routes in the OSPF process.

To redistribute routes, use the following command in CONFIGURATION- ROUTER-OSPF mode.

Command Syntax	Command Mode	Usage
<code>redistribute {bgp connected rip ospf <process-id> static} [metric <i>metric-value</i> metric-type <i>type-value</i>] [route-map <i>map-name</i>] [tag <i>tag-value</i>]</code>	CONFIG-ROUTER-OSPF-id	<p>Specify which routes will be redistributed into OSPF process. Configure the following required and optional parameters:</p> <ul style="list-style-type: none">• <code>bgp</code>, <code>connected</code>, <code>ospf <process-id></code>, or <code>static</code>: enter one of the keywords to redistribute those routes.• <code>metric <i>metric-value</i></code> range: 0 to 4294967295.• <code>metric-type <i>metric-type</i></code>: 1 for OSPF external route type 1 or 2 for OSPF external route type 2.• <code>route-map <i>map-name</i></code>: enter a name of a configured route map.• <code>tag <i>tag-value</i></code> range: 0 to 4294967295.

To view the current OSPF configuration, use the `show running-config ospf` command in EXEC mode or the `show config` command in ROUTER OSPF mode.

Figure 20-17. show config Command Example

```
FTOS(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
FTOS(conf-router_ospf)#
```

Troubleshooting OSPFv2

FTOS has several tools to make troubleshooting easier. Be sure to check the following, as these are typical issues that interrupt an OSPFv2 process. This is not a comprehensive list, just some examples of typical troubleshooting checks:

- Has OSPF been enabled globally?
- Is the OSPF process active on the interface?
- Are adjacencies established correctly?
- Are the interfaces configured for Layer 3 correctly?
- Is the router in the correct area type?
- Have the routes been included in the OSPF database?
- Have the OSPF routes been included in the routing table (not just the OSPF database)?

Some useful troubleshooting commands are:

- show interfaces
- show protocols
- debug IP OSPF events and/or packets
- show neighbors
- show routes

To view the state of all the enabled OSPFv2 processes, use the `show running-config ospf` command (Figure 20-18).

Command Syntax	Command Mode	Usage
<code>show running-config ospf</code>	EXEC Privilege	View the summary of all OSPF process IDs enables on the router.

Figure 20-18. show running-config ospf Command Example

```
FTOS#show run ospf
!
router ospf 3
!
router ospf 4
  router-id 4.4.4.4
  network 4.4.4.0/28 area 1
!
router ospf 5
!
router ospf 6
!
router ospf 7
  mib-binding
!
router ospf 8
!
  default-information originate always
  router-id 10.10.10.10
FTOS#
```

To get general route and links status information, use the following commands in EXEC Privilege mode.

Command Syntax	Command Mode	Usage
show ip route summary	EXEC Privilege	View the summary information of the IP routes
show ip ospf database	EXEC Privilege	View the summary information for the OSPF database

To view the OSPFv2 configuration for a neighboring router, use the following command in EXEC Privilege mode.

Command Syntax	Command Mode	Usage
show ip ospf neighbor	EXEC Privilege	View the configuration of OSPF neighbors connected to the local router.

To view the OSPFv2 configuration for LSA throttling, use the following command in EXEC Privilege mode.

Command Syntax	Command Mode	Usage
show ip ospf timers rate-limit	EXEC Privilege	View the LSAs currently in the queue.

To configure the debugging options of an OSPFv2 process, use the following command in EXEC Privilege mode.

Command Syntax	Command Mode	Usage
debug ip ospf <i>process-id</i> [event packet spf database-timers rate-limit]	EXEC Privilege	<p>View debug messages.</p> <p>To view debug messages for a specific OSPF process ID, enter debug ip ospf <i>process-id</i>.</p> <p>If you do not enter a process ID, the command applies to the first OSPF process.</p> <p>To view debug messages for a specific operation, enter one of the optional keywords:</p> <ul style="list-style-type: none"> • event: view OSPF event messages • packet: view OSPF packet information. • spf: view shortest path first (spf) information. • database-timers rate-limit: view the LSAs currently in the queue

Sample Configurations for OSPFv2


The following configurations are examples for enabling OSPFv2. These are not comprehensive directions. They are intended to give you a some guidance with typical configurations.

You can copy and paste from these examples to your CLI. Be sure you make the necessary changes to support your own IP addresses, interfaces, names, etc.

Basic OSPFv2 Router Topology

The following shows a sample basic OSPFv2 topology.

Figure 20-19. Basic Topology and CLI Commands for OSPFv2



```

router ospf 11111
 network 10.0.11.0/24 area 0
 network 10.0.12.0/24 area 0
 network 192.168.100.0/24 area 0
 !
 interface TenGigabitEthernet 1/1
 ip address 10.1.11.1/24
 no shutdown
 !
 interface TenGigabitEthernet 1/2
 ip address 10.2.12.2/24
 no shutdown
 !
 interface Loopback 10
 ip address 192.168.100.100/24
 no shutdown
  
```

```

router ospf 33333
 network 192.168.100.0/24 area 0
 network 10.0.13.0/24 area 0
 network 10.0.23.0/24 area 0
 !
 interface Loopback 30
 ip address 192.168.100.100/24
 no shutdown
 !
 interface TenGigabitEthernet 3/1
 ip address 10.1.13.3/24
 no shutdown
 !
 interface TenGigabitEthernet 3/2
 ip address 10.2.13.3/24
 no shutdown
  
```

```

router ospf 22222
 network 192.168.100.0/24 area 0
 network 10.2.21.0/24 area 0
 network 10.2.22.0/24 area 0
 !
 interface Loopback 20
 ip address 192.168.100.20/24
 no shutdown
 !
 interface TenGigabitEthernet 2/1
 ip address 10.2.21.2/24
 no shutdown
 !
 interface TenGigabitEthernet 2/2
 ip address 10.2.22.2/24
 no shutdown
  
```


Port Monitoring

Port monitoring is a feature that copies all incoming or outgoing packets on one port and forwards (mirrors) them to another port. The source port is the monitored port (MD) and the destination port is the monitoring port (MG). Port monitoring functionality is different between platforms, but the behavior is the same, with highlighted exceptions.

This chapter contains the following sections:

- [Important Points to Remember](#)
- [Port Monitoring](#)
- [Configuring Port Monitoring](#)

Important Points to Remember

- Port monitoring is supported on physical ports only; virtual local area network (VLAN) and port-channel interfaces do not support port monitoring.
- The monitored (source, MD) and monitoring ports (destination, MG) must be on the same switch.
- In general, a monitoring port should have no ip address and no shutdown as the only configuration; FTOS permits a limited set of commands for monitoring ports. To display these commands, use the command ?.
- A monitoring port also may not be a member of a VLAN.
- There may only be one destination port in a monitoring session.
- A source port (MD) can only be monitored by one destination port (MG). If you try to assign a monitored port to more than one monitoring port, the following error is displayed ([Message 1](#)).

Message 1 Assign a Monitored Port to More than One Monitoring Port

```

FTOS(conf)#mon ses 1
FTOS(conf-mon-sess-1)#source tengig 0/0 destination tengig 0/60 direction both
FTOS(conf-mon-sess-1)#do show mon ses
  SessionID      Source          Destination      Direction      Mode      Type
  -----      -
Port-based      1              TenGig 0/0      TenGig 0/60    both      interface
FTOS(conf-mon-sess-1)#mon ses 2
FTOS(conf-mon-sess-2)#source tengig 0/0 destination tengig 0/61 direction both
% Error: MD port is already being monitored.

```



Note: There is no limit to the number of monitoring sessions per system, provided that there are only four destination ports per port-pipe. If each monitoring session has a unique destination port, the maximum number of session is four per port-pipe.

Port Monitoring

The MXL 10/40GbE Switch supports multiple source-destination statements in a monitor session, but there may only be one destination port in a monitoring session ([Message 2](#)).

Message 2 One Destination Port in a Monitoring Session Error Message

```
% Error: Only one MG port is allowed in a session.
```

The number of source ports FTOS allows within a port-pipe is equal to the number of physical ports in the port-pipe (n). However, n number of ports may only have four different destination ports ([Message 3](#)).

In [Figure 21-1](#), ports 0/13, 0/14, 0/15, and 0/16 all belong to the same port-pipe. They are pointing to four different destinations (0/1, 0/2, 0/3, and 0/37). Now it is not possible for another source port from the same port-pipe (for example, 0/17) to point to another new destination (for example, 0/4). If you attempt to configure another destination, [Message 3](#) appears. However, you can configure another monitoring session that uses one of previously used destination ports ([Figure 21-2](#)).

Figure 21-1. Number of Monitoring Ports

```
FTOS#show mon session
  SessionID      Source      Destination      Direction      Mode      Type
  -----
      0      TenGig 0/13      TenGig 0/1      rx      interface      Port-based
     10      TenGig 0/14      TenGig 0/2      rx      interface      Port-based
     20      TenGig 0/15      TenGig 0/3      rx      interface      Port-based
     30      TenGig 0/16      TenGig 0/37     rx      interface      Port-based
FTOS(conf)#mon ses 300
FTOS(conf-mon-sess-300)#source tengig 0/17 destination tengig 0/4 direction tx
% Error: Exceeding max MG ports for this MD port pipe.
FTOS(conf-mon-sess-300)#
FTOS(conf-mon-sess-300)#source tengig 0/17 destination tengig 0/1 direction tx
FTOS(conf-mon-sess-300)#do show mon session
  SessionID      Source      Destination      Direction      Mode      Type
  -----
      0      TenGig 0/13      TenGig 0/1      rx      interface      Port-based
     10      TenGig 0/14      TenGig 0/2      rx      interface      Port-based
     20      TenGig 0/15      TenGig 0/3      rx      interface      Port-based
     30      TenGig 0/16      TenGig 0/37     rx      interface      Port-based
    300      TenGig 0/17      TenGig 0/1      tx      interface      Port-based
FTOS(conf-mon-sess-300)#
```

Figure 21-2. Number of Monitoring Ports

```
FTOS(conf)#mon ses 300
FTOS(conf-mon-sess-300)#source tengig 0/17 destination tengig 0/4 direction tx
% Error: Exceeding max MG ports for this MD port pipe.
FTOS(conf-mon-sess-300)#
FTOS(conf-mon-sess-300)#source tengig 0/17 destination tengig 0/1 direction tx
FTOS(conf-mon-sess-300)#do show mon session
```

SessionID	Source	Destination	Direction	Mode	Type
0	TenGig 0/13	TenGig 0/1	rx	interface	Port-based
10	TenGig 0/14	TenGig 0/2	rx	interface	Port-based
20	TenGig 0/15	TenGig 0/3	rx	interface	Port-based
30	TenGig 0/16	TenGig 0/37	rx	interface	Port-based
300	TenGig 0/17	TenGig 0/1	tx	interface	Port-based

In Figure 21-3, 0/25 and 0/26 belong to Port-pipe 1. This port-pipe again has the same restriction of only four destination ports, new or used.

Figure 21-3. Number of Monitoring Ports

```
FTOS(conf-mon-sess-300)#do show mon session
```

SessionID	Source	Destination	Direction	Mode	Type
0	TenGig 0/13	TenGig 0/1	rx	interface	Port-based
10	TenGig 0/14	TenGig 0/2	rx	interface	Port-based
20	TenGig 0/15	TenGig 0/3	rx	interface	Port-based
30	TenGig 0/16	TenGig 0/37	rx	interface	Port-based
100	TenGig 0/25	TenGig 0/38	tx	interface	Port-based
110	TenGig 0/26	TenGig 0/39	tx	interface	Port-based
300	TenGig 0/17	TenGig 0/1	tx	interface	Port-based

```
FTOS(conf-mon-sess-300)#
```

A source port may only be monitored by one destination port (Message 3), but a destination port may monitor more than one source port.

Message 3 One Destination Port in a Monitoring Session Error Message

```
% Error: Exceeding max MG ports for this MD port pipe.
```

Message 4 One Destination Port per Source Port Error Message

```
% Error: MD port is already being monitored.
```



FTOS Behavior: All monitored frames are tagged if the configured monitoring direction is transmit (TX), regardless of whether the monitored port (MD) is a Layer 2 or Layer 3 port. If the MD port is a Layer 2 port, the frames are tagged with the VLAN ID of the VLAN to which the MD belongs. If the MD port is a Layer 3 port, the frames are tagged with VLAN ID 4095. If the MD port is in a Layer 3 VLAN, the frames are tagged with the respective Layer 3 VLAN ID. For example, in the configuration *source tengig 6/0 destination tengig 6/1 direction tx*, if the MD port tengigabitethernet 6/0 is an untagged member of any VLAN, all monitored frames that the MG port tengigabitethernet 6/1 receives are tagged with the VLAN ID of the MD port. Similarly, if BPDUs are transmitted, the MG port receives them tagged with the VLAN ID 4095. This behavior might result in a difference between the number of egress packets on the MD port and monitored packets on the MG port.



FTOS Behavior: The MXL 10/40GbE Switch continues to mirror outgoing traffic even after an MD participating in Spanning Tree Protocol (STP) transitions from the forwarding to blocking.

Configuring Port Monitoring

To configure port monitoring, use the following example:

Step	Task	Command Syntax	Command Mode
1	Verify that the intended monitoring port has no configuration other than no shutdown (Figure 21-4).	show interface	EXEC Privilege
2	Create a monitoring session using the command monitor session from CONFIGURATION mode (Figure 21-4).	monitor session	CONFIGURATION
3	Specify the source and destination port and direction of traffic (Figure 21-4).	source	MONITOR SESSION

To display monitor sessions, use the show monitor session command from EXEC Privilege mode (Figure 21-4).

Figure 21-4. Configuring Port-based Monitoring

```

FTOS(conf-if-te-1/2)#show config
!
interface TenGigabitEthernet 1/2
  no ip address
  no shutdown
FTOS(conf-if-te-1/2)#exit
FTOS(conf)#monitor session 0
FTOS(conf-mon-sess-0)#source tengig 1/1 dest tengig 1/2 direction rx
FTOS(conf-mon-sess-0)#exit
FTOS(conf)#do show monitor session 0

```

SessionID	Source	Destination	Direction	Mode	Type
-----	-----	-----	-----	----	----
0	TenGig 1/1	TenGig 1/2	rx	interface	Port-based


```

FTOS(conf)#

```

In Figure 21-5, the host and server are exchanging traffic which passes through interface tengigabitethernet 1/1. Interface tengigabitethernet 1/1 is the monitored port and tengigabitethernet 1/2 is the monitoring port, which is configured to only monitor traffic received on tengigabitethernet 1/1 (host-originated traffic).

Figure 21-5. Port Monitoring Example



Port Monitoring 001

Private VLANs (PVLAN)

For syntax details about the commands described in this chapter, refer to the Private VLANs (PVLAN) Commands chapter in the *FTOS Command Reference Guide*.

This chapter contains the following sections:

- [Private VLAN Concepts](#)
- [Private VLAN Commands](#)
- [Private VLAN Configuration Task List](#)
- [Private VLAN Configuration Example](#)
- [Inspecting the Private VLAN Configuration](#)

Private VLANs (PVLANS) extend the Dell Force10 operating software (FTOS) security suite by providing Layer 2 isolation between ports within the same VLAN. A private VLAN partitions a traditional VLAN into subdomains identified by a *primary* and *secondary VLAN* pair. Private VLANs block all traffic to isolated ports except traffic from promiscuous ports. Traffic received from an isolated port is forwarded only to promiscuous ports or trunk ports.

Example uses of PVLANS:

- A hotel can use an isolated VLAN in a private VLAN to provide internet access for its guests, while stopping direct access between the guest ports.
- A service provider can provide Layer 2 security for customers and use the IP addresses more efficiently by using a separate community VLAN per customer, while at the same time using the same IP subnet address space for all community and isolated VLANs mapped to the same primary VLAN. In more detail, community VLANs are especially useful in the service provider environment because multiple customers are likely to maintain servers that must be strictly separated in customer-specific groups. A set of servers owned by a customer could comprise a community VLAN, so that those servers could communicate with each other, and would be isolated from other customers. Another customer might have another set of servers in another community VLAN. Another customer might want an isolated VLAN, which has one or more ports that are also isolated from each other.

Private VLAN Concepts

The VLAN types in a PVLAN include:

Community VLAN—a type of secondary VLAN in a primary VLAN:

- Ports in a community VLAN can communicate with each other.
- Ports in a community VLAN can communicate with all promiscuous ports in the primary VLAN.
- A community VLAN can only contain ports configured as host.

Isolated VLAN—a type of secondary VLAN in a primary VLAN:

- Ports in an isolated VLAN cannot talk directly to each other.
- Ports in an isolated VLAN can only communicate with promiscuous ports in the primary VLAN.
- An isolated VLAN can only contain ports configured as host.

Primary VLAN—the base VLAN of a private VLAN:

- A switch can have one or more primary VLANs, or none.
- A primary VLAN has one or more secondary VLANs.
- A primary VLAN and each of its secondary VLANs decrement the available number of VLAN IDs in the switch.
- A primary VLAN has one or more promiscuous ports.
- A primary VLAN might have one or more trunk ports, or none.

Secondary VLAN—a subdomain of the primary VLAN. There are two types of secondary VLAN—community VLAN and isolated VLAN.

PVLAN port types:

- **Community port:** a port that belongs to a community VLAN and is allowed to communicate with other ports in the same community VLAN and with promiscuous ports.
- **Host port:** (in the context of a private VLAN) a port in a secondary VLAN:
 - You must first assign the port that role in INTERFACE mode.
 - A port assigned the host role cannot be added to a regular VLAN.
- **Isolated port:** a port that, in Layer 2, can only communicate with promiscuous ports that are in the same PVLAN.
- **Promiscuous port:** a port that is allowed to communicate with any other port type in the PVLAN:
 - A promiscuous port can be part of more than one primary VLAN.
 - A promiscuous port cannot be added to a regular VLAN.
- **Trunk port:** carries traffic between the switches:
 - A trunk port in a PVLAN is always tagged.
 - Primary or secondary VLAN traffic is carried by the trunk port in tagged mode. The tag on the packet helps identify the VLAN to which the packet belongs.
 - A trunk port can also belong to a regular VLAN (non-private VLAN).

Each of the port types can be any type of physical Ethernet port, including port channels (LAGs). For details about port channels, refer to [Port Channel Interfaces](#) in [Interfaces](#).

For an introduction to VLANs, refer to [Layer 2](#).

Private VLAN Commands

The commands dedicated to supporting the PVLANs feature are:

Table 22-1. Private VLAN Commands

Task	Command Syntax	Command Mode
Enable/disable Layer 3 communication between secondary VLANs.	[no] ip local-proxy-arp Note: Even after you disable ip-local-proxy-arp (no ip-local-proxy-arp) in a secondary VLAN, Layer 3 communication may happen between some secondary VLAN hosts, until the ARP timeout happens on those secondary VLAN hosts.	INTERFACE VLAN
Set the mode of the selected VLAN to community, isolated, or primary.	[no] private-vlan mode {community isolated primary}	INTERFACE VLAN
Map secondary VLANs to the selected primary VLAN.	[no] private-vlan mapping secondary-vlan vlan-list	INTERFACE VLAN
Display type and status of PVLAN interfaces.	show interfaces private-vlan [interface interface]	EXEC EXEC Privilege
Display PVLANs and/or interfaces that are part of a PVLAN.	show vlan private-vlan [community interface isolated primary primary_vlan interface interface]	EXEC EXEC Privilege
Display primary-secondary VLAN mapping.	show vlan private-vlan mapping	EXEC EXEC Privilege
Set the PVLAN mode of the selected port.	switchport mode private-vlan {host promiscuous trunk}	INTERFACE



Note: Secondary VLANs are Layer 2 VLANs, so even if they are operationally down while primary VLANs are operationally up, Layer 3 traffic is still transmitted across the secondary VLANs.

To provide PVLAN data, the outputs of the following commands are augmented in FTOS version 7.8.1.0:

- show arp: refer to the IP Routing Commands chapter in the *FTOS Command Reference Guide*.
- show vlan: refer to the Layer 2 Commands chapter in the *FTOS Command Reference Guide*.

Private VLAN Configuration Task List

The following sections contain the procedures that configure a PVLAN:

- [Creating PVLAN Ports](#)
- [Creating a Primary VLAN](#)
- [Creating a Community VLAN](#)
- [Creating an Isolated VLAN](#)

Creating PVLAN Ports

PVLAN ports are those that are assigned to the Private VLAN. To assign PVLAN ports, follow these steps:

Step	Command Syntax	Command Mode	Purpose
1	<code>interface <i>interface</i></code>	CONFIGURATION	Access INTERFACE mode for the port that you want to assign to a PVLAN.
2	<code>no shutdown</code>	INTERFACE	Enable the port.
3	<code>switchport</code>	INTERFACE	Set the port in Layer 2 mode.
4	<code>switchport mode private-vlan {host promiscuous trunk}</code>	INTERFACE	Select the PVLAN mode: <ul style="list-style-type: none"> • host (port in isolated or community VLAN) • promiscuous (intra-VLAN communication port) • trunk (inter-switch PVLAN hub port)

For interface details, refer to [Enable a Physical Interface in Interfaces](#).



Note: Interfaces that are configured as PVLAN ports cannot be added to regular VLANs. Conversely, “regular” ports (ports not configured as PVLAN ports) cannot be added to PVLANS.

[Figure 22-1](#) shows the use of the `switchport mode private-vlan` command on a port and on a port channel.

Figure 22-1. `switchport mode private-vlan` Command Example

```
FTOS#conf
FTOS(conf)#interface TenGigabitEthernet 2/1
FTOS(conf-if-te-2/1)#switchport mode private-vlan promiscuous

FTOS(conf)#interface TenGigabitEthernet 2/2
FTOS(conf-if-te-2/2)#switchport mode private-vlan host

FTOS(conf)#interface TenGigabitEthernet 2/3
FTOS(conf-if-te-2/3)#switchport mode private-vlan trunk

FTOS(conf)#interface TenGigabitEthernet 2/2
FTOS(conf-if-te-2/2)#switchport mode private-vlan host
```

Creating a Primary VLAN

A primary VLAN is a port-based VLAN that is specifically enabled as a primary VLAN to contain the promiscuous ports and PVLAN trunk ports for the private VLAN. A primary VLAN also contains a mapping to secondary VLANs, which are comprised of community VLANs and isolated VLANs. To create a primary VLAN, follow these steps:

Step	Command Syntax	Command Mode	Purpose
1	<code>interface vlan <i>vlan-id</i></code>	CONFIGURATION	Access INTERFACE VLAN mode for the VLAN you want to assign the PVLAN interfaces.
2	<code>no shutdown</code>	INTERFACE VLAN	Enable the VLAN.
3	<code>private-vlan mode primary</code>	INTERFACE VLAN	Set PVLAN mode of the selected VLAN to primary.
4	<code>private-vlan mapping secondary-vlan <i>vlan-list</i></code>	INTERFACE VLAN	Map secondary VLANs to the selected primary VLAN. The list of secondary VLANs can be: <ul style="list-style-type: none"> Specified in comma-delimited (<i>VLAN-ID,VLAN-ID</i>) or hyphenated-range format (<i>VLAN-ID-VLAN-ID</i>). Specified with this command even before they have been created. Amended by specifying the new secondary VLAN to be added to the list.
5	<code>tagged <i>interface</i></code> or <code>untagged <i>interface</i></code>	INTERFACE VLAN	Add promiscuous ports as tagged or untagged interfaces. Add PVLAN trunk ports to the VLAN only as tagged interfaces. Interfaces can be entered singly or in range format, either comma-delimited (<i>slot/port,port,port</i>) or hyphenated (<i>slot/port-port</i>). You can only add promiscuous ports or PVLAN trunk ports to the PVLAN (no host or regular ports).
6	<code>ip address <i>ip address</i></code>	INTERFACE VLAN	(OPTIONAL) Assign an IP address to the VLAN.
7	<code>ip local-proxy-arp</code>	INTERFACE VLAN	(OPTIONAL) Enable/disable Layer 3 communication between secondary VLANs.



Note: If a promiscuous or host port is untagged in a VLAN and it receives a tagged packet in the same VLAN, the packet is NOT dropped.

Creating a Community VLAN

A community VLAN is a secondary VLAN of the primary VLAN in a Private VLAN. The ports in a community VLAN can talk to each other and with the promiscuous ports in the primary VLAN. To create a community VLAN, follow these steps:

Step	Command Syntax	Command Mode	Purpose
1	<code>interface vlan <i>vlan-id</i></code>	CONFIGURATION	Access INTERFACE VLAN mode for the VLAN that you want to make a community VLAN.
2	<code>no shutdown</code>	INTERFACE VLAN	Enable the VLAN.
3	<code>private-vlan mode community</code>	INTERFACE VLAN	Set PVLAN mode of the selected VLAN to community.
4	<code>tagged <i>interface</i></code> or <code>untagged <i>interface</i></code>	INTERFACE VLAN	Add one or more host ports to the VLAN. The interfaces can be entered singly or in range format, either comma-delimited (<i>slot/port,port,port</i>) or hyphenated (<i>slot/port-port</i>). You can only add host (isolated) ports to the VLAN.

Creating an Isolated VLAN

An isolated VLAN is a secondary VLAN of a primary VLAN. Its ports can only talk with the promiscuous ports in that primary VLAN. To create an isolated VLAN, follow these steps:

Step	Command Syntax	Command Mode	Purpose
1	<code>interface vlan <i>vlan-id</i></code>	CONFIGURATION	Access INTERFACE VLAN mode for the VLAN that you want to make an isolated VLAN.
2	<code>no shutdown</code>	INTERFACE VLAN	Enable the VLAN.
3	<code>private-vlan mode isolated</code>	INTERFACE VLAN	Set PVLAN mode of the selected VLAN to isolated.
4	<code>tagged <i>interface</i></code> or <code>untagged <i>interface</i></code>	INTERFACE VLAN	Add one or more host ports to the VLAN. The interfaces can be entered singly or in range format, either comma-delimited (<i>slot/port,port,port</i>) or hyphenated (<i>slot/port-port</i>). You can only add ports defined as host to the VLAN.

To configure the PVLAN member VLANs (primary, community, and isolated VLANs), use the following commands in VLAN INTERFACE mode (Figure 22-2).

Figure 22-2. Configuring VLANs for a Private VLAN

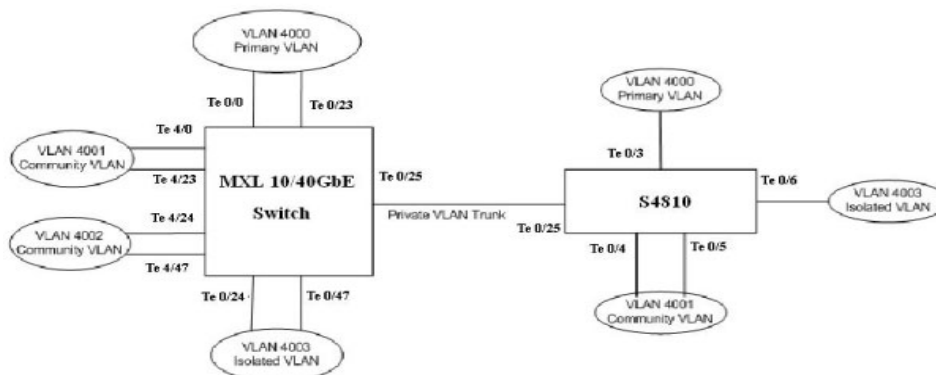
```
FTOS#conf
FTOS(conf)# interface vlan 10
FTOS(conf-vlan-10)# private-vlan mode primary
FTOS(conf-vlan-10)# private-vlan mapping secondary-vlan 100-101
FTOS(conf-vlan-10)# untagged TenGig 2/1
FTOS(conf-vlan-10)# tagged TenGig 2/3

FTOS(conf)# interface vlan 101
FTOS(conf-vlan-101)# private-vlan mode community
FTOS(conf-vlan-101)# untagged TenGig 2/10

FTOS(conf)# interface vlan 100
```

Private VLAN Configuration Example

Figure 22-3. Sample Private VLAN Topology



The following configuration is based on the [Figure 22-3](#):

On MXL 10/40GbE Switch:

- TenGig 0/0 and TenGig 0/23 are configured as promiscuous ports, assigned to the primary VLAN, VLAN 4000.
- TenGig 0/25 is configured as a PVLAN trunk port, also assigned to the primary VLAN 4000.
- TenGig 0/24 and TenGig 0/47 are configured as host ports and assigned to the isolated VLAN, VLAN 4003.
- TenGig 4/0 and TenGig 0/23 are configured as host ports and assigned to the community VLAN, VLAN 4001.
- TenGig 4/24 and TenGig 4/47 are configured as host ports and assigned to community VLAN 4002.

The results are:

- The ports in community VLAN 4001 can communicate directly with each other and with promiscuous ports.
- The ports in community VLAN 4002 can communicate directly with each other and with promiscuous ports.
- The ports in isolated VLAN 4003 can only communicate with the promiscuous ports in the primary VLAN 4000.
- All the ports in the secondary VLANs (both community and isolated VLANs) can only communicate with ports in the other secondary VLANs of that PVLAN over Layer 3, and only when you use the `ip local-proxy-arp` command in the primary VLAN.



Note: Even after you disable the `ip-local-proxy-arp` command (using the `no ip-local-proxy-arp` command) in a secondary VLAN, Layer 3 communication may happen between some secondary VLAN hosts until the ARP timeout happens on those secondary VLAN hosts.

In parallel, on S4810:

- Te 0/3 is a promiscuous port and Te 0/25 is a PVLAN trunk port, assigned to the primary VLAN 4000.
- Te 0/4-6 are host ports. Te 0/4 and Te 0/5 are assigned to the community VLAN 4001, while Te 0/6 is assigned to the isolated VLAN 4003.

The results are:

- The S4810 ports would have the same intra-switch communication characteristics as described above for the MXL 10/40GbE Switch.
- For transmission between switches, tagged packets originating from host PVLAN ports in one secondary VLAN and destined for host PVLAN ports in the other switch travel through the promiscuous ports in the local VLAN 4000 and then through the trunk ports (0/25 in each switch).

Inspecting the Private VLAN Configuration

The standard methods of inspecting configurations also apply in PVLANS:

- Within `INTERFACE` and `INTERFACE VLAN` modes, to display the specific interface configuration, use the `show config` command.
- To inspect the running-config, use the `grep` pipe option (`show running-config | grep string`). You can display a specific part of the running-config. [Figure 22-6](#) shows the PVLAN parts of the running-config from the switch in the topology diagram shown in [Figure 22-3](#).

- You can also use one of three show commands that are specific to the Private VLAN feature:
 - show interfaces private-vlan [interface *interface*]: Display the type and status of the configured PVLAN interfaces. Refer to the example output in the Security chapter of the *FTOS Command Reference Guide*.
 - show vlan private-vlan [community | *interface* | isolated | primary | *primary_vlan* | interface *interface*]: Display the configured PVLANS or interfaces that are part of a PVLAN. Figure 22-4 shows the results of using the command without command options on the switch in the topology diagram shown in Figure 22-3.
 - show vlan private-vlan mapping: Display the primary-secondary VLAN mapping. Refer to the example output from the MXL Switch in Figure 22-4.
- Two show commands revised to display PVLAN data are:
 - show arp
 - show vlan: Refer to the revised output in Figure 22-5.

Figure 22-4. show vlan private-vlan Command Example

```
FTOS#show vlan private-vlan

Primary Secondary Type      Active Ports
-----
20          30      Primary  Yes    Te 1/1,5
          40      Community Yes    Te 1/2
          40      Isolated Yes    Te 1/3
FTOS#
```

In Figure 22-5, note the addition of the PVLAN codes—P, I, and C—in the left column:

Figure 22-5. show vlan Command Example

```
FTOS#show vlan

Codes: * - Default VLAN, G - GVRP VLANs, R - Remote Port Mirroring VLANs, P - Primary, C -
Community, I - Isolated
Q: U - Untagged, T - Tagged
  x - Dot1x untagged, X - Dot1x tagged
  G - GVRP tagged, M - Vlan-stack, H - VSN tagged
  i - Internal untagged, I - Internal tagged, v - VLT untagged, V - VLT tagged

  NUM      Status      Description      Q Ports
*  1        Active
P  20        Active
C  30        Active          T Te 1/1,5
I  40        Active          T Te 1/2
FTOS#
```

Figure 22-6. running-config Command Example of PVLAN Configuration


```
FTOS#show vlan
!
interface TenGigabitEthernet 1/1
no ip address
switchport
switchport mode private-vlan promiscuous
no keepalive
no shutdown
!
interface TenGigabitEthernet 1/2
no ip address
switchport
switchport mode private-vlan host
no shutdown
!
interface TenGigabitEthernet 1/3
no ip address
switchport
switchport mode private-vlan host
no shutdown
!
interface TenGigabitEthernet 1/5
no ip address
switchport
switchport mode private-vlan trunk
no shutdown
interface Vlan 20
private-vlan mode primary
private-vlan mapping secondary-vlan 30,40
no ip address
tagged TenGigabitEthernet 1/1,5
shutdown
!
interface Vlan 30
private-vlan mode community
no ip address
tagged TenGigabitEthernet 1/2
no shutdown
!
```


Per-VLAN Spanning Tree Plus (PVST+)

Overview

Per-VLAN spanning tree plus (PVST+) is a variation of spanning tree—developed by a third party—that allows you to configure a separate spanning tree instance for each VLAN (Figure 23-1). For more information about spanning tree, refer to [Spanning Tree Protocol \(STP\)](#).

Figure 23-1. Per-VLAN Spanning Tree



The Dell Force10 operating software (FTOS) supports three other variations of spanning tree ([Table 23-1](#)).

Table 23-1. FTOS Supported Spanning Tree Protocols

Dell Force10 Term	IEEE Specification
Spanning Tree Protocol (STP)	802.1d
Rapid Spanning Tree Protocol (RSTP)	802.1w
Multiple Spanning Tree Protocol (MSTP)	802.1s
Per-VLAN Spanning Tree Plus (PVST+)	Third Party

Implementation Information

- The FTOS implementation of PVST+ is based on RPVST.
- The FTOS implementation of PVST+ uses IEEE 802.1s costs as the default costs ([Table 23-2](#)). Other implementations use IEEE 802.1d costs as the default costs. If you are using Dell Force10 systems in a multi-vendor network, verify that the costs are values you intended.
- You can enable PVST+ on 255 VLANs.

Configure Per-VLAN Spanning Tree Plus

Configuring PVST+ is a four-step process:

1. Configure interfaces for Layer 2.
2. Place the interfaces in VLANs.
3. Enable PVST+.
4. Optionally, for load balancing, select a non-default bridge-priority for a VLAN.

Related Configuration Tasks

- [Modify Global PVST+ Parameters](#)
- [Enable BPDU Filtering globally](#)
- [Configure an EdgePort](#)
- [Flush MAC Addresses after a Topology Change](#)
- [Preventing Network Disruptions with BPDU Guard](#)
- [SNMP Traps for Root Elections and Topology Changes](#)
- [PVST+ in Multi-vendor Networks](#)
- [PVST+ Extended System ID](#)
- [PVST+ Sample Configurations](#)

Enable PVST+

When you enable PVST+, FTOS instantiates STP on each active VLAN. To enable PVST+ globally, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Enter PVST context.	protocol spanning-tree pvst	PROTOCOL PVST
2	Enable PVST+.	no disable	PROTOCOL PVST

Disable PVST+

To disable PVST+, use the following commands.

Task	Command Syntax	Command Mode
Disable PVST+ globally.	disable	PROTOCOL PVST
Disable PVST+ on an interface, or remove a PVST+ parameter configuration.	no spanning-tree pvst	INTERFACE


Display your PVST+ configuration by entering the show config command from PROTOCOL PVST context (Figure 23-2).

Figure 23-2. Display the PVST+ Configuration

```
FTOS(conf-pvst)#show config verbose
!  
protocol spanning-tree pvst  
no disable  
vlan 100 bridge-priority 4096
```

Influence PVST+ Root Selection

In Figure 23-1, all VLANs use the same forwarding topology because R2 is elected the root and all TenGigabitEthernet ports have the same cost. Figure 23-3 changes the bridge priority of each bridge so that a different forwarding topology is generated for each VLAN. This behavior demonstrates how you can use PVST+ to achieve load balancing.

Figure 23-3. Load Balancing with PVST+

The bridge with the bridge value for bridge priority is elected root. Because all bridges use the default priority (until configured otherwise), the lowest MAC address is used as a tie-breaker. Assign bridges a low non-default value for bridge priority to increase the likelihood that it is selected as the STP root.

Task	Command Syntax	Command Mode
Assign a bridge priority. Range: 0 to 61440 Default: 32768	vlan bridge-priority	PROTOCOL PVST

Display the PVST+ forwarding topology by entering the `show spanning-tree pvst [vlan vlan-id]` command from EXEC Privilege mode (Figure 23-4).

Figure 23-4. Display the PVST+ Forwarding Topology

```
FTOS(conf-if-te-5/41)#do show spanning-tree pvst vlan 2
VLAN 2
Root Identifier has priority 32768, Address 001e.c9f1.00f3
Root Bridge hello time 2, max age 20, forward delay 15
Bridge Identifier has priority 32768, Address 001e.c9f1.00f3
Configured hello time 2, max age 20, forward delay 15
Bpdu filter disabled globally
We are the root of VLAN 2
Current root has priority 32768, Address 001e.c9f1.00f3
Number of topology changes 2, last change occurred 00:14:39 ago on Po 23

Port 24 (Port-channel 23) is designated Forwarding
Port path cost 1600, Port priority 128, Port Identifier 128.24
Designated root has priority 32768, address 001e.c9f1.00:f3
Designated bridge has priority 32768, address 001e.c9f1.00:f3
Designated port id is 128.24 , designated path cost 0
Number of transitions to forwarding state 1
BPDU sent 449, received 0
The port is not in the Edge port mode, bpdu filter is disabled

Port 450 (TenGigabitEthernet 5/41) is disabled Discarding
Port path cost 2000, Port priority 128, Port Identifier 128.450
Designated root has priority 32768, address 001e.c9f1.00:f3
Designated bridge has priority 32768, address 001e.c9f1.00:f3
Designated port id is 128.450 , designated path cost 0
Number of transitions to forwarding state 0
BPDU sent 0, received 0
The port is not in the Edge port mode, bpdu filter is disabled

Port 459 (TenGigabitEthernet 5/50) is designated Forwarding
Port path cost 2000, Port priority 128, Port Identifier 128.459
Designated root has priority 32768, address 001e.c9f1.00:f3
Designated bridge has priority 32768, address 001e.c9f1.00:f3
Designated port id is 128.459 , designated path cost 0
Number of transitions to forwarding state 1
BPDU sent 457, received 0
The port is not in the Edge port mode, bpdu filter is disabled
```

Modify Global PVST+ Parameters

The root bridge sets the values for forward-delay and hello-time and overwrites the values set on other PVST+ bridges.

- **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 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 PVST+ topology.

To change PVST+ parameters, use the following commands on the Root Bridge:


Task	Command Syntax	Command Mode
Change the forward-delay parameter. <ul style="list-style-type: none"> • Range: 4 to 30 • Default: 15 seconds 	vlan forward-delay	PROTOCOL PVST
Change the hello-time parameter. <p>Note: With large configurations (especially those with more ports), Dell Force10 recommends that you increase the hello-time.</p> Range: 1 to 10 Default: 2 seconds	vlan hello-time	PROTOCOL PVST
Change the max-age parameter. Range: 6 to 40 Default: 20 seconds	vlan max-age	PROTOCOL PVST

To view the values for the global PVST+ parameters, use the show spanning-tree pvst command (Figure 23-4).

Enable BPDU Filtering globally

The enabling of BPDU Filtering stops transmitting of BPDUs on the operational port fast enabled ports by default. When BPDUs are received, the spanning tree is automatically prepared. By default global bpdud filtering is disabled.

Figure 23-5. BPDU Filtering enabled globally



Task	Command Syntax	Command Mode
Enable BPDU Filter globally to filter transmission of BPDU port fast enabled interfaces.	edge-port bpdu filter default	PROTOCOL PVST

Modify Interface PVST+ Parameters

To increase or decrease the probability that a port becomes a forwarding port, you can adjust two interface parameters:

- **Port cost** is a value that is based on the interface type. The greater the port cost, the less likely the port is selected to be a forwarding port.
- **Port priority** influences the likelihood that a port is selected to be a forwarding port in case that several ports have the same port cost.

Table 23-2 lists the default values for port cost by interface.

Table 23-2. PVST+ Default Port Cost Values

Port Cost	Default Value
1000-Mb/s Ethernet interfaces	20000
10-Gigabit Ethernet interfaces	2000
40-Gigabit Ethernet interfaces	1400
Port Channel with one 10-Gigabit Ethernet interface	2000
Port Channel with one 40-Gigabit Ethernet interface	1400
Port Channel with two 10-Gigabit Ethernet interfaces	1800
Port Channel with two 40-Gigabit Ethernet interfaces	600



Note: The FTOS implementation of PVST+ uses IEEE 802.1s costs as the default costs. Other implementations use IEEE 802.1d costs as the default costs. If you are using Dell Force10 systems in a multi-vendor network, verify that the costs are values you intended.

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 200000 Default: refer to Table 23-2 .	spanning-tree pvst vlan cost	INTERFACE
Change the port priority of an interface. Range: 0 to 240, in increments of 16 Default: 128	spanning-tree pvst vlan priority	INTERFACE

To view the values for interface PVST+ parameters, use the show spanning-tree pvst command ([Figure 23-4](#)).

Configure an EdgePort

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 it should behave otherwise; it does not go through the Learning and Listening states.

The bpduguard shutdown-on-violation option causes the interface hardware to be shutdown when it receives a BPDU. Only after you implement the bpduguard option is the interface placed in an Error Disabled state. When receiving the BPDU, the physical interface remains up and spanning-tree drops packets in the hardware after a BPDU violation. BPDUs are dropped in the software after receiving the BPDU violation. 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 you enable it on an interface connected to a network.

The enabling of BPDU Filtering stops sending and receiving of BPDUs on the port fast enabled ports. When both BPDU guard and BPDU filter are enabled on the port, BPDU filter takes the higher precedence. By default, BPDU filtering on an interface is disabled.

To enable EdgePort on an interface, use the following command:

Task	Command Syntax	Command Mode
Enable EdgePort on an interface.	spanning-tree pvst edge-port [bpduguard [shutdown-on-violation] bpdufilter]	INTERFACE

To view the EdgePort status of each interface, use the show spanning-tree pvst command ([Figure 23-4](#)).



FTOS Behavior: Regarding the bpduguard shutdown-on-violation behavior:

- 1 If the interface to be shutdown is a port channel, all the member ports are disabled in the hardware.
- 2 When you add a physical port to a port channel already in Error Disable state, the new member port id also disabled in the hardware.
- 3 When you remove a physical port from a port channel in Error Disable state, the Error Disabled state is cleared on this physical port (the physical port is enabled in the hardware).
- 4 You can clear the Error Disabled state with any of the following methods:
 - Perform a shutdown command on the interface.
 - Disable the shutdown-on-violation command on the interface (no spanning-tree pvst edge-port [bpduguard | [shutdown-on-violation]]).
 - Disable spanning tree on the interface (no spanning-tree in INTERFACE mode).
 - Disabling global spanning tree (no spanning-tree in CONFIGURATION mode).

PVST+ in Multi-vendor Networks

Some non-Dell Force10 systems that have hybrid ports participating in PVST+, transmit two kinds of BPDUs: an 802.1D BPDU and an untagged PVST+ BPDU.


Dell Force10 systems do not expect PVST+ BPDU (tagged or untagged) on an untagged port. If this happens, FTOS places the port in the Error Disable state. This behavior might result in the network not converging. To prevent FTOS from executing this action, use the `no spanning-tree pvst err-disable cause invalid-pvst-bpdu` command. After you configure this command, if the port receives a PVST+ BPDU, the BPDU is dropped and the port remains operational.

PVST+ Extended System ID

In [Figure 23-6](#), ports P1 and P2 are untagged members of different VLANs. These ports are untagged because the hub is VLAN unaware. There is no data loop in the above scenario; however, you can employ PVST+ to avoid potential misconfigurations.

If you enable PVST+ on the Dell Force10 switch in this network, P1 and P2 receive BPDUs from each other. Ordinarily, the Bridge ID in the frame matches the Root ID, a loop is detected, and the rules of convergence require that P2 move to the Blocking state because it has the lowest port ID.

To keep both ports in Forwarding state, use Extend System ID. Extend System ID augments the Bridge ID with a VLAN ID to differentiate BPDUs on each VLAN so that PVST+ does not detect a loop, and both ports can remain in Forwarding state.

Figure 23-6. PVST+ with Extend System ID

Task	Command Syntax	Command Mode
Augment the Bridge ID with the VLAN ID.	extend system-id	PROTOCOL PVST

```

FTOS(conf-pvst)#do show spanning-tree pvst vlan 5 brief

VLAN 5
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32773, Address 0001.e832.73f7
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID  Priority 32773 (priority 32768 sys-id-ext 5), Address 0001.e832.73f7
We are the root of Vlan 5
Configured hello time 2, max age 20, forward delay 15
...

```

PVST+ Sample Configurations

Figure 23-7, Figure 23-8, and Figure 23-9 provide the running configurations for the topology shown in Figure 23-3.

Figure 23-7. PVST+ Sample Configuration: R1 Running-Configuration

```
interface TenGigabitEthernet 1/22
  no ip address
  switchport
  no shutdown
!
interface TenGigabitEthernet 1/32
  no ip address
  switchport
  no shutdown
!
protocol spanning-tree pvst
  no disable
  vlan 100 bridge-priority 4096
interface Vlan 100
  no ip address
  tagged TenGigabitEthernet 1/22,32
  no shutdown
!
interface Vlan 200
  no ip address
  tagged TenGigabitEthernet 1/22,32
  no shutdown
!
interface Vlan 300
  no ip address
  tagged TenGigabitEthernet 1/22,32
  no shutdown
!
protocol spanning-tree pvst
  no disable
  vlan 100 bridge-priority 4096
```

Figure 23-8. PVST+ Sample Configuration: R2 Running-Configuration

```

interface TenGigabitEthernet 2/12
  no ip address
  switchport
  no shutdown
!
interface TenGigabitEthernet 2/32
  no ip address
  switchport
  no shutdown
!
interface Vlan 100
  no ip address
  tagged TenGigabitEthernet 2/12,32
  no shutdown
!
interface Vlan 200
  no ip address
  tagged TenGigabitEthernet 2/12,32
  no shutdown
!
interface Vlan 300
  no ip address
  tagged TenGigabitEthernet 2/12,32
  no shutdown
!
protocol spanning-tree pvst
  no disable
  vlan 200 bridge-priority 4096

```

Figure 23-9. PVST+ Sample Configuration: R3 Running-Configuration

```

interface TenGigabitEthernet 3/12
  no ip address
  switchport
  no shutdown
!
interface TenGigabitEthernet 3/22
  no ip address
  switchport
  no shutdown
!
interface Vlan 100
  no ip address
  tagged TenGigabitEthernet 3/12,22
  no shutdown
!
interface Vlan 200
  no ip address
  tagged TenGigabitEthernet 3/12,22
  no shutdown
!
interface Vlan 300
  no ip address
  tagged TenGigabitEthernet 3/12,22
  no shutdown
!
protocol spanning-tree pvst
  no disable
  vlan 300 bridge-priority 4096

```


Quality of Service (QoS)

Overview

Differentiated service is accomplished by classifying and queuing traffic and assigning priorities to those queues.

The MXL Switch traffic has four data queues per port. All queues are serviced using the Weighted Round Robin scheduling algorithm. You can only manage queuing prioritization on egress. (Figure 24-1)


 **Note:** When you enable DCB, the egress QoS features in the output QoS policy-map (such as service-class bandwidth-percentage and bandwidth-percentage), the default bandwidth allocation ratio for egress queues and strict-priority may not work as intended. This is to provide compatibility with DCBX. Hence, it is recommended to have the DCB disabled when you wish to apply these features exclusively.

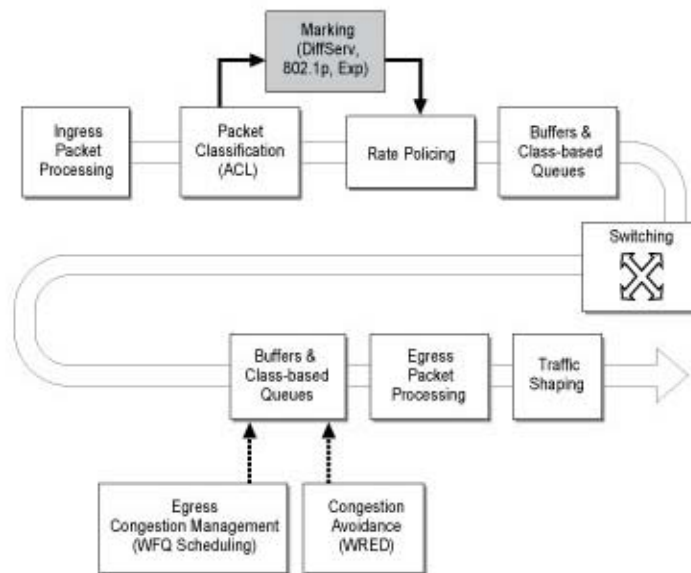
Table 24-1. FTOS Support for Port-based, Policy-based, and Multicast QoS Features

Feature	Direction
Port-Based QoS Configurations	Ingress + Egress
Set dot1p Priorities for Incoming Traffic	Ingress
Honor dot1p Priorities on Ingress Traffic	
Configure Port-based Rate Policing	
Configure Port-based Rate Shaping	Egress
Policy-Based QoS Configurations	Ingress + Egress
Classify Traffic	Ingress
Create a Layer 3 Class Map	
Set DSCP Values for Egress Packets Based on Flow	
Create a Layer 2 Class Map	
Create a QoS Policy	Ingress + Egress

Table 24-1. FTOS Support for Port-based, Policy-based, and Multicast QoS Features

Feature	Direction
Create an Input QoS Policy	Ingress
Configure Policy-Based Rate Policing	
Set a DSCP Value for Egress Packets	
Set a dot1p Value for Egress Packets	
Create an Output QoS Policy	Egress
Configure Policy-Based Rate Shaping	
Allocate Bandwidth to the Queue	
Configure a Scheduler to Queue	
Specify WRED Drop Precedence	
Create Policy Maps	Ingress + Egress
Create Input Policy Maps	Ingress
Honor DSCP Values on Ingress Packets	
Honoring dot1p Values on Ingress Packets	
Create Output Policy Maps	Egress
Specify an Aggregate QoS Policy	
QoS Rate Adjustment	
Strict-Priority Queueing	—
Weighted Random Early Detection	Egress
Create WRED Profiles	

Figure 24-1. Dell Force10 QoS Architecture



Implementation Information

The Dell Force10 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 Headers*
- RFC 2475, *An Architecture for Differentiated Services*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*

You cannot configure port-based and policy-based QoS on the same interface.

Port-Based QoS Configurations

You can configure the following QoS features on an interface:



Note: You cannot simultaneously use egress rate shaping and ingress rate policing on the same VLAN.

- [Set dot1p Priorities for Incoming Traffic](#)
- [Configure Port-based Rate Policing](#)
- [Configure Port-based Rate Shaping](#)

Set dot1p Priorities for Incoming Traffic

Change the priority of incoming traffic on the interface using the dot1p-priority command from INTERFACE mode (Figure 24-2). The Dell Force10 operating software (FTOS) places marked traffic in the corresponding queue as shown in Table 24-2. If you set a dot1p priority for a port-channel, all port-channel members are configured with the same value. You cannot assign a dot1p value to individual interfaces in a port-channel.



FTOS Behavior: The MXL Switch distributes eight dot1p priorities across four data queues.

Table 24-2. dot1p-priority values and queue numbers

dot1p	Queue Number
0	0
1	0
2	0
3	1
4	2
5	3
6	3
7	3

Figure 24-2. Configuring dot1p Priority on an Interface

```
FTOS#config
FTOS(conf)#interface tengigabitethernet 1/0
FTOS(conf-if)#switchport
FTOS(conf-if)#dot1p-priority 1
FTOS(conf-if)#end
FTOS#
```



Note: The dot1p-priority command marks all incoming traffic on an interface with a specified dot1p priority and maps all incoming traffic to the corresponding queue (Table 24-2). When you enable PFC and/or ETS on an interface, incoming traffic with a specified dot1p priority can be distributed across different queues. Therefore, when you use PFC and ETS to manage data center traffic, it is not recommended that you use the dot1p-priority command to set a queue assignment. See [Data Center Bridging \(DCB\)](#) for more information.

Honor dot1p Priorities on Ingress Traffic

By default, FTOS does not honor dot1p priorities on ingress traffic. To honor dot1p priorities on ingress traffic, use the service-class dynamic dot1p command from INTERFACE mode (Figure 24-3). You can configure this feature on physical interfaces and port-channels, but you cannot configure it on individual interfaces in a port channel.

On the MXL Switch, you can configure service-class dynamic dot1p from CONFIGURATION mode, which applies the configuration to all interfaces. A CONFIGURATION mode service-class dynamic dot1p entry supersedes any INTERFACE entries. For more information, refer to [Mapping dot1p Values to Service Queues](#).



Note: You cannot configure service-policy input and service-class dynamic dot1p on the same interface.

Figure 24-3. service-class dynamic dot1p Command Example

```
FTOS#config t
FTOS(conf)#interface tengigabitethernet 1/0
FTOS(conf-if)#service-class dynamic dot1p
FTOS(conf-if)#end
FTOS#
```

Priority-Tagged Frames on the Default VLAN

Priority-tagged frames are 802.1Q tagged frames with VLAN ID 0. For VLAN classification, these packets are treated as untagged. However, the dot1p value is still honored when you configure service-class dynamic dot1p or trust dot1p.

When priority-tagged frames ingress an untagged port or hybrid port the frames are classified to the default VLAN of the port, and to a queue according to their dot1p priority if you configure service-class dotp or trust dot1p. When priority-tagged frames ingress a tagged port, the frames are dropped because for a tagged port the default VLAN is 0.



FTOS Behavior: Hybrid ports can receive untagged, tagged, and priority tagged frames. The rate metering calculation may be inaccurate for untagged ports because an internal assumption is made that all frames are treated as tagged. Internally, the ASIC adds a 4-bytes tag to received untagged frames. Though these 4-bytes are not part of the untagged frame received on the wire, they are included in the rate metering calculation resulting in metering inaccuracy.

Configure Port-based Rate Policing

To configure rate policing ingress traffic on an interface, use the rate police command from INTERFACE mode ([Figure 24-4](#)). If the interface is a member of a VLAN, you may specify the VLAN for which ingress packets are policed.

Figure 24-4. Rate Policing Ingress Traffic

```
FTOS#config t
FTOS(conf)#interface tengigabitethernet 1/0
FTOS(conf-if)#rate police 100 40 peak 150 50
FTOS(conf-if)#end
FTOS#
```

Configure Port-based Rate Shaping

Rate shaping buffers, rather than drops, traffic that exceeds the specified rate until the buffer is exhausted. If any stream exceeds the configured bandwidth on a continuous basis, it can consume all of the buffer space that is allocated to the port.

- To apply rate shaping to outgoing traffic on a port, use the rate shape command from INTERFACE mode (Figure 24-5).
- To apply rate shaping to a queue, use the rate-shape command from QoS Policy mode.

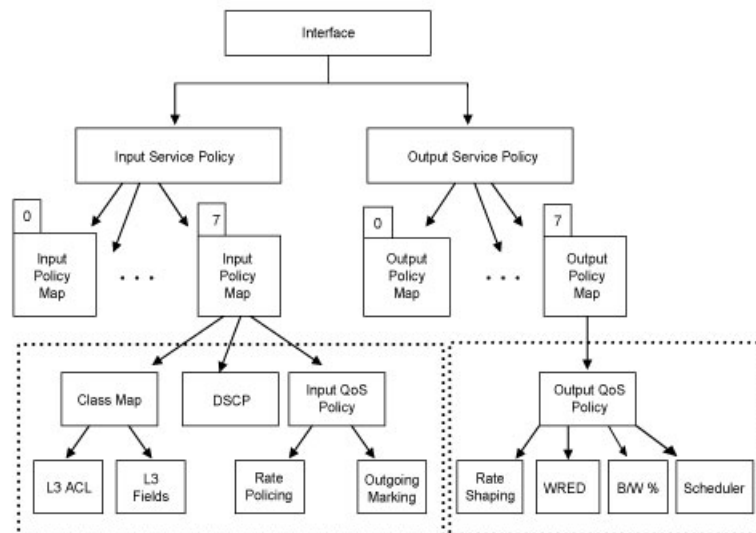
Figure 24-5. Applying Rate Shaping to Outgoing Traffic

```
FTOS#config
FTOS(conf)#interface tengigabitethernet 1/0
FTOS(conf-if)#rate shape 500 50
FTOS(conf-if)#end
FTOS#
```

Policy-Based QoS Configurations

Policy-based QoS configurations consist of the components shown in Figure 24-6.

Figure 24-6. Constructing Policy-based QoS Configurations



Classify Traffic

Class maps differentiate traffic so that you can apply separate quality of service policies to each class. For both class maps, Layer 2 and Layer 3, FTOS matches packets against match criteria in the order that you configure them.

Create a Layer 3 Class Map

A Layer 3 class map differentiates ingress packets based on the DSCP value or IP precedence, and characteristics defined in an IP access control list (ACL). You may specify more than one DSCP and IP precedence value, but only one value must match to trigger a positive match for the class map.

1. To create a match-any class map, use the `class-map match-any` command or to create a match-all class map, use the `class-map match-all` command from CONFIGURATION mode (Figure 24-7).
2. After you create a class-map, FTOS places you in CLASS MAP mode. From this mode, specify your match criteria using the `match ip` command (Figure 24-7). Match-any class maps allow up to five ACLs. Match-all class-maps allow only one ACL.
3. After you specify your match criteria, link the class-map to a queue using the `service-queue` command from POLICY MAP mode (Figure 24-7).

Figure 24-7. Using the Order Keyword in ACLs

```
FTOS(conf)#ip access-list standard acl1
FTOS(conf-std-nacl)#permit 20.0.0.0/8
FTOS(conf-std-nacl)#exit
FTOS(conf)#ip access-list standard acl2
FTOS(conf-std-nacl)#permit 20.1.1.0/24 order 0
FTOS(conf-std-nacl)#exit
FTOS(conf)#class-map match-all cmap1
FTOS(conf-class-map)#match ip access-group acl1
FTOS(conf-class-map)#exit
FTOS(conf)#class-map match-all cmap2
FTOS(conf-class-map)#match ip access-group acl2
FTOS(conf-class-map)#exit
FTOS(conf)#policy-map-input pmap
FTOS(conf-policy-map-in)#service-queue 4 class-map cmap1
FTOS(conf-policy-map-in)#service-queue 1 class-map cmap2
FTOS(conf-policy-map-in)#exit
FTOS(conf)#interface tengig 1/0
FTOS(conf-if-te-1/0)#service-policy input pmap
```

Create a Layer 2 Class Map

All class maps are Layer 3 by default; you can create a Layer 2 class map by specifying the option `layer2` with the `class-map` command. A Layer 2 class map differentiates traffic according to the 802.1p value and/or characteristics defined in a MAC ACL.

1. To create a match-any class map, use the `class-map match-any` command or to create a match-all class map, use the `class-map match-all` command from CONFIGURATION mode, and enter the keyword `layer2`.
2. After you create a class-map, FTOS places you in CLASS MAP mode. From this mode, specify your match criteria using the `match mac` command. Match-any class maps allow up to five access-lists. Match-all class-maps allow only one access list. You can match against only one VLAN ID.
3. After you specify your match criteria, link the class-map to a queue using the `service-queue` command from POLICY MAP mode.

Determine the Order in Which You Use ACLs to Classify Traffic

When you link class-maps to queues using the `service-queue` command, FTOS matches the class-maps according to queue priority (queue numbers closer to 0 have lower priorities). For example, in [Figure 24-7](#), class-map `cmap2` is matched against ingress packets before `cmap1`.

ACLs `acl1` and `acl2` have overlapping rules because the address range `20.1.1.0/24` is within `20.0.0.0/8`. Therefore, (without the keyword `order`) packets within the range `20.1.1.0/24` match positive against `cmap1` and are buffered in queue 4, though you intended for these packets to match positive against `cmap2` and be buffered in queue 1.

In cases such as these, where class-maps with overlapping ACL rules are applied to different queues, use the keyword `order` to specify the order in which you want to apply ACL rules ([Figure 24-7](#)). The order can range from 0 to 254. FTOS writes to the content addressable memory (CAM) ACL rules with lower order numbers (order numbers closer to 0) before rules with higher order numbers so that packets are matched as you intended. By default, all ACL rules have an order of 254.

Set DSCP Values for Egress Packets Based on Flow

Match-any Layer 3 flows may have several match criteria. All flows that match at least one of the match criteria are mapped to the same queue because they are in the same class map. Setting a DSCP value from QOS-POLICY-IN mode (refer to [Set a DSCP Value for Egress Packets on page 424](#)) assigns the *same* DSCP value to all of the matching flows in the class-map. The flow-based DSCP marking feature allows you to assign *different* DSCP to each match criteria CLASS-MAP mode using the `set-ip-dscp` option with the `match` command so that matching flows within a class map can have *different* DSCP values ([Figure 24-8](#)). The values you set from CLASS-MAP mode override the QoS input policy DSCP value, and packets matching the rule are marked with the specified value.

Figure 24-8. Marking Flows in the Same Queue with Different DSCP Values

```
FTOS#show run class-map
!
class-map match-any example-flowbased-dscp
  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

FTOS#show run qos-policy-input
!
qos-policy-input flowbased
  set ip-dscp 3
```

Display Configured Class Maps and Match Criteria

To display all class-maps or a specific class map, use the `show qos class-map` command from EXEC Privilege mode.



FTOS Behavior: An explicit *deny any* rule in a Layer 3 ACL used in a (match any or match all) class-map creates a *default to Queue 0* entry in the CAM, which causes unintended traffic classification.

Create a QoS Policy

There are two types of QoS policies: input and output.

Input QoS policies regulate Layer 3 and Layer 2 ingress traffic. The regulation mechanisms for input QoS policies are rate policing and setting priority values. There are two types of input QoS policies: Layer 3 and Layer 2.

- Layer 3 QoS input policies allow you to rate police and set a DSCP.
- Layer 2 QoS input policies allow you to rate police and set a dot1p value.

Output QoS policies regulate Layer 3 egress traffic. The regulation mechanisms for output QoS policies are rate shaping and weighted random early detection (WRED).



Note: When changing a service-queue configuration in a QoS policy map, all QoS rules are deleted and re-added automatically to ensure that the order of the rules is maintained. As a result, the Matched Packets value shown in the `show qos statistics` command is reset.

Create an Input QoS Policy

To create an input QoS policy, follow these steps:

1. Create a Layer 3 input QoS policy using the `qos-policy-input` command from CONFIGURATION mode. Create a Layer 2 input QoS policy by specifying the keyword `layer2` after the `qos-policy-input` command.
2. After you create an input QoS policy, do one or more of the following:
 - [Configure Policy-Based Rate Policing](#)
 - [Set a DSCP Value for Egress Packets](#)
 - [Set a dot1p Value for Egress Packets](#)

Configure Policy-Based Rate Policing

To rate police ingress traffic, use the `rate-police` command from QOS-POLICY-IN mode.

Set a DSCP Value for Egress Packets

You can set the DSCP value for egress packets based on ingress QoS classification. The 6 bits that are used or DSCP are also used to identify the queue in which traffic is buffered.

Figure 24-9. Marking DSCP Values for Egress Packets

```
FTOS#config
FTOS(conf)#qos-policy-input my-input-qos-policy
FTOS(conf-qos-policy-in)#set ip-dscp 34

FTOS(conf-qos-policy-in)#show config
!
qos-policy-input my-input-qos-policy
  set ip-dscp 34
FTOS(conf-qos-policy-in)#end

FTOS#
```

Set a dot1p Value for Egress Packets

To set a dot1p value for egress packets, use the `set mac-dot1p` command from QOS-POLICY-IN mode. When you set a dot1p value, FTOS displays an informational message advising you of the queue to which you should apply the QoS policy (using the command `service-queue` from POLICY-MAP-IN mode).

Create an Output QoS Policy

To create an output QoS policy, follow these steps:

1. Create an output QoS policy using the `qos-policy-output` command from CONFIGURATION mode.
2. After you configure an output QoS policy, do one or more of the following
 - [Configure Policy-Based Rate Shaping](#)
 - [Allocate Bandwidth to the Queue](#)
 - [Configure a Scheduler to Queue](#)
 - [Specify WRED Drop Precedence](#)

Configure Policy-Based Rate Shaping

To rate shape egress traffic, use the rate-shape command from QOS-POLICY-OUT mode.

Allocate Bandwidth to the Queue

To allocate bandwidth, use the bandwidth-percentage command in QOS-POLICY-OUT mode. FTOS recommends that you pre-calculate your bandwidth requirements before creating them. Make sure you apply the QoS policy to all the four queues and that the sum of the bandwidths allocated through them is exactly 100.

When you apply the QoS policies through output policy map and if the sum of the bandwidth percentages configured is below or above 100, then the actual bandwidth is allocated proportionally. If the sum of allocated bandwidth is less than 100, the unused bandwidth is allotted to un-allocated queues. If the sums of allocated bandwidth exceed 100, then 1% of the bandwidth is derived for unassigned queues from assigned queues.

Configure a Scheduler to Queue

By default, the MXL Switch schedules packets for egress based on weighted round robin (WRR). Note that the bandwidth and scheduler cannot be configured at the same time. Policy-level scheduler assigned to queue is applied to both unicast and multicast traffic.

Specify WRED Drop Precedence

To specify a WRED profile to yellow and/or green traffic, use the wred command from QOS-POLICY-OUT mode. For more information, refer to [Apply a WRED Profile to Traffic](#).

Create Policy Maps

There are two types of policy maps: input and output.

Create Input Policy Maps

There are two types of input policy-maps: Layer 3 and Layer 2.

1. To create a Layer 3 input policy map, use the policy-map-input command from CONFIGURATION mode. Create a Layer 2 input policy map by specifying the keyword layer2 with the policy-map-input command.
2. After you create an input policy map, do one or more of the following:
 - [Apply a Class-Map or Input QoS Policy to a Queue](#)
 - [Apply an Input QoS Policy to an Input Policy Map](#)
 - [Honor DSCP Values on Ingress Packets](#)
 - [Honoring dot1p Values on Ingress Packets](#)
 - [Fall Back to trust diffserve or dot1p](#)

3. Apply the input policy map to an interface.

Apply a Class-Map or Input QoS Policy to a Queue

To assign an input QoS policy to a queue, use the `service-queue` command from POLICY-MAP-IN mode.

Apply an Input QoS Policy to an Input Policy Map

To apply an input QoS policy to an input policy map, use the `policy-aggregate` command from POLICY-MAP-IN mode.

Honor DSCP Values on Ingress Packets

FTOS provides the ability to honor DSCP values on ingress packets using the trust DSCP feature. To enable the trust DSCP feature, use the `trust diffserv` command from POLICY-MAP-IN mode. [Table 24-3](#) lists the standard DSCP definitions and indicates to which queues FTOS maps the DSCP values. When you configure trust DSCP, the matched packets and matched bytes counters are not incremented in the `show qos statistics` command.



Note: Packets with DSCP value of 63 are dropped.

Table 24-3. Default DSCP to Queue Mapping

DSCP/CP hex range (XXX)	DSCP Definition	Traditional IP Precedence	Internal Queue ID	DSCP/CP decimal
111XXX		Network Control	3	48–63
110XXX		Internetwork Control	3	
101XXX	EF (Expedited Forwarding)	CRITIC/ECP	2	32–47
100XXX	AF4 (Assured Forwarding)	Flash Override	2	
011XXX	AF3	Flash	1	16–31
010XXX	AF2	Immediate	1	
001XXX	AF1	Priority	0	0–15
000XXX	BE (Best Effort)	Best Effort	0	

Honoring dot1p Values on Ingress Packets

FTOS provides the ability to honor dot1p values on ingress packets with the trust dot1p feature. To enable trust dot1p, use the trust dot1p command from POLICY-MAP-IN mode. [Table 24-4](#) lists the queue to which the classified traffic is sent based on the dot1p value.

Table 24-4. Default dot1p to Queue Mapping

dot1p	Queue ID
0	0
1	0
2	0
3	1
4	2
5	3
6	3
7	3

The dot1p value is also honored for frames on the default VLAN. For more information, refer to [Priority-Tagged Frames on the Default VLAN](#).

Fall Back to trust diffserve or dot1p

When using QoS service policies with multiple class maps, you can configure FTOS to use the incoming DSCP or dot1p marking as a secondary option for packet queuing in the event that no match occurs in the class maps.

When class-maps are used, traffic is matched against each class-map sequentially from first to last. The sequence is based on the priority of the rules, as follows:

1. rules with lowest priority, or in the absence of a priority configuration,
2. rules of the next numerically higher queue

By default, if no match occurs, the packet is queued to the default queue, Queue 0.

In the following configuration, packets are classified to queues using the three class maps:

Figure 24-10. Configuration Example

```

!
policy-map-input input-policy
  service-queue 1 class-map qos-BE1
  service-queue 3 class-map qos-AF3
  service-queue 4 class-map qos-AF4
  trust diffserv fallback
!
class-map match-any qos-AF3
  match ip dscp 24
  match ip access-group qos-AF3-ACL
!
class-map match-any qos-AF4
  match ip dscp 32
  match ip access-group qos-AF4-ACL
!
class-map match-all qos-BE1
  match ip dscp 0
  match ip access-group qos-BE1-ACL

```

The packet classification logic for the above configuration is as follows:

1. Match packets against match-any qos-AF4. If a match exists, queue the packet as AF4 in Queue 4, and if no match exists, go to the next class map.
2. Match packets against match-any qos-AF3. If a match exists, queue the packet as AF3 in Queue 3, and if no match exists, go to the next class map.
3. Match packets against match-all qos-BE1. If a match exists, queue the packet as BE1, and if no match exists, queue the packets to the default queue, Queue 0.

You can optionally classify packets using their DSCP marking, instead of placing packets in Queue 0, if no match occurs. In the above example, if no match occurs against match-all qos-BE1, the classification logic continues:

4. Queue the packet according to the DSCP marking. The DSCP to Queue mapping will be as per the [Table 24-3](#).

The behavior is similar for trust dot1p fallback in a Layer2 input policy map; the dot1p-to-queue mapping is according to [Table 24-4](#).

To enable Fall Back to trust diffserve or dot1p:

Task	Command Syntax	Command Mode
Classify packets according to their DSCP value as a secondary option in case no match occurs against the configured class maps.	trust {diffserve dot1p} fallback	POLICY-MAP-IN

Mapping dot1p Values to Service Queues

All traffic is, by default, mapped to the same queue, Queue 0. If you honor dot1p on ingress, you can create service classes based the queuing strategy in [Table 24-4](#) using the service-class dynamic dot1p command from INTERFACE mode. Apply this queuing strategy globally by entering this command from CONFIGURATION mode.

- All dot1p traffic is mapped to Queue 0 unless you enable the service-class dynamic dot1p command on an interface or globally.
- Layer 2 or Layer 3 service policies supersede dot1p service classes.

Guaranteeing Bandwidth to dot1p-Based Service Queues

To guarantee a minimum bandwidth to queues globally from CONFIGURATION mode, use the service-class bandwidth-weight command. The command is applied in the same way as the bandwidth-weight command in an output QoS policy (refer to [Allocate Bandwidth to the Queue](#)). The bandwidth-percentage command in QOS-POLICY-OUT mode supersedes the service-class bandwidth-weight command.

Apply an Input Policy Map to an Interface

To apply an input policy map to an interface, use the service-policy input command from INTERFACE mode. Specify the keyword layer2 if the policy map you are applying is a Layer 2 policy map; in this case, the INTERFACE must be in Switchport mode. You can apply the same policy map to multiple interfaces, and you can modify a policy map after you apply it.

- You cannot apply an input Layer 2 QoS policy on an interface you also configure with the vlan-stack access command.
- If you apply a service policy that contains an ACL to more than one interface, FTOS uses ACL optimization to conserve CAM space. (See [CAM Optimization](#) for details)

Create Output Policy Maps

1. To create an output policy map, use the policy-map-output command from CONFIGURATION mode.
2. After you create an output policy map, do one or more of the following:
 - [Apply an Output QoS Policy to a Queue](#)
 - [Specify an Aggregate QoS Policy](#)
 - [Apply an Output Policy Map to an Interface](#)
3. Apply the policy map to an interface.

Apply an Output QoS Policy to a Queue

To apply an output QoS policy to queues, use the service-queue command from INTERFACE mode.

Specify an Aggregate QoS Policy

To specify an aggregate QoS policy, use the policy-aggregate command from POLICY-MAP-OUT mode.

Apply an Output Policy Map to an Interface

To apply an output policy map to an interface, use the `service-policy output` command from `INTERFACE` mode. You can apply the same policy map to multiple interfaces, and you can modify a policy map after you apply it.

QoS Rate Adjustment

The Ethernet packet format consists of:

- Preamble: 7-bytes Preamble
- Start Frame Delimiter (SFD): 1 byte
- Destination MAC Address: 6 bytes
- Source MAC Address: 6 bytes
- Ethernet Type/Length: 2 bytes
- Payload: (variable)
- Cyclic Redundancy Check (CRC): 4 bytes
- Inter-frame Gap (IFG): (variable)

By default, for rate policing and shaping, FTOS does not include the Preamble, SFD, or the IFG fields. These fields are overhead; only the fields from MAC Destination Address to the CRC are used for forwarding and are included in these rate metering calculations. You can optionally include overhead fields in rate metering calculations by enabling QoS Rate Adjustment.

QoS rate adjustment is disabled by default with the `no qos-rate-adjust` parameter listed in the `running-configuration`.

Task	Command Syntax	Command Mode
Include a specified number of bytes of packet overhead to include in rate policing and shaping calculations. For example, to include the Preamble and SFD, enter <code>qos-rate-adjust 8</code> . For variable length overhead fields, you must know the number of bytes you want to include.	<code>qos-rate-adjust <i>overhead-bytes</i></code> Default: Disabled Range: 1-31	CONFIGURATION

Strict-Priority Queueing

To assign strict-priority to one unicast queue, 1 to 3, use the `strict-priority` command from CONFIGURATION mode. Strict-priority means that FTOS dequeues all packets from the assigned queue before servicing any other queues.

- The `strict-priority` command supersedes the `bandwidth-percentage` command percentage configurations.
- A queue with strict-priority can starve other queues in the same port-pipe.
- If more than two strict priority queues are configured, the strict priority queue with a higher queue number is scheduled first.


Weighted Random Early Detection

The weighted random early detection (WRED) congestion avoidance mechanism drops packets to prevent buffering resources from being consumed.

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. You can apply a WRED profile to a policy-map so that you can prevent specified traffic from consuming too much of the BTM resources.

WRED uses a profile to specify the 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 are dropped randomly at an exponential rate until the maximum threshold is reached (Figure 24-11); this is the *early detection* part of WRED. If the maximum threshold—2000KB, for example—is reached, all incoming packets are dropped until less than 2000KB of buffer space is consumed by the specified traffic.

Figure 24-11. Packet Drop Rate for WRED



fnC0045mp

You can create a custom WRED profile or use one of the five pre-defined profiles.

Table 24-5. Pre-defined WRED Profiles

Default Profile Name	Minimum Threshold	Maximum Threshold	Maximum Drop Rate
wred_drop	0	0	100
wred_teng_y	467	4671	100
wred_teng_g	467	4671	50
wred_fortyg_y	467	4671	50
wred_fortyg_g	467	4671	25

Create WRED Profiles

To create a WRED profile, follow these steps:

1. To create a WRED profile, use the `wred` command from CONFIGURATION mode.
2. The `wred` command places you in WRED mode. From this mode, specify minimum and maximum threshold values using the `threshold` command.

Apply a WRED Profile to Traffic

After you create a WRED profile, you must specify to which traffic FTOS should apply the profile.

FTOS assigns a color (also called drop precedence)—red, yellow, or green—to each packet based on the DSCP value before queuing it. DSCP is a 6-bit field. If the fourth bit from MSB in DSCP is 0, then it is assigned as green packet. If it is 1, then it is assigned as Yellow packet (except DSCP of 111111, which is a Red packet and is always dropped). If you do not configure FTOS to honor DSCP values on ingress, all traffic defaults to green (refer to [Honor DSCP Values on Ingress Packets](#)).

To assign a WRED profile to either yellow or green traffic, use the `wred` command from QOS-POLICY-OUT mode.

Display Default and Configured WRED Profiles

To display default and configured WRED profiles and their threshold values, use the `show qos wred-profile` command from EXEC mode ([Figure 24-12](#)).

Figure 24-12. Displaying WRED Profiles

```
FTOS#show qos wred-profile

Wred-profile-name      min-threshold  max-threshold  max-drop-rate
wred_drop              0              0              100
wred_teng_y            467            4671           100
wred_teng_g            467            4671           50
```

Display WRED Drop Statistics

To display the number of packets FTOS dropped by the WRED profile, use the `show qos statistics` command from EXEC Privilege mode (Figure 24-13).

Figure 24-13. show qos statistics Command Example

```
FTOS#show qos statistics wred-profile
Interface Te 0/20
Drop-statistic  Dropped Pkts
Green           11234
Yellow         12484
Out of Profile  0
FTOS#
```


Routing Information Protocol (RIP)

Routing Information Protocol (RIP) is based on a distance-vector algorithm. RIP tracks distances or hop counts to nearby routers when establishing network connections.

- [Overview](#)
- [Implementation Information](#)
- [Configuration Information](#)
- [RIP Configuration Example](#)

RIP protocol standards are listed in the [Standards Compliance](#) chapter.

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 learns where nodes in a network are located by automatically constructing a routing data table. The routing table is established after RIP sends out one or more broadcast signals to all adjacent nodes in a network. Hop counts of these signals are tracked and entered into the routing table, which defines where nodes in the network are located.

The information that is used to update the routing table is sent as either a request or a response message. In RIPv1, automatic updates to the routing table are performed as either one-time requests or periodic responses (every 30 seconds). RIP transports its responses or requests by means of the user datagram protocol (UDP) over 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 variable length subnet mask (VLSM) or classless inter-domain routing (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

The Dell Force10 operating software (FTOS) supports both versions of RIP and allows you to configure one version globally and the other version or both versions on the interfaces. [Table 25-1](#) displays the defaults for RIP in FTOS.

Table 25-1. 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

By default, RIP is disabled in FTOS. To configure RIP, you must use commands in two modes: ROUTER and INTERFACE. Commands executed in ROUTER RIP mode configure RIP globally; commands executed in INTERFACE mode configure RIP features on that interface only.

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 are mandatory and optional configuration tasks for RIP.

- [Enable RIP Globally](#) (mandatory)
- [Configure RIP on Interfaces](#) (optional)
- [Control RIP Routing Updates](#) (optional)
- [Set the Send and Receive Version](#) (optional)

- [Generate a Default Route](#) (optional)
- [Control Route Metrics](#) (optional)
- [Summarize Routes](#) (optional)
- [Control Route Metrics](#)
- [Debug RIP](#)

For a complete listing of all commands related to RIP, refer to the *FTOS Command Reference Guide*.

Enable RIP Globally

By default, RIP is not enabled in FTOS. To enable RIP, use the following commands in sequence, starting in 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 <i>ip-address</i>	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 designating the networks that the system is to exchange RIP information, ensure that all devices on that network are configured to exchange RIP information.

The FTOS default is to send RIPv1, and to receive RIPv1 and RIPv2. To change the RIP version globally, use the version command in ROUTER RIP mode.

After you enable RIP, to view the global RIP configuration, use the show running-config command in EXEC mode or the show config command ([Figure 25-1](#)) in ROUTER RIP mode.

Figure 25-1. show config Command Example (ROUTER RIP Mode)

```
FTOS(conf-router_rip)#show config
!
router rip
 network 10.0.0.0
FTOS(conf-router_rip)#
```

When the RIP process has learned the RIP routes, use the `show ip rip database` command in EXEC mode to view those routes (Figure 25-2).

Figure 25-2. show ip rip database Command Example (Partial)

```
FTOS#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
```

To disable RIP globally, use the `no router rip` command in CONFIGURATION mode.

Configure RIP on Interfaces

When you enable RIP globally on the system, interfaces meeting certain conditions start receiving RIP routes. By default, interfaces that you enable and configure 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 to 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 ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
<code>neighbor ip-address</code>	ROUTER RIP	Define a specific router to exchange RIP information between it and the Dell Force10 system. You can use this command multiple times to exchange RIP information with as many RIP networks as you want.
<code>passive-interface interface</code>	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. A prefix list is applied to incoming or outgoing routes. Those routes must meet the conditions of the prefix list; if not, FTOS drops the route. 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, refer to [Access Control Lists \(ACLs\)](#).

To apply prefix lists to incoming or outgoing RIP routes, use the following commands in ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
<code>distribute-list prefix-list-name in</code>	ROUTER RIP	Assign a configured prefix list to all incoming RIP routes.
<code>distribute-list prefix-list-name out</code>	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. To include OSPF, static, or directly connected routes in the RIP process, use the redistribute command.

To add routes from other routing instances or protocols, use any of the following commands in ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
redistribute {connected static} [metric <i>metric-value</i>] [route-map <i>map-name</i>]	ROUTER RIP	Include directly connected or user-configured (static) routes in RIP. <ul style="list-style-type: none"> <i>metric</i> range: 0 to 16 <i>map-name</i>: name of a configured route map.
redistribute ospf <i>process-id</i> [match external {1 2} match internal] [metric <i>value</i>] [route-map <i>map-name</i>]	ROUTER RIP	Include specific OSPF routes in RIP. Configure the following parameters: <ul style="list-style-type: none"> <i>process-id</i> range: 1 to 65535 <i>metric</i> range: 0 to 16 <i>map-name</i>: name of a configured route map.

To view the current RIP configuration, use the show running-config command in EXEC mode or the show config command in ROUTER RIP mode.

Set the Send and Receive Version

To specify the RIP version, use the version command in 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 INTERFACE mode.

To change the RIP version globally in FTOS, use the following command in ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
version {1 2}	ROUTER RIP	Set the RIP version sent and received on the system.

To set one RIP version globally, use the system command. 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.

To see whether the version command is configured, use the show config command in ROUTER RIP mode. To view the routing protocols configuration, use the show ip protocols command in EXEC mode.

Figure 25-3 shows an example of the RIP configuration after you use the version command to set RIPv2 in ROUTER RIP mode. After you set the version command in ROUTER RIP mode, the interface (TenGigabitEthernet 0/0) participating in the RIP process is also set to send and receive RIPv2.

Figure 25-3. show ip protocols Command Example

```

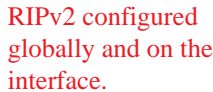
FTOS#show ip protocols

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
  TenGigabitEthernet 0/0  2    2
Routing for Networks:
  10.0.0.0

Routing Information Sources:
Gateway          Distance      Last Update

Distance: (default is 120)

FTOS#
  
```



To configure the interfaces to send or receive different RIP versions from the RIP version configured globally, use either of the following commands in 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 25-4 shows the command syntax for sending both RIPv1 and RIPv2 and receiving only RIPv2.

Figure 25-4. Configuring an Interface to Send Both Versions of RIP

```

FTOS(conf-if)#ip rip send version 1 2
FTOS(conf-if)#ip rip receive version 2
  
```

The show ip protocols command example [Figure 25-5](#) confirms that both versions are sent out that interface. This interface no longer sends and receives the same RIP versions as FTOS does globally.

Figure 25-5. show ip protocols Command Example

```
FTOS#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
Tengigabitethernet 0/0  2     1 2
Routing for Networks:
  10.0.0.0

Routing Information Sources:
Gateway          Distance    Last Update

Distance: (default is 120)

FTOS#
```

RIPv2 configured globally

Different RIP versions configured for this interface

Generate a Default Route

When the traffic's network is not explicitly listed in the routing table, traffic is forwarded to the default route. Default routes are not enabled in RIP unless specified. To generate a default route into RIP, use the default-information originate command in ROUTER RIP mode. In FTOS, default routes received in RIP updates from other routes are advertised if you configured the default-information originate command.

To configure FTOS to generate a default route, use the following command in ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
default-information originate [always] [metric <i>value</i>] [route-map <i>route-map-name</i>]	ROUTER RIP	Specify the generation of a default route in RIP. Configure the following parameters: <ul style="list-style-type: none"> • <i>always</i>: enter this keyword to always generate a default route. • <i>value</i> range: 1 to 16. • <i>route-map-name</i>: name of a configured route map.

To confirm that the default route configuration is completed, use the show config command in ROUTER RIP mode.

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 ROUTER RIP mode is enabled and summarizes RIP routes up to the classful network boundary.

If you must perform routing between discontinuous subnets, disable automatic summarization. With automatic route summarization disabled, subnets are advertised.

The autosummary command requires no other configuration commands. To disable automatic route summarization, use the no autosummary command in ROUTER RIP mode.



Note: If you enable the ip split-horizon command on an interface, the system does not advertise the summarized address.

Control Route Metrics

As a distance-vector protocol, RIP 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.

Exercise caution when applying an offset command to routers on a broadcast network, as the router using the offset command is modifying RIP advertisements before sending out those advertisements.

The distance command also allows you to manipulate route metrics. Use the command to 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 ROUTER RIP mode:

Command Syntax	Command Mode	Purpose
<code>distance weight [ip-address mask [access-list-name]]</code>	ROUTER RIP	Apply a weight to all routes or a specific route and ACL. Configure the following parameters: <ul style="list-style-type: none">• <i>weight</i> range: 1 to 255 (default is 120)• <i>ip-address mask</i>: the IP address in dotted decimal format (A.B.C.D), and the mask in slash format (/x).• <i>access-list-name</i>: name of a configured IP ACL.
<code>offset-list access-list-name {in out} offset [interface]</code>	ROUTER RIP	Apply an additional number to the incoming or outgoing route metrics. Configure the following parameters: <ul style="list-style-type: none">• <i>prefix-list-name</i>: the name of an established prefix list to determine which incoming routes will be modified• <i>offset</i> range: 0 to 16.• <i>interface</i>: the type, slot, and number of an interface.

To view configuration changes, use the show config command in ROUTER RIP mode.

Debug RIP

To enable RIP debugging, use the debug ip rip command. When you enable debugging, you can view information about RIP protocol changes or RIP routes (Figure 25-6).

To enable RIP debugging, use the following command in EXEC privilege mode:

Command Syntax	Command Mode	Purpose
debug ip rip [<i>interface</i> database events trigger]	EXEC privilege	Enable debugging of RIP.

Figure 25-6. debug ip rip Command Example

```
FTOS#debug ip rip
RIP protocol debug is ON
FTOS#
```

To disable RIP, use the no debug ip rip command.

RIP Configuration Example

The example in this section shows the command sequence to configure RIPv2 on the two routers shown in Figure 25-7—“Core 2” and “Core 3”. The host prompts used in the example reflect those names. The examples are divided into the following groups of command sequences:

- [Configuring RIPv2 on Core 2](#)
- [Core 2 Output](#)
- [RIP Configuration on Core 3](#)
- [Core 3 RIP Output](#)
- [RIP Configuration Summary](#)

Figure 25-7. RIP Topology Example



Configuring RIPv2 on Core 2

Figure 25-8. Configuring RIPv2 on Core 2

```
Core2(conf-if-te-2/31)#
Core2(conf-if-te-2/31)#router rip
Core2(conf-router_rip)#ver 2
Core2(conf-router_rip)#network 10.200.10.0
Core2(conf-router_rip)#network 10.300.10.0
Core2(conf-router_rip)#network 10.11.10.0
Core2(conf-router_rip)#network 10.11.20.0
Core2(conf-router_rip)#show config
!
router rip
 network 10.0.0.0
 version 2
Core2(conf-router_rip)#
```

Core 2 Output

The examples in this section are:

- To display the Core 2 RIP database, use the `show ip rip database` command (Figure 25-9).
- To display the Core 2 RIP setup, use the `show ip route` command (Figure 25-10).
- To display the Core 2 RIP activity, use the `show ip protocols` command (Figure 25-11).

Figure 25-9. Example of RIP Configuration Response from Core 2

```

Core2(conf-router_rip)#end
00:12:24: %RPM0-P:CP %SYS-5-CONFIG_I: Configured from console by console
Core2#show ip rip database
Total number of routes in RIP database: 7
10.11.30.0/24
    [120/1] via 10.11.20.1, 00:00:03, TenGigabitEthernet 2/31
10.300.10.0/24      directly connected,TenGigabitEthernet 2/42
10.200.10.0/24      directly connected,TenGigabitEthernet 2/41
10.11.20.0/24       directly connected,TenGigabitEthernet 2/31
10.11.10.0/24       directly connected,TenGigabitEthernet 2/11
10.0.0.0/8          auto-summary
192.168.1.0/24
    [120/1] via 10.11.20.1, 00:00:03, TenGigabitEthernet 2/31
192.168.1.0/24      auto-summary
192.168.2.0/24
    [120/1] via 10.11.20.1, 00:00:03, TenGigabitEthernet 2/31
192.168.2.0/24      auto-summary

Core2#

```

Figure 25-10. show ip route Command Example to Show RIP Configuration on Core 2

```

Core2#show ip route

Codes: C - connected, 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, L1 - IS-IS level-1,
       L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
       > - non-active route, + - summary route

Gateway of last resort is not set

Destination          Gateway                Dist/Metric Last Change
-----
C    10.11.10.0/24       Direct, TenGig 2/11    0/0      00:02:26
C    10.11.20.0/24       Direct, TenGig 2/31    0/0      00:02:02
R    10.11.30.0/24       via 10.11.20.1, TenGig 2/31    120/1    00:01:20
C    10.200.10.0/24      Direct, TenGig 2/41    0/0      00:03:03
C    10.300.10.0/24      Direct, TenGig 2/42    0/0      00:02:42
R    192.168.1.0/24      via 10.11.20.1, TenGig 2/31    120/1    00:01:20
R    192.168.2.0/24      via 10.11.20.1, TenGig 2/31    120/1    00:01:20
Core2#
R    192.168.1.0/24      via 10.11.20.1, TenGig 2/31    120/1    00:05:22
R    192.168.2.0/24      via 10.11.20.1, TenGig 2/31    120/1    00:05:22

Core2#

```


Figure 25-11. show ip protocols Command Example to Show RIP Configuration Activity on Core 2

```
Core2#show ip protocols
Routing Protocol is "RIP"
  Sending updates every 30 seconds, next due in 17
  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
    TenGigabitEthernet 2/42    2    2
    TenGigabitEthernet 2/41    2    2
    TenGigabitEthernet 2/31    2    2
    TenGigabitEthernet 2/11    2    2
  Routing for Networks:
    10.300.10.0
    10.200.10.0
    10.11.20.0
    10.11.10.0

  Routing Information Sources:
  Gateway          Distance    Last Update
  10.11.20.1        120         00:00:12

  Distance: (default is 120)

Core2#
```

RIP Configuration on Core 3

Figure 25-12. RIP Configuration on Core 3

```
Core3(conf-if-te-3/21)#router rip
Core3(conf-router_rip)#version 2
Core3(conf-router_rip)#network 192.168.1.0
Core3(conf-router_rip)#network 192.168.2.0
Core3(conf-router_rip)#network 10.11.30.0
Core3(conf-router_rip)#network 10.11.20.0
Core3(conf-router_rip)#show config
!
router rip
 network 10.0.0.0
 network 192.168.1.0
 network 192.168.2.0
 version 2
Core3(conf-router_rip)#
```

Core 3 RIP Output

The examples in this section are:

- To display the Core 3 RIP database, use the show ip rip database command (Figure 25-13).
- To display the Core 3 RIP setup, use the show ip route command (Figure 25-14).
- To display the Core 3 RIP activity, use the show ip protocols command (Figure 25-15).

Figure 25-13. show ip rip database Command Example for Core 3 RIP Setup

```
Core3#show ip rip database
Total number of routes in RIP database: 7
10.11.10.0/24
    [120/1] via 10.11.20.2, 00:00:13, TenGigabitEthernet 3/21
10.200.10.0/24
    [120/1] via 10.11.20.2, 00:00:13, TenGigabitEthernet 3/21
10.300.10.0/24
    [120/1] via 10.11.20.2, 00:00:13, TenGigabitEthernet 3/21
10.11.20.0/24      directly connected,TenGigabitEthernet 3/21
10.11.30.0/24     directly connected,TenGigabitEthernet 3/11
10.0.0.0/8        auto-summary
192.168.1.0/24    directly connected,TenGigabitEthernet 3/43
192.168.1.0/24    auto-summary
192.168.2.0/24    directly connected,TenGigabitEthernet 3/44
192.168.2.0/24    auto-summary
Core3#
```

Figure 25-14. show ip routes Command Example for Core 3 RIP Setup

```
Core3#show ip routes

Codes: C - connected, 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,
       > - non-active route, + - summary route

Gateway of last resort is not set

      Destination            Gateway                      Dist/Metric  Last Change
      -----            -
R    10.11.10.0/24          via 10.11.20.2, TenGig 3/21      120/1       00:01:14
C    10.11.20.0/24          Direct, TenGig 3/21              0/0         00:01:53
C    10.11.30.0/24          Direct, TenGig 3/11              0/0         00:06:00
R    10.200.10.0/24         via 10.11.20.2, TenGig 3/21      120/1       00:01:14
R    10.300.10.0/24         via 10.11.20.2, TenGig 3/21      120/1       00:01:14
C    192.168.1.0/24         Direct, TenGig 3/43              0/0         00:06:53
C    192.168.2.0/24         Direct, TenGig 3/44              0/0         00:06:26
Core3#
```

Figure 25-15. show ip protocols Command Example to Show RIP Configuration Activity on Core 3

```
Core3#show ip protocols

Routing Protocol is "RIP"
  Sending updates every 30 seconds, next due in 6
  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
    TenGigabitEthernet 3/21  2    2
    TenGigabitEthernet 3/11  2    2
    TenGigabitEthernet 3/44  2    2
    TenGigabitEthernet 3/43  2    2
  Routing for Networks:
    10.11.20.0
    10.11.30.0
    192.168.2.0
    192.168.1.0

  Routing Information Sources:
  Gateway          Distance    Last Update
  10.11.20.2        120         00:00:22

  Distance: (default is 120)

Core3#
```

RIP Configuration Summary

Figure 25-16. Summary of Core 2 RIP Configuration Using Output of show run Command

```
!  
interface TenGigabitEthernet 2/11  
ip address 10.11.10.1/24  
no shutdown  
!  
interface TenGigabitEthernet 2/31  
ip address 10.11.20.2/24  
no shutdown  
!  
interface TenGigabitEthernet 2/41  
ip address 10.200.10.1/24  
no shutdown  
!  
interface TenGigabitEthernet 2/42  
ip address 10.300.10.1/24  
no shutdown  
  
router rip  
version 2  
10.200.10.0  
10.300.10.0  
10.11.10.0  
10.11.20.0
```

Figure 25-17. Summary of Core 3 RIP Configuration Using Output of show run Command

```
!  
interface TenGigabitEthernet 3/11  
ip address 10.11.30.1/24  
no shutdown  
!  
interface TenGigabitEthernet 3/21  
ip address 10.11.20.1/24  
no shutdown  
!  
interface TenGigabitEthernet 3/43  
ip address 192.168.1.1/24  
no shutdown  
!  
interface TenGigabitEthernet 3/44  
ip address 192.168.2.1/24  
no shutdown  
!  
router rip  
version 2  
network 10.11.20.0  
network 10.11.30.0  
network 192.168.1.0  
network 192.168.2.0
```


Remote Monitoring (RMON)

Overview

This chapter describes remote monitoring (RMON). This chapter includes the following sections:

- [Implementation](#)
- [Fault Recovery](#)

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 Dell Force10 Ethernet interfaces.

RMON operates with the simple network management protocol (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 management information base (MIBs).

Implementation

You must configure SNMP prior to setting up RMON. For a complete SNMP implementation description, refer to [Simple Network Management Protocol \(SNMP\)](#).

Configuring RMON requires using the RMON command line interface (CLI) and includes the following tasks:

- [Set the RMON Alarm](#)
- [Configure an RMON Event](#)
- [Configure RMON Collection Statistics](#)
- [Configure RMON Collection History](#)
- [Enable an RMON MIB Collection History Group](#)

RMON implements the following standard request for comment (RFCs) (for more information, refer to [RFC and I-D Compliance](#)):

- 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.


Set the 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 these commands:

Command Syntax	Command Mode	Purpose
<pre>[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]</pre>	CONFIGURATION	<p>Set an alarm on any MIB object. Use the <code>no</code> form of this command to disable the alarm.</p> <p>Configure the alarm using the following optional parameters:</p> <ul style="list-style-type: none">• <i>number</i>: Alarm number, must be an integer from 1 to 65,535, the value must be unique in the RMON alarm table.• <i>variable</i>: 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 <code>rmon alarm</code> command and 64 bits for the <code>rmon hc-alarm</code> command.• <i>interval</i>: Time in seconds the alarm monitors the MIB variable, the value must be between 5 to 3,600.• <i>delta</i>: Tests the change between MIB variables, this is the <i>alarmSampleType</i> in the RMON alarm table.• <i>absolute</i>: Tests each MIB variable directly, this is the <i>alarmSampleType</i> in the RMON alarm table.• <i>rising-threshold value</i>: Value at which the rising-threshold alarm is triggered or reset. For the <code>rmon alarm</code> command, this is a 32-bits value, for <code>rmon hc-alarm</code> command, this is a 64-bits value.• <i>event-number</i>: Event number to trigger when the rising threshold exceeds its limit. This value is identical to the <i>alarmRisingEventIndex</i> in the <i>alarmTable</i> of the RMON MIB. If there is no corresponding rising-threshold event, the value should be zero.• <i>falling-threshold value</i>: Value at which the falling-threshold alarm is triggered or reset. For the <code>rmon alarm</code> command, this is a 32-bits value, for <code>rmon hc-alarm</code> command, this is a 64bits value.• <i>event-number</i>: Event number to trigger when the falling threshold exceeds its limit. This value is identical to the <i>alarmFallingEventIndex</i> in the <i>alarmTable</i> of the RMON MIB. If there is no corresponding falling-threshold event, the value should be zero.• <i>owner string</i>: (Optional) Specifies an owner for the alarm, this is the <i>alarmOwner</i> object in the <i>alarmTable</i> of the RMON MIB. Default is a null-terminated string.

To configure an RMON alarm, use the `rmon alarm` command (Figure 26-1).

Figure 26-1. 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.

Configure an RMON Event

To add an event in the RMON event table, use the `rmon event` command in GLOBAL CONFIGURATION mode (Figure 26-2). To disable RMON on the interface, use the `no rmon event` command:

Command Syntax	Command Mode	Purpose
<code>[no] rmon event number [log] [trap community] [description string] [owner string]</code>	CONFIGURATION	<p><i>number</i>: 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.</p> <p><i>log</i>: (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.</p> <p><i>trap community</i>: (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.</p> <p><i>description string</i>: (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.</p> <p><i>owner string</i>: (Optional) Owner of this event, which is identical to the eventOwner in the eventTable of the RMON MIB. Default is a null-terminated string.</p>

Figure 26-2. rmon event Command Example

```
FTOS(conf)#rmon event 1 log trap eventtrap description "High ifOutErrors" owner nms1
```

The configuration in [Figure 26-2](#) 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”.

Configure RMON Collection Statistics

To enable RMON MIB statistics collection on an interface, use the RMON collection statistics command in CONFIGURATION INTERFACE (conf-if) mode. To remove a specified RMON statistics collection, use the no RMON collection statistics command.

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

The command in [Figure 26-3](#) enables RMON statistics collection on the interface with an ID value of 20 and an owner of “john.”

Figure 26-3. rmon collection statistics Command Example

```
FTOS(conf-if-te-0/40)#rmon collection statistics controlEntry 20 owner john
```

Configure RMON Collection History

To enable the RMON MIB history group of statistics collection on an interface, use the `rmon collection history` command in `CONFIGURATION INTERFACE (conf-if)` mode. To remove a specified RMON history group of statistics collection, use the `no rmon collection history` command.

Command Syntax	Command Mode	Purpose
<code>[no] rmon collection history {controlEntry integer} [owner owner-string] [buckets bucket-number] [interval seconds]</code>	CONFIGURATION INTERFACE (conf-if)	<p><i>controlEntry</i>: Specifies the RMON group of statistics using a value.</p> <p><i>integer</i>: 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.</p> <p><i>owner</i>: (Optional) Specifies the name of the owner of the RMON group of statistics. The default is a null-terminated string.</p> <p><i>owner-string</i>: (Optional) Records the name of the owner of the RMON group of statistics.</p> <p><i>buckets</i>: (Optional) Specifies the maximum number of buckets desired for the RMON collection history group of statistics.</p> <p><i>bucket-number</i>: (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).</p> <p><i>interval</i>: (Optional) Specifies the number of seconds in each polling cycle.</p> <p><i>seconds</i>: (Optional) The number of seconds in each polling cycle. The value is ranged from 5 to 3,600 (Seconds). The default is 1,800 as defined in RFC-2819.</p>

Enable an RMON MIB Collection History Group

The command in [Figure 26-4](#) 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.

Figure 26-4. rmon collection history Command Example

```
FTOS(conf-if-te-0/40)#rmon collection history controlEntry 20 owner john
```


Rapid Spanning Tree Protocol (RSTP)

Overview

Rapid spanning tree protocol (RSTP) is a Layer 2 protocol—specified by IEEE 802.1w—that is essentially the same as the spanning-tree protocol (STP) but provides faster convergence and interoperability with switches configured with STP and multiple spanning tree protocol (MSTP).

FTOS supports three other variations of spanning tree ([Table 27-1](#)).

Table 27-1. FTOS Supported Spanning Tree Protocols

Dell Force10 Term	IEEE Specification
Spanning Tree Protocol (STP)	802.1d
Rapid Spanning Tree Protocol (RSTP)	802.1w
Multiple Spanning Tree Protocol (MSTP)	802.1s
Per-VLAN Spanning Tree Plus (PVST+)	Third Party

Configuring Rapid Spanning Tree

Configuring rapid spanning tree is a two-step process:

1. Configure interfaces for Layer 2.
2. Enable Rapid Spanning Tree Protocol.

Related Configuration Tasks

- [Add and Remove Interfaces](#)
- [Modify Global Parameters](#)
- [Enable BPDU Filtering globally](#)
- [Modify Interface Parameters](#)
- [Configure an EdgePort](#)
- [Preventing Network Disruptions with BPDU Guard](#)
- [Influence RSTP Root Selection](#)

- SNMP Traps for Root Elections and Topology Changes
- Fast Hellos for Link State Detection
- Flush MAC Addresses after a Topology Change

Important Points to Remember

- RSTP is disabled by default.
- FTOS supports only one RST instance.
- All interfaces in virtual local area networks (VLANs) and all enabled interfaces in Layer 2 mode are automatically added to the RST topology.
- Avoid using the range command to add a large group of ports to a large group of VLANs; adding a group of ports to a range of VLANs sends multiple messages to the RSTP task. When using the range command, Dell Force10 recommends limiting the range to five ports and 40 VLANs.

Configure Interfaces for Layer 2 Mode


All interfaces on all bridges that participates in RST must be in Layer 2 and enabled.

Figure 27-1. Configuring Interfaces for Layer 2 Mode

```

FT05(conf)#interface range tengigabitethernet 1/1 - 4
FT05(conf-if-range-te-1/1-4)#switchport
FT05(conf-if-range-te-1/1-4)#no shutdown
FT05(conf-if-range-te-1/1-4)#show config
!
interface TenGigabitEthernet 1/1
no ip address
switchport
no shutdown
!
interface TenGigabitEthernet 1/2
no ip address
switchport
no shutdown
!
interface TenGigabitEthernet 1/3
no ip address
switchport
no shutdown
!
interface TenGigabitEthernet 1/4
no ip address
switchport
no shutdown
FT05(conf-if-range-te-1/1-4)#

```



To configure and enable the interfaces for Layer 2, use the following commands:

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

To verify that an interface is in Layer 2 mode and enabled, use the show config command from INTERFACE mode.

Figure 27-2. Verifying Layer 2 Configuration

```

FTOS(conf-if-te-1/1)#show config
!
interface TenGigabitEthernet 1/1
  no ip address
  switchport
  no shutdown
FTOS(conf-if-te-1/1)#
  
```

← Indicates that the interface is in Layer 2 mode

Enable Rapid Spanning Tree Protocol Globally

You must enable RSTP globally on all participating bridges; it is not enabled by default.

To enable RSTP globally for all Layer 2 interfaces, use the following commands:

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, use the disable command from PROTOCOL SPANNING TREE RSTP mode.

To verify that RSTP is enabled, use the show config command from PROTOCOL SPANNING TREE RSTP mode (Figure 27-3).

Figure 27-3. Verifying RSTP is Enabled


```
FTOS(conf-rstp)#show config
!
protocol spanning-tree rstp
no disable
FTOS(conf-rstp)#
```

Indicates that Rapid Spanning Tree is enabled

When you enable RST, all physical and port-channel interfaces that are enabled and in Layer 2 mode are automatically part of the RST topology (Figure 27-4).

- Only one path from any bridge to any other bridge is enabled.
- Bridges block a redundant path by disabling one of the link ports.

Figure 27-4. Rapid Spanning Tree Enabled Globally



```
Port 684 (TenGigabitEthernet 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
```

To view the interfaces participating in RST, use the `show spanning-tree rstp` command from EXEC privilege mode (Figure 27-5). If a physical interface is part of a port channel, only the port channel is listed in the command output.

Figure 27-5. show spanning-tree rstp Command Example

```
FTOS#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 occurred 00:02:17 ago on TenGig 1/26

Port 377 (TenGigabitEthernet 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, bpdu filter is disabled

Port 378 (TenGigabitEthernet 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, bpdu filter is disabled

Port 379 (TenGigabitEthernet 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, bpdu filter is disabled

Port 380 (TenGigabitEthernet 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, bpdu filter is disabled

FTOS#
```

To confirm that a port is participating in RST, use the `show spanning-tree rstp brief` command from EXEC privilege mode (Figure 27-6).

Figure 27-6. show spanning-tree rstp brief Command Example

```
FTOS#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
Bpdu filter disabled globally
```

Interface Name	PortID	Prio	Cost	Sts	Cost	Designated Bridge ID	PortID
TenGig 3/1	128.681	128	20000	BLK	20000	32768 0001.e80b.88bd	128.469
TenGig 3/2	128.682	128	20000	BLK	20000	32768 0001.e80b.88bd	128.470
TenGig 3/3	128.683	128	20000	FWD	20000	32768 0001.e801.cbb4	128.379
TenGig 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	Bpdu Filter
TenGig 3/1	Altr	128.681	128	20000	BLK	20000	P2P	No	No
TenGig 3/2	Altr	128.682	128	20000	BLK	20000	P2P	No	No
TenGig 3/3	Root	128.683	128	20000	FWD	20000	P2P	No	No
TenGig 3/4	Altr	128.684	128	20000	BLK	20000	P2P	No	No

```
FTOS#
```

Add and Remove Interfaces

- To add an interface to the RST topology, configure it for Layer 2 and it is automatically added. If you previously disabled RSTP on the interface using the command `no spanning-tree 0`, re-enable it using the `spanning-tree 0` command.
- To remove an interface from the RST topology, use the `no spanning-tree 0` command. For bridge protocol data units (BPDU) filtering behavior, refer to [Removing an Interface from the Spanning Tree Group](#).

Modify Global Parameters

You can modify the RST 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 RST 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: Dell Force10 recommends that only experienced network administrators change the RST group parameters. Poorly planned modification of the RSTG parameters can negatively impact network performance.

Table 27-2 lists the default values for RSTP.

Table 27-2. RSTP Default Values

RSTP Parameter		Default Value
Forward Delay		15 seconds
Hello Time		2 seconds
Max Age		20 seconds
Port Cost	40-Gigabit Ethernet interfaces	1400
	10-Gigabit Ethernet interfaces	2000
	Port Channel with two 40-Gigabit Ethernet interfaces	600
	Port Channel with two 10-Gigabit Ethernet interfaces	1800
Port Priority		128

To change these parameters, use the following commands, on the root bridge:


Task	Command Syntax	Command Mode
Change the forward-delay parameter. <ul style="list-style-type: none"> • Range: 4 to 30 • Default: 15 seconds 	<code>forward-delay <i>seconds</i></code>	PROTOCOL SPANNING TREE RSTP
Change the hello-time parameter. Note: With large configurations (especially those with more ports) Dell Force10 recommends that you increase the hello-time. Range: 1 to 10 Default: 2 seconds	<code>hello-time <i>seconds</i></code>	PROTOCOL SPANNING TREE RSTP
Change the max-age parameter. Range: 6 to 40 Default: 20 seconds	<code>max-age <i>seconds</i></code>	PROTOCOL SPANNING TREE RSTP

To view the current values for global parameters, use the `show spanning-tree rstp` command from EXEC privilege mode (Figure 27-5).

Enable BPDU Filtering globally

The enabling of BPDU Filtering stops transmitting of BPDUs on the operational port fast enabled ports by default. When BPDUs are received, the spanning tree is automatically prepared. By default global bpdu filtering is disabled.

Figure 27-7. BPDU Filtering enabled globally



Task	Command Syntax	Command Mode
Enable BPDU Filter globally to filter transmission of BPDU port fast enabled interfaces.	edge-port bpdu filter default	PROTOCOL SPANNING TREE RSTP

Modify 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 27-2](#). The greater the port cost, the less likely the port is selected to be a forwarding port.
- **Port priority** influences the likelihood that a port is 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: refer to Table 27-2 .	spanning-tree rstp cost <i>cost</i>	INTERFACE
Change the port priority of an interface. Range: 0 to 240 Default: 128	spanning-tree rstp priority <i>priority-value</i>	INTERFACE

To view the current values for interface parameters, use the show spanning-tree rstp command from EXEC privilege mode ([Figure 27-5](#)).

Configure an EdgePort

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. The bpduguard shutdown-on-violation option causes the interface hardware to be shutdown when it receives a BPDU. When you implement only bpduguard, although the interface is placed in an Error Disabled state when receiving the BPDU, the physical interface remains up and spanning-tree drops packets in the hardware after a BPDU violation. BPDUs are dropped in the software after receiving the BPDU violation. 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 you enable it on an interface connected to a network.

Enabling BPDU Filtering on an interface, stops sending and receiving of BPDUs on the port fast enabled ports. When BPDU guard and BPDU filter is enabled on the port, BPDU filter takes the highest precedence. By default bpduguard on an interface is disabled.

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 [bpduguard [shutdown-on-violation] bpduguard]	INTERFACE

To verify that EdgePort is enabled on a port, use the show spanning-tree rstp command from EXEC privilege mode or the show config command from INTERFACE mode. Dell Force10 recommends using the show config command ([Figure 27-8](#)).



FTOS Behavior: Regarding bpduguard shutdown-on-violation behavior:

- 1 If the interface to be shutdown is a port channel, all the member ports are disabled in the hardware.
- 2 When a physical port is added to a port channel already in the Error Disable state, the new member port is also disabled in the hardware.
- 3 When a physical port is removed from a port channel in the Error Disable state, the error disabled state is cleared on this physical port (the physical port is enabled in the hardware).
- 4 You can clear the Error Disabled state with any of the following methods:
 - Perform an shutdown command on the interface.
 - Disable the shutdown-on-violation command on the interface (no spanning-tree rstp edge-port [bpduguard | [shutdown-on-violation]]).
 - Disable spanning tree on the interface (no spanning-tree in INTERFACE mode).
- 5 Disabling global spanning tree (no spanning-tree in CONFIGURATION mode).

Figure 27-8. EdgePort Enabled on Interface

```
FTOS(conf-if-te-2/0)#show config
!
interface TenGigabitEthernet 2/0
no ip address
switchport
spanning-tree rstp edge-port ← Indicates the interface is in EdgePort mode
shutdown
FTOS(conf-if-te-2/0)#
```

Influence RSTP Root Selection

The RSTP determines the root bridge, but you can assign one bridge a lower priority to increase the likelihood so that it is 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. <i>priority-value</i> 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.	<code>bridge-priority <i>priority-value</i></code>	PROTOCOL SPANNING TREE RSTP

A console message appears when a new root bridge has been assigned. [Figure 27-9](#) shows the console message after you use the bridge-priority command to make R2 the root bridge.

Figure 27-9. bridge-priority Command Example

```

FTOS(conf-rstp)#bridge-priority 4096
FTOS(conf-rstp)#2d0h22mi: %STKUNIT3-M:CP: %SPANMGR-5-STP_ROOT_CHANGE: RSTP root changed. My
Bridge ID: 4096:001e.c9f1.00cf Old Root: 32768:0001.e88a.fdb3 New Root: 4096:001e.c9f1.00cf
    
```

↑
↑
Old root bridge ID
New root bridge ID

SNMP Traps for Root Elections and Topology Changes

To enable SNMP traps for RSTP, MSTP, and PVST+ collectively, use the `snmp-server enable traps xstp` command.

Fast Hellos for Link State Detection

Use RSTP fast hellos to achieve sub-second link-down detection so that convergence is triggered faster. The standard RSTP link-state detection mechanism does not offer the same low link-state detection speed.

RSTP fast hellos decrease the hello interval to the order of milliseconds and all timers derived from the hello timer are adjusted accordingly. This feature does not inter-operate with other vendors, and is available only for RSTP.

To configure the hello time, use the following command:

Task	Command Syntax	Command Mode
Configure a hello time on the order of milliseconds.	<code>hello-time milli-second <i>interval</i></code> Range: 50 - 950 milliseconds	PROTOCOL RSTP

```

FTOS(conf-rstp)#do show spanning-tree rstp brief
Executing IEEE compatible Spanning Tree Protocol
Root ID      Priority 0, Address 0001.e811.2233
Root Bridge hello time 50 ms, max age 20, forward delay 15
Bridge ID    Priority 0, Address 0001.e811.2233
We are the root
Configured hello time 50 ms, max age 20, forward delay 15
    
```



Note: The hello time is encoded in BPDUs in increments of 1/256ths of a second. The standard minimum hello time in seconds is 1 second, which is encoded as 256. Millisecond hello times are encoded using values less than 256; the millisecond hello time equals $(x/1000)*256$. When you configure the millisecond hellos, the default hello interval of two seconds is still used for edge ports; the millisecond hello interval is not used.

Security

This chapter describes the following:

- AAA Accounting
- AAA Authentication
- AAA Authorization
- RADIUS
- TACACS+
- Protection from TCP Tiny and Overlapping Fragment Attacks
- SCP and SSH
- Telnet
- VTY Line and Access-Class Configuration

For details about all the commands described in this chapter, refer to the Security Commands chapter in the *FTOS Command Reference Guide*.

AAA Accounting

AAA accounting is part of the AAA security model (accounting, authentication, and authorization), which includes services for authentication, authorization, and accounting. For details about commands related to AAA security, refer to the Security chapter in the *FTOS Command Reference Guide*.

AAA accounting enables tracking of services that users are accessing and the amount of network resources being consumed by those services. When you enable AAA accounting, the network server reports user activity to the security server in the form of accounting records. Each accounting record is comprised of accounting attribute/value (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 sections present the AAA accounting configuration tasks:

- Enable AAA Accounting (mandatory)
- Suppress AAA Accounting for Null Username Sessions (optional)
- Configure Accounting of EXEC and Privilege-Level Command Usage (optional)

- [Configure AAA Accounting for Terminal Lines](#) (optional)
- [Monitor AAA Accounting](#) (optional)

Enable AAA Accounting

To create a record for any or all of the accounting functions monitored, use the `aaa accounting` command. To enable AAA accounting, perform the following task in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
<pre>aaa accounting {system exec command level} {default name} {start-stop wait-start stop-only} {tacacs+}</pre>	CONFIGURATION	<p>Enable AAA accounting and create a record for monitoring the accounting function.</p> <p>The variables are:</p> <ul style="list-style-type: none"> • <code>system</code>—sends accounting information of any other AAA configuration. • <code>exec</code>—sends accounting information when a user has logged in to the EXEC mode. • <code>command level</code>—sends accounting of commands executed at the specified privilege level. • <code>default name</code>—Enter the name of a list of accounting methods. • <code>start-stop</code>—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. • <code>wait-start</code>—ensures that the TACACS+ security server acknowledges the start notice before granting the user's process request. • <code>stop-only</code>—Use for minimal accounting; instructs the TACACS+ server to send a stop record accounting notice at the end of the requested user process. • <code>tacacs+</code> —Designate the security service. Currently, FTOS supports only TACACS+

Suppress AAA Accounting for Null Username Sessions

When you activate AAA accounting, 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
<pre>aaa accounting suppress null-username</pre>	CONFIGURATION	Prevent accounting records from being generated for users whose username string is NULL

Configure Accounting of EXEC and Privilege-Level Command Usage

The network access server monitors the accounting functions defined in the terminal access controller access control system (TACACS+) attribute/value (AV) pairs.

In [Figure 28-1](#), AAA accounting is set to track all usage of EXEC commands and commands on privilege level 15.

Figure 28-1. AAA Accounting Tracking All Usage of EXEC Commands

```
FTOS(conf)#aaa accounting exec default start-stop  
tacacs+  
FTOS(conf)#aaa accounting command 15 default  
start-stop tacacs+
```

System accounting can use only the default method list: `aaa accounting system default start-stop tacacs+`.

Configure AAA Accounting for Terminal Lines

To enable accounting with a named method list for a specific terminal line (where `com15` and `execAcct` are the method list names), use the accounting commands and accounting exec commands ([Figure 28-2](#)).

Figure 28-2. accounting and accounting exec Command Example

```
FTOS(conf-line-vty)# accounting commands 15 com15  
FTOS(conf-line-vty)# accounting exec execAcct
```

Monitor AAA Accounting

The Dell Force10 operating software (FTOS) does not support periodic interim accounting because the periodic command can cause heavy congestion when many users are logged into the network.

No specific show command exists for TACACS+ accounting. To obtain accounting records displaying information about users currently logged in, perform the following task in Privileged EXEC mode ([Figure 28-3](#)):

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

Figure 28-3. show accounting Command Example

```

FTOS#show accounting
Active accounted actions on tty2, User admin Priv 1
  Task ID 1, EXEC Accounting record, 00:00:39 Elapsed, service=shell
Active accounted actions on tty3, User admin Priv 1
  Task ID 2, EXEC Accounting record, 00:00:26 Elapsed, service=shell

```

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 Dell Force10 implementation, the system acts as a remote authentication dial-in service (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.

Dell Force10 uses local usernames/passwords (stored on the Dell Force10 system) or 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.



Note: If a console user logs in with RADIUS authentication, the privilege level will be applied from the RADIUS server if the privilege level is configured for that user in RADIUS whether or not RADIUS authorization is configured.

Configuration Task List for AAA Authentication

The following sections provide the configuration tasks:

- [Configure Login Authentication for Terminal Lines](#)
- [Configure AAA Authentication Login Methods](#)
- [Enable AAA Authentication](#)
- [AAA Authentication—RADIUS](#)

For a complete listing of all commands related to login authentication, refer to the Security chapter in the *FTOS Command Reference Guide*.

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.

Configure AAA Authentication Login Methods

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

Step	Command Syntax	Command Mode	Purpose
1	aaa authentication login {method-list-name default} <i>method1</i> [... <i>method4</i>]	CONFIGURATION	Define an authentication method-list (<i>method-list-name</i>) or specify the default. The default method-list is applied to all terminal lines. Possible methods are: <ul style="list-style-type: none">• enable—use the password defined by the <code>enable secret</code> or <code>enable password</code> command in CONFIGURATION mode.• line—use the password defined by the <code>password</code> command in 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 <code>radius-server</code> host command.• tacacs+—use the TACACS+ server(s) configured with the <code>tacacs-server</code> host command.
2	line {aux 0 console 0 vty <i>number</i> [... <i>end-number</i>]}	CONFIGURATION	Enter LINE mode.
3	login authentication { <i>method-list-name</i> default}	LINE	Assign a <i>method-list-name</i> or the default list to the terminal line.



FTOS Behavior: If you use a method list on the console port in which RADIUS or TACACS+ is the last authentication method, and the server is not reachable, FTOS allows access even though the username and password credentials cannot be verified. Only the console port behaves this way, and does so to ensure that users are not locked out of the system in the event that network-wide issue prevents access to these servers.

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



Note: Dell Force10 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 secure shell (SSH).

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

Enable AAA Authentication

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

Command Syntax	Command Mode	Purpose
<code>aaa authentication enable {<i>method-list-name</i> default} <i>method1</i> [... <i>method4</i>]</code>	CONFIGURATION	<ul style="list-style-type: none"> <code>default</code>—Uses the listed authentication methods that follows this argument as the default list of methods when a user logs in. <code><i>method-list-name</i></code>—Character string used to name the list of enable authentication methods activated when a user logs in. <code><i>method1</i> [... <i>method4</i>]</code>—Any of the following: RADIUS, TACACS, enable, line, none.

If you do not set the default list, only the local enable is checked. This has the same effect as issuing the `aaa authentication enable default enable` command.

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	<code>aaa authentication enable default radius tacacs</code>	CONFIGURATION	To enable RADIUS and to set up TACACS as backup.
2	<code>radius-server host x.x.x.x key some-password</code>	CONFIGURATION	To establish host address and password.
3	<code>tacacs-server host x.x.x.x key some-password</code>	CONFIGURATION	To establish host address and password.

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

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


To use local authentication for enable secret on the console, while using remote authentication on virtual terminal line (VTY) lines, use the following commands:

```
FTOS(conf)# aaa authentication enable mymethodlist radius tacacs
FTOS(conf)# line vty 0 9
FTOS(conf-line-vty)# enable authentication mymethodlist
```

Server-Side Configuration

TACACS+: When using TACACS+, Dell 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\$. When using RADIUS authentication, 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 username.

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 system is one method of protecting the system 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 system.

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 three privilege levels and you can configure 13 more. The three pre-configured levels are:

- **Privilege level 1**—is the default level for 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. For more information about configuring user names, refer to [Configure a Username and Password](#).

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 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 system and when you enter the enable command. You can 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 has the configuration tasks for privilege levels and passwords.

- [Configure a Username and Password](#) (mandatory)
- [Configure the Enable Password Command](#) (mandatory)
- [Configure Custom Privilege Levels](#) (mandatory)
- [Specify the LINE Mode Password and Privilege](#) (optional)
- [Enable and Disable Privilege Levels](#) (optional)

For a complete listing of all commands related to FTOS privilege levels and passwords, refer to the Security chapter in the *FTOS Command Reference Guide*.

Configure a Username and Password

In FTOS, you can assign a specific username to limit user access to the system.

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

Command Syntax	Command Mode	Purpose
<code>username <i>name</i> [access-class <i>access-list-name</i>] [nopassword password [<i>encryption-type</i>] <i>password</i>] [privilege <i>level</i>]</code>	CONFIGURATION	Assign a user name and password. Configure the optional and required parameters: <ul style="list-style-type: none">• <i>name</i>: Enter a text string up to 63 characters long.• <i>access-class access-list-name</i>: Enter the name of a configured IP ACL.• <i>nopassword</i>: Do not require the user to enter a password.• <i>encryption-type</i>: Enter 0 for plain text or 7 for encrypted text.• <i>password</i>: Enter a string.• <i>privilege level</i> range: 0 to 15.

To view usernames, use the show users command in EXEC Privilege mode.

Configure the Enable Password Command

To configure FTOS, you must use the enable command to enter 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 CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
<code>enable password [level <i>level</i>] [<i>encryption-mode</i>] <i>password</i></code>	CONFIGURATION	Configure a password for a privilege level. Configure the optional and required parameters: <ul style="list-style-type: none">• <i>level level</i>: Specify a level 0 to 15. Level 15 includes all levels.• <i>encryption-type</i>: Enter 0 for plain text or 7 for encrypted text.• <i>password</i>: Enter a string. To change only the password for the enable command, configure only the <i>password</i> parameter.

To view the configuration for the enable secret command, use the show running-config command in 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 command line interfaces (CLIs).

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, you can change the default level or you can reset their privilege level back to the default.

- Assign the launch keyword (for example, `configure`) for the keyword's command mode.

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 CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	<code>username <i>name</i> [access-class <i>access-list-name</i>] [privilege <i>level</i>] [nopassword password] [<i>encryption-type</i>] <i>password</i></code>	CONFIGURATION	Assign a user name and password. Configure the optional and required parameters: <ul style="list-style-type: none"> • <i>name</i>: Enter a text string (up to 63 characters). • <i>access-class access-list-name</i>: Enter the name of a configured IP ACL. • <i>privilege level</i>: range: 0 to 15. • <i>nopassword</i>: Do not require the user to enter a password. • <i>encryption-type</i>: Enter 0 for plain text or 7 for encrypted text. • <i>password</i>: Enter a string (up to 32 characters). The first character of the password must be a letter. You cannot use spaces in the password.
2	<code>enable password [level <i>level</i>] [<i>encryption-mode</i>] <i>password</i></code>	CONFIGURATION	Configure a password for privilege level. Configure the optional and required parameters: <ul style="list-style-type: none"> • <i>level level</i>: Specify a level 0 to 15. Level 15 includes all levels. • <i>encryption-type</i>: Enter 0 for plain text or 7 for encrypted text. • <i>password</i>: Enter a string (up to 25 characters). To change only the password for the <code>enable</code> command, configure only the <i>password</i> parameter.
3	<code>privilege <i>mode</i> {level <i>level</i> <i>command</i> reset <i>command</i>}</code>	CONFIGURATION	Configure level and commands for a mode or reset a command's level. Configure the following required and optional parameters: <ul style="list-style-type: none"> • <i>mode</i>: Enter a keyword for the modes (<code>exec</code>, <code>configure</code>, <code>interface</code>, <code>line</code>, <code>route-map</code>, <code>router</code>) • <i>level level</i>: range: 0 to 15. Levels 0, 1 and 15 are pre-configured. Levels 2 to 14 are available for custom configuration. • <i>command</i>: A FTOS CLI keyword (up to 5 keywords allowed). • <i>reset</i>: Return the command to its default privilege mode.

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

Figure 28-4 is an example of a configuration to allow a user “john” to view only EXEC mode commands and all snmp-server commands. Because the snmp-server commands are “enable” level commands and, by default, found in CONFIGURATION mode, you must also assign the launch command for CONFIGURATION mode, configure, to the same privilege level as the snmp-server commands.

Figure 28-4. Configuring a Custom Privilege Level

```
FTOS(conf)#username john privilege 8 password john
FTOS(conf)#enable password level 8 notjohn
FTOS(conf)#privilege exec level 8 configure
FTOS(conf)#privilege config level 8 snmp-server
FTOS(conf)#end
FTOS#show running-config
Current Configuration ...
!Version E8-3-16-0

hostname FTOS

!
enable password level 8 notjohn
enable password FTOS
!
username admin password 0 admin
```

← 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 because it is needed to reach CONFIGURATION mode where the snmp-server commands are located. The snmp-server commands, in CONFIGURATION mode, are assigned to privilege level 8.

Figure 28-5 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 EXEC Privilege mode, “john” can access only the commands listed. In CONFIGURATION mode, “john” can access only the snmp-server commands.

Figure 28-5. User john’s Login and the List of Available Commands

```
apollo% telnet 172.31.1.53
Trying 172.31.1.53...
Connected to 172.31.1.53.
Escape character is '^]'.
Login: john
Password:
FTOS#show priv
Current privilege level is 8
FTOS#?
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
FTOS#confi
FTOS(conf)#?
end              Exit from Configuration mode
```

Specify the LINE Mode Password and Privilege

You can specify a password authentication of all users on different *terminal* lines. The user's privilege level is the same as the privilege level assigned to the terminal line.

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

Command Syntax	Command Mode	Purpose
<code>privilege level <i>level</i></code>	LINE	Configure a custom privilege level for the terminal lines. <ul style="list-style-type: none"> <i>level level</i> range: 0 to 15. Levels 0, 1, and 15 are pre-configured. Levels 2 to 14 are available for custom configuration.
<code>password [<i>encryption-type</i>] <i>password</i></code>	LINE	Specify either a plain text or encrypted password. Configure the following optional and required parameters: <ul style="list-style-type: none"> <i>encryption-type</i>: Enter 0 for plain text or 7 for encrypted text. <i>password</i>: Enter a text string up to 25 characters long.

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

Enable and Disable Privilege Levels

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

To move to a lower privilege level, enter the disable command followed by the level-number you wish to set for the user in EXEC Privilege mode. If you enter the disable command 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 Dell Force10 system). The system 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, you can view the error 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 the user datagram protocol (UDP) as the transport protocol between the RADIUS server host and the client.

For more information about 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 you can specify it 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 you enable authorization, 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 the following 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 access control list (ACL). If an ACL is configured on the RADIUS server, and if that ACL is present, a 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 can configure the system through the RADIUS server to automatically execute a command when you connect to a specific line. To do this, use the auto-command command. The auto-command is executed when the user is authenticated and before the prompt appears to the user.

Set Access to Privilege Levels through RADIUS

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

Configuration Task List for RADIUS

To authenticate users using RADIUS, you must specify at least one RADIUS server so that the system can 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](#) (mandatory)
- [Apply the Method List to Terminal Lines](#) (mandatory except when using default lists)
- [Specify a RADIUS Server Host](#) (mandatory)
- [Set the Global Communication Parameters for all RADIUS Server Hosts](#) (optional)
- [Monitor RADIUS](#) (optional)

For a complete listing of all FTOS commands related to RADIUS, refer to the Security chapter in the *FTOS Command Reference Guide*.



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, 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 command in LINE mode or the show running-config command in EXEC Privilege mode.

Define an AAA Method List to be Used for RADIUS

To configure RADIUS to authenticate or authorize users on the system, you must create an AAA method list. Default method lists do not need to be explicitly applied to the line, so they are not mandatory.

To create a method list, enter one of the following commands in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
aaa authentication login <i>method-list-name</i> 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 { <i>method-list-name</i> 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 in CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
line {aux 0 console 0 vty <i>number</i> [<i>end-number</i>]}	CONFIGURATION	Enter the LINE mode.
login authentication { <i>method-list-name</i> 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 <i>methodlist</i>	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 CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
radius-server host { <i>hostname</i> <i>ip-address</i> } [auth-port <i>port-number</i>] [retransmit <i>retries</i>] [timeout <i>seconds</i>] [key [<i>encryption-type</i>] <i>key</i>]	CONFIGURATION	<p>Enter the host name or IP address of the RADIUS server host. Configure the optional communication parameters for the specific host:</p> <ul style="list-style-type: none"> auth-port <i>port-number</i> range: 0 to 65335. Enter a UDP port number. The default is 1812. retransmit <i>retries</i> range: 0 to 100. Default is 3. timeout <i>seconds</i> range: 0 to 1000. Default is 5 seconds. key [<i>encryption-type</i>] <i>key</i>: Enter 0 for plain text or 7 for encrypted text, and a string for the key. The key can be up to 42 characters long. This key must match the key configured on the RADIUS server host. <p>If you do not configure these optional parameters, the global default values for all RADIUS host are applied.</p>

To specify multiple RADIUS server hosts, configure the radius-server host command multiple times. If you configure multiple RADIUS server hosts, 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 the Global Communication Parameters for all RADIUS Server Hosts](#).

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

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

Set the 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 system. However, if you configure both global and specific host parameters, 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 CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
radius-server deadtime <i>seconds</i>	CONFIGURATION	<p>Set a time interval after which a RADIUS host server is declared dead.</p> <ul style="list-style-type: none"> <i>seconds</i> range: 0 to 2147483647. Default: 0 seconds

Command Syntax	Command Mode	Purpose
radius-server key [<i>encryption-type</i>] <i>key</i>	CONFIGURATION	Configure a key for all RADIUS communications between the system and RADIUS server hosts. <ul style="list-style-type: none"> <i>encryption-type</i>: Enter 7 to encrypt the password. Enter 0 to keep the password as plain text. <i>key</i>: Enter a string. The key can be up to 42 characters long. You cannot use spaces in the key.
radius-server retransmit <i>retries</i>	CONFIGURATION	Configure the number of times FTOS retransmits RADIUS requests. <ul style="list-style-type: none"> <i>retries</i> range: 0 to 100. Default is 3 retries.
radius-server timeout <i>seconds</i>	CONFIGURATION	Configure the time interval the system waits for a RADIUS server host response. <ul style="list-style-type: none"> <i>seconds</i> range: 0 to 1000. Default is 5 seconds.

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

Monitor RADIUS

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

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

TACACS+

FTOS supports the 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:

- [Choose TACACS+ as the Authentication Method](#)
- [Monitor TACACS+](#)
- [TACACS+ Remote Authentication and Authorization](#)
- [TACACS+ Remote Authentication and Authorization](#)
- [Specify a TACACS+ Server Host](#)

- [Choose TACACS+ as the Authentication Method](#)

For a complete listing of all commands related to TACACS+, refer to the Security chapter in the *FTOS Command Reference Guide*.

Choose TACACS+ as the 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 system 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 CONFIGURATION mode:

Step	Command Syntax	Command Mode	Purpose
1	<code>tacacs-server host {ip-address host}</code>	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	<code>aaa authentication login {method-list-name default} tacacs+ [...method3]</code>	CONFIGURATION	Enter a text string (up to 16 characters long) as the name of the method list you wish to use with the TACAS+ authentication method The <code>tacacs+</code> method should not be the last method specified.
3	<code>line {aux 0 console 0 vty number [end-number]}</code>	CONFIGURATION	Enter the LINE mode.
4	<code>login authentication {method-list-name default}</code>	LINE	Assign the <i>method-list</i> to the terminal line.

To view the configuration, use the `show config` command in LINE mode or the `show running-config tacacs+` command in 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 28-6](#), the TACACS+ is incorrect, but the user is still authenticated by the secondary method.

Figure 28-6. Failed Authentication

```

FTOS(conf)#
FTOS(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+
FTOS(conf)#
FTOS(conf)#do show run tacacs+
!
tacacs-server key 7 d05206c308f4d35b
tacacs-server host 10.10.10.10 timeout 1
FTOS(conf)#tacacs-server key angeline
FTOS(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)
FTOS(conf)#username angeline password angeline
FTOS(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 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 sizes. If you have configured remote authorization, 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, at least sees the login prompt. If the access class denies the connection, FTOS closes the Telnet session immediately.

Figure 28-7 shows how to configure access-class from a TACACS+ server. This causes the configured access-class on the VTY line to be ignored. If you have configured a deny10 ACL on the TACACS+ server, FTOS downloads it and applies it. If the user is found to be coming from the 10.0.0.0 subnet, FTOS also immediately closes the Telnet connection. Note that no matter where the user is coming from, they see the login prompt.

Figure 28-7. Specify a TACACS+ Server Host

```
FTOS#
FTOS(conf)#
FTOS(conf)#ip access-list standard deny10
FTOS(conf-std-nacl)#permit 10.0.0.0/8
FTOS(conf-std-nacl)#deny any
FTOS(conf)#
FTOS(conf)#aaa authentication login tacacsmethod tacacs+
FTOS(conf)#aaa authentication exec tacacsauthorization tacacs+
FTOS(conf)#tacacs-server host 25.1.1.2 key FTOS
FTOS(conf)#
FTOS(conf)#line vty 0 9
FTOS(conf-line-vty)#login authentication tacacsmethod
FTOS(conf-line-vty)#authorization exec tacacauthor
FTOS(conf-line-vty)#
FTOS(conf-line-vty)#access-class deny10
FTOS(conf-line-vty)#end
```

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

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

Command Syntax	Command Mode	Purpose
<code>tacacs-server host {hostname ip-address} [port port-number] [timeout seconds] [key key]</code>	CONFIGURATION	<p>Enter the host name or IP address of the TACACS+ server host. Configure the optional communication parameters for the specific host:</p> <ul style="list-style-type: none"> <code>port port-number</code> range: 0 to 65335. Enter a TCP port number. The default is 49. <code>timeout seconds</code> range: 0 to 1000. Default is 10 seconds. <code>key key</code>: Enter a string for the key. The key can be up to 42 characters long. This key must match a key configured on the TACACS+ server host. This parameter should be the last parameter configured. <p>If these optional parameters are not configured, the default global values are applied.</p>

To specify multiple TACACS+ server hosts, configure the `tacacs-server host` command multiple times. If you configure multiple TACACS+ server hosts, 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 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:
FTOS#
FTOS#
```

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 AAA authorization commands configure the system to check both the EXEC mode and CONFIGURATION mode commands. To enable only EXEC mode command checking, use the `no aaa authorization config-commands` command.

If rejected by the AAA server, the command is not added to the running config, and messages similar to [Message 1](#) are displayed.

Message 1 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 )
```

Protection from TCP Tiny and Overlapping Fragment Attacks

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 the ACL. RFC 1858 and 3128 proposes a countermeasure to the problem. This countermeasure is configured into the stack units and enabled by default.

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 information about command syntax, refer to the Security chapter in the *FTOS Command Line Interface Reference Guide*.

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



Note: The Windows-based WinSCP client software is not supported for secure copying between a PC and an FTOS-based system. Unix-based SCP client software is supported.

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

Command Syntax	Command Mode	Purpose
<code>ssh {hostname} [-l username -p port-number -v {1 2}]</code>	EXEC Privilege	Open an SSH connection specifying the hostname, username, port number, and version of the SSH client. <ul style="list-style-type: none"> <code>hostname</code> is the IP address or hostname of the remote device. Enter an IPv4 address in dotted decimal format (A.B.C.D).

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

Command Syntax	Command Mode	Purpose
<code>ip ssh server {enable port port-number}</code>	CONFIGURATION	Configure the Dell Force10 system as an SCP/SSH server.

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

Command Syntax	Command Mode	Purpose
<code>ip ssh server version {1 2}</code>	CONFIGURATION	Configure the Dell Force10 system as an SSH server that uses only version 1 or 2.

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

Command Syntax	Command Mode	Purpose
<code>show ip ssh</code>	EXEC Privilege	Display SSH connection information.

Figure 28-8 shows the use of the `ip ssh server version 2` command to enable SSH version 2, and the `show ip ssh` command to confirm the setting.

Figure 28-8. Specifying an SSH version

```
FTOS(conf)#ip ssh server version 2
FTOS(conf)#do show ip ssh
SSH server          : disabled.
SSH server version  : v2.
Password Authentication : enabled.
Hostbased Authentication : disabled.
RSA Authentication  : disabled.
```

To disable SSH server functions, use the `no ip ssh server enable` command.

Using SCP with SSH to Copy a Software Image

To use SCP to copy a software image through an SSH connection from one switch to another, follow these steps:

Step	Task	Command Syntax	Command Mode
1	On Chassis One, set the SSH port number (port 22 by default).	<code>ip ssh server port <i>number</i></code>	CONFIGURATION
2	On Chassis One, enable SSH.	<code>ip ssh server enable</code>	CONFIGURATION
3	On Chassis Two, invoke SCP.	<code>copy scp: flash:</code>	CONFIGURATION
4	On Chassis Two, in response to prompts, enter the path to the desired file and enter the port number specified in Step 1.		EXEC Privilege

Figure 28-9 shows the use of SCP and SSH to copy a software image from one switch running SSH Server on UDP port 99 to the local switch:

Figure 28-9. Using SCP to copy from an SSH Server on another Switch

```
FTOS#copy scp: flash:
Address or name of remote host []: 10.10.10.1
Port number of the server [22]: 99
Source file name []: test.cfg
User name to login remote host: admin
Password to login remote host:
```

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.

Secure Shell Authentication

SSH is disabled by default. Enable it using the `ip ssh server enable` command.

SSH supports three methods of authentication:

- [SSH Authentication by Password](#)
- [RSA Authentication of SSH](#)
- [Host-Based SSH Authentication](#)

Important Points to Remember for SSH Authentication

- If you enable more than one method, the order in which the methods are preferred is based on the `ssh_config` file on the Unix machine.
- When you enable all the three authentication methods, password authentication is the backup method when the RSA method fails.
- The `known_hosts` and `known_hosts2` files are generated when a user tries to SSH using version 1 or version 2, respectively.

SSH Authentication by Password

Authenticate an SSH client by prompting for a password when attempting to connect to the Dell Force10 system. This is the simplest methods of authentication and uses SSH version 1.

To enable SSH password authentication, use the `ip ssh password-authentication enable` command from CONFIGURATION mode. To view your SSH configuration, use the `show ip ssh` command from EXEC Privilege mode ([Figure 28-10](#)).

Figure 28-10. Enabling SSH Password Authentication

```
FTOS(conf)#ip ssh server enable
                % Please wait while SSH Daemon initializes ... done.
FTOS(conf)#ip ssh password-authentication enable
FTOS#sh ip ssh
SSH server      : enabled.
Password Authentication : enabled.
Hostbased Authentication : disabled.
RSA Authentication : disabled.
```

RSA Authentication of SSH

To authenticates an SSH client based on an RSA key using RSA authentication, follow these steps. This method uses SSH version 2:

Step	Task	Command Syntax	Command Mode
1	On the SSH client (Unix machine), generate an RSA key (Figure 28-11).		
	Figure 28-11. Generating RSA Keys		
	<pre>admin@Unix_client#ssh-keygen -t rsa Generating public/private rsa key pair. Enter file in which to save the key (/home/admin/.ssh/id_rsa): /home/admin/.ssh/id_rsa already exists. Overwrite (y/n)? y Enter passphrase (empty for no passphrase): Enter same passphrase again: Your identification has been saved in /home/admin/.ssh/id_rsa. Your public key has been saved in /home/admin/.ssh/id_rsa.pub.</pre>		
2	Copy the public key <i>id_rsa.pub</i> to the Dell Force10 system.		
3	Disable password authentication if enabled.	no ip ssh password-authentication enable	CONFIGURATION
4	Enable RSA authentication.	ip ssh rsa-authentication enable	EXEC Privilege
5	Bind the public keys to RSA authentication.	ip ssh rsa-authentication my-authorized-keys flash://public_key	EXEC Privilege

Host-Based SSH Authentication

Authenticate a particular host. This method uses SSH version 2.

To configure host-based authentication, use the following steps:

Step	Task	Command Syntax	Command Mode
1	Configure RSA Authentication. Refer to RSA Authentication of SSH above.		
2	Create shosts by copying the public RSA key to the file <i>shosts</i> in the directory <i>.ssh</i> , and write the IP address of the host to the file (Figure 28-12).	<code>cp /etc/ssh/ssh_host_rsa_key.pub /.ssh/shosts</code>	
<p>Figure 28-12. Creating shosts</p> <pre> admin@Unix_client# cd /etc/ssh admin@Unix_client# ls moduli sshd_config ssh_host_dsa_key.pub ssh_host_key.pub ssh_host_rsa_key.pub ssh_config ssh_host_dsa_key ssh_host_key ssh_host_rsa_key admin@Unix_client# cat ssh_host_rsa_key.pub ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEA8K7jLZRVfjgHJzUOmXxuIbZx/ AyWhVgJDQh39k8v3e8eQvLnHBIsqIL8jVy1QHhUeb7GaDlJVEDAMz30myqQbJgXBBRTWgBpLWwL/ doyUXFufjiL9YmoVTkbKcFmxJEMkE3JyHanEi7hg34LChjk9hL1by8cYZP2kYS2lnSyQWk= admin@Unix_client# ls id_rsa id_rsa.pub shosts admin@Unix_client# cat shosts 10.16.127.201, ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEA8K7jLZRVfjgHJzUOmXxuIbZx/AyW hVgJDQh39k8v3e8eQvLnHBIsqIL8jVy1QHhUeb7GaDlJVEDAMz30myqQbJgXBBRTWgBpLWwL/ </pre>			
3	Create a list of IP addresses and usernames that are permitted to SSH in a file called <i>rhosts</i> (Figure 28-13).		
<p>Figure 28-13. Creating rhosts</p> <pre> admin@Unix_client# ls id_rsa id_rsa.pub rhosts shosts admin@Unix_client# cat rhosts 10.16.127.201 admin </pre>			
4	Copy the file <i>shosts</i> and <i>rhosts</i> to the Dell Force10 system.		
5	Disable password authentication and RSA authentication, if configured	<ul style="list-style-type: none"> no ip ssh password-authentication no ip ssh rsa-authentication 	<ul style="list-style-type: none"> CONFIGURATION EXEC Privilege
6	Enable host-based authentication.	<code>ip ssh hostbased-authentication enable</code>	CONFIGURATION
7	Bind <i>shosts</i> and <i>rhosts</i> to host-based authentication.	<code>ip ssh pub-key-file flash://filename</code> <code>ip ssh rhostsfile flash://filename</code>	CONFIGURATION

Client-based SSH Authentication

To set SSH from the chassis to the SSH client, use the `ssh ip_address` command. This method uses SSH version 1 or version 2. If the SSH port is a non-default value, to change the default port number, use the `ip ssh server port number` command. You may only change the port number when SSH is disabled. You must then still use the `-p` option with the `ssh` command.

Figure 28-14. Client-Based SSH Authentication

```
FTOS#ssh 10.16.127.201 ?
-l           User name option
-p           SSH server port option (default 22)
-v           SSH protocol version
```

Troubleshooting SSH

- You may not bind `id_rsa.pub` to RSA authentication while logged in using the console. In this case, [Message 2](#) appears.

Message 2 RSA Authentication Error

```
%Error: No username set for this term.
```

- You must enable host-based authentication on the server (Dell Force10 system) and the client (Unix machine). [Message 3](#) appears if you attempt to log in using SSH and host-based is disabled on the client. In this case, verify that host-based authentication is set to “Yes” in the `ssh_config` file (root permission is required to edit this file).

Message 3 Host-Based Authentication Error

```
permission denied (host based)
```

- If the IP address in the RSA key does not match the IP address from which you attempt to log in, [Message 4](#) appears. In this case, verify that the name and IP address of the client is contained in the file `/etc/hosts`.

Message 4 RSA Authentication Error

```
getname info 8 failed
```

Telnet

To use Telnet with SSH, you must first enable SSH, as described above.

By default, the Telnet daemon is enabled. To disable the Telnet daemon, use the `[no] ip telnet server enable` command, or disable Telnet in the startup config (Figure 28-15).

Figure 28-15. [no] ip telnet server enable Command Example

```
FTOS(conf)#ip telnet server enable
FTOS(conf)#no ip telnet server enable
```

VTY Line and Access-Class Configuration

Various methods are available to restrict VTY access in FTOS. These depend on which authentication scheme you use—line, local, or remote (Table 28-1).

Table 28-1. VTY Access

Authentication Method	VTY Access-Class Support?	Username Access-Class Support?	Remote Authorization Support?
Line	YES	NO	NO
Local	NO	YES	NO
TACACS+	YES	NO	YES (with FTOS 5.2.1.0 and later)
RADIUS	YES	NO	YES (with FTOS 6.1.1.0 and later)

FTOS provides several ways to configure access classes for VTY lines, including:

- [VTY Line Local Authentication and Authorization](#)
- [VTY Line Remote Authentication and Authorization](#)

VTY Line Local Authentication and Authorization

FTOS retrieves the access class from the local database. To use this feature, follow these steps:

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

You can assign line authentication on a per-VTY basis; it is a simple password authentication using an access-class as authorization.

Local authentication is configured globally. You configure access classes on a per-user basis.

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



Note: If a VTY user logs in with RADIUS authentication, the privilege level is applied from the RADIUS server only if you configure RADIUS authentication.

Figure 28-16 shows how to allow or deny a Telnet connection to a user. Users see a login prompt, even if they cannot login. No access class is configured for the VTY line. It defaults from the local database.

Figure 28-16. Example Access-Class Configuration Using Local Database

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



Note: For more information, refer to [Access Control Lists \(ACLs\)](#).

VTY Line Remote Authentication and Authorization

FTOS 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. Figure 28-17 shows how to deny incoming connections from subnet 10.0.0.0 without displaying a login prompt. The example uses TACACS+ as the authentication mechanism.

Figure 28-17. Example Access Class Configuration Using TACACS+ Without Prompt

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

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 28-18](#)). [Figure 28-18](#) shows how to deny incoming connections from subnet 10.0.0.0 without displaying a login prompt.

Figure 28-18. Example Access Class Configuration Using TACACS+ Without Prompt

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


sFlow

This chapter contains the following sections:

- [Enable and Disable sFlow](#)
- [sFlow Show Commands](#)
- [Specify Collectors](#)
- [Polling Intervals](#)
- [Sampling Rate](#)
- [Back-Off Mechanism](#)
- [sFlow on LAG ports](#)
- [Extended sFlow](#)

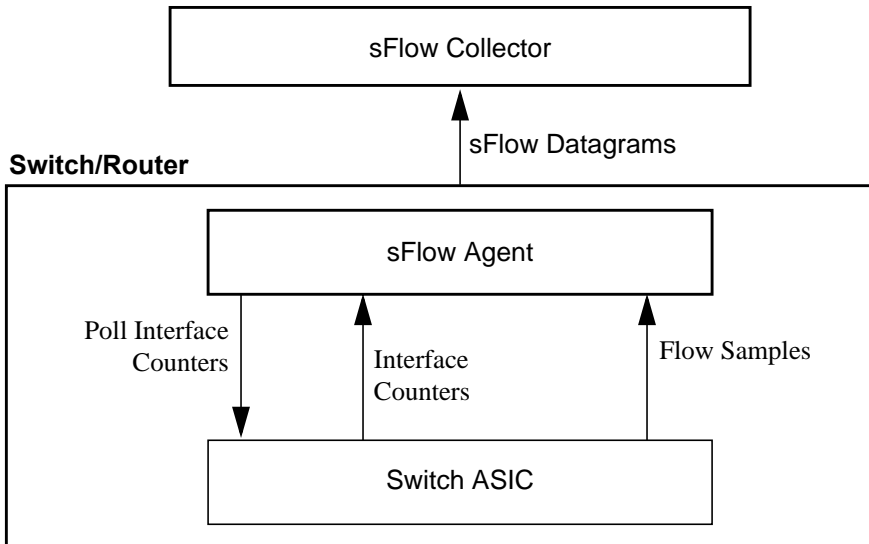
Overview

The Dell Force10 operating software (FTOS) supports sFlow version 5. sFlow is a standard-based sampling technology embedded within switches and routers which you can use to monitor network traffic (Figure 29-1). 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 application-specific integrated circuit (ASIC). sFlow collector analyses the sFlow datagrams received from different devices and produces a network-wide view of traffic flows.

Figure 29-1. sFlow Traffic Monitoring System

Implementation Information

The Dell Force10 sFlow is designed so that the hardware sampling rate is per stack unit port-pipe and is decided based on all the ports in that port-pipe. If you do not enable sFlow on any port specifically, the global sampling rate is downloaded to that port and is used to calculate the port-pipe's lowest sampling rate. This design supports the possibility that you might configure sFlow on that port in the future. Back-off is triggered based on the port-pipe's hardware sampling rate.

For example, if port 1 in the port-pipe has sFlow configured with a 16384 sampling rate while port 2 in the port-pipe has sFlow configured but no sampling rate set, FTOS applies a global sampling rate of 512 to port 2. The hardware sampling rate on the port-pipe is then set at 512 because that is the lowest configured rate on the port-pipe. When a high traffic situation occurs, a back-off is triggered and the hardware sampling rate is backed-off from 512 to 1024. Note that port 1 maintains its sampling rate of 16384; port 1 is unaffected because it maintains its configured sampling rate of 16384.

To avoid the back-off, either increase the global sampling rate or configure all the stack unit ports with the desired sampling rate even if some ports have no sFlow configured.

Important Points to Remember

- The FTOS implementation of the sFlow management information base (MIB) supports sFlow configuration using the `snmpset` command.
- FTOS exports all sFlow packets to the collector. A small sampling rate can equate to a large number of exported packets. A backoff mechanism automatically applies to reduce this amount. Some sampled packets may be dropped when the exported packet rate is high and the backoff mechanism is about to or is starting to take effect. The `dropEvent` counter, in the sFlow packet, is always zero.
- The community list and local preference fields are not filled in the extended gateway element in the sFlow datagram.

- The 802.1P source priority field is not filled in extended switch element in the sFlow datagram.
- Only the Destination and Destination Peer AS number are packed in the dst-as-path field in extended gateway element.
- If the packet being sampled is redirected using policy-based routing (PBR), the sFlow datagram may contain incorrect extended gateway and/or router information.
- The source VLAN field in the extended switch element is not packed in case of routed packet.
- The destination VLAN field in the extended switch element is not packed in case of multicast packet.
- On the MXL Switch, up to 700 packets can be sampled and processed per second.

Enable and Disable sFlow

By default, sFlow is disabled globally on the system. To enable sFlow globally, use the `sflow enable` command in CONFIGURATION mode. To disable sFlow globally, use the `no sflow enable` command.

Command Syntax	Command Mode	Usage
[no] sflow enable	CONFIGURATION	Enable sFlow globally.

Enable and Disable 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. This command line interface (CLI) is supported on physical ports and link aggregation group (LAG) ports. To disable sFlow on an interface, use the `no sflow enable` command.

Command Syntax	Command Mode	Usage
[no] sflow enable	INTERFACE	Enable sFlow on an interface.

sFlow Show Commands

FTOS includes the following sFlow display commands:

- [Show sFlow Globally](#)
- [Show sFlow on an Interface](#)
- [Show sFlow on a Stack Unit](#)

Show sFlow Globally

To view sFlow statistics, use the following command (Figure 29-2):

Command Syntax	Command Mode	Purpose
show sflow	EXEC	Display sFlow configuration information and statistics.

Figure 29-2. show sflow Command Example

```
FTOS#show sflow
sFlow services are enabled ← Indicates sFlow is globally 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
FTOS#
```

Show sFlow on an Interface

To view sFlow information on a specific interface, use the following command (Figure 29-3):

Command Syntax	Command Mode	Purpose
show sflow interface <i>interface-name</i>	EXEC	Display sFlow configuration information and statistics on a specific interface.

Figure 29-3. show sflow interface Command Example

```
FTOS#show sflow interface tengigabitethernet 1/16
Tengig 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
```

Show sFlow on a Stack Unit

To view sFlow statistics on a specified stack unit, use the following command (Figure 29-4):

Command Syntax	Command Mode	Purpose
<code>show sflow stack-unit <i>unit-number</i></code>	EXEC	Display sFlow configuration information and statistics on the specified interface.

Figure 29-4. show sflow stack unit Command Example

```
FTOS#show sflow stack-unit 1
Stack-Unit 1
  Samples rcvd from h/w           :0
  Total UDP packets exported      :0
  UDP packets exported via RPM    :0
  UDP packets dropped             :0
FTOS#
```

Specify Collectors

The `sflow collector` command allows identification of sFlow collectors to which sFlow datagrams are forwarded. You can specify up to two sFlow collectors. If you specify two collectors, the samples are sent to both.

To identify sFlow collectors, use the following command:

Command Syntax	Command Mode	Usage
<code>sflow collector <i>ip-address</i> agent-addr <i>ip-address</i> [<i>number</i> [<i>max-datagram-size number</i>]] [<i>max-datagram-size number</i>]</code>	CONFIGURATION	Identify 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.

You can configure the polling interval globally (in CONFIGURATION mode) or by interface (in INTERFACE mode) by executing the following interval command:

Command Syntax	Command Mode	Usage
<code>sflow polling-interval <i>interval</i></code> <i>value</i>	CONFIGURATION or INTERFACE	Change 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 numbers and re-enter the command.

For more information about values in power-of-2, refer to [Sub-Sampling](#).

You can configure the sample rate globally or by interface using the following sample rate command:

Command Syntax	Command Mode	Usage
<code>[no] sflow sample-rate</code> <i>sample-rate</i>	CONFIGURATION or INTERFACE	Change the global or interface sampling rate. Rate must be entered in factors of 2 (for example, 4096, 8192). <i>sample-rate</i> range: 256 to 8388608

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.

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 the 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 Tengig 1/0 and 1/1 are in a port-pipe, and they are configured with a sampling rate of 4096 on interface Tengig 1/0, and 8192 on Tengig 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. Configures interface Tengig 1/0 to a sub-sampling rate of 1 to achieve an actual rate of 4096.

3. Configures interface Tengig 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. The following equation shows the relationship between the actual sampling rate, the sub-sampling rate, and the hardware sampling rate for an interface:

$$\text{Actual sampling rate} = \text{sub-sampling rate} * \text{hardware sampling rate}$$

Note: There is an 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 overloaded with flow samples under high-traffic conditions. In such a scenario, a binary back-off mechanism is triggered, which doubles the sampling-rate (halves the number of samples per second) for all interfaces. The backoff mechanism continues to double the sampling-rate until CPU condition is cleared. This is as per sFlow version 5 draft. After the back-off changes the sample-rate, you 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. To view the actual sampling-rate of the interface and the configured sample-rate, use 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

The MXL switch supports extended-switch information processing *only*.

Extended sFlow packs additional information in the sFlow datagram depending on the type of sampled packet. You can enable the following options:

- `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 border gateway protocol (BGP) next-hop.

To enable extended sFlow, use the `sflow [extended-switch] [extended-router] [extended-gateway] enable` command. By default, packing of any of the extended information in the datagram is disabled.

To confirm that extended information packing is enabled, use the show sflow to confirm that extended information packing is enabled (Figure 29-5).

Figure 29-5. Confirming that Extended sFlow is Enabled

```
FTOS#show sflow
sFlow services are enabled
Global default sampling rate: 4096
Global default counter polling interval: 15
Global extended information enabled: switch
1 collectors configured
Collector IP addr: 10.10.10.3, Agent IP addr: 10.10.0.0, UDP port: 6343
77 UDP packets exported
0 UDP packets dropped
165 sFlow samples collected
69 sFlow samples dropped due to sub-sampling
Stackunit 1 Port set 0 H/W sampling rate 8192
Tengig 1/16: configured rate 8192, actual rate 8192, sub-sampling rate 1
Tengig 1/17: configured rate 16384, actual rate 16384, sub-sampling rate 2
Stackunit 3 Port set 1 H/W sampling rate 16384
Tengig 3/40: configured rate 16384, actual rate 16384, sub-sampling rate 1
```

Extended sFlow settings show all 3 types are enabled





Figure 29-6 shows that none of the extended information is enabled.

Figure 29-6. Confirming that Extended sFlow is Disabled

```
FTOS#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
```

Indicates no Extended sFlow types enabled.



Simple Network Management Protocol (SNMP)

Protocol Overview

Network management stations use the Simple Network Management Protocol (SNMP) to retrieve or alter management data from network elements. A datum of management information is called a *managed object*; the value of a managed object can be static or variable. Network elements store managed objects in a database called a *Management Information Base* (MIB).

MIBs are hierarchically structured and use object identifiers to address managed objects, but managed objects also have a textual name called an *object descriptor*.



Note: On Dell Force10 routers, standard and private SNMP MIBs are supported, including all Get and a limited number of Set operations (such as **set vlan** and **copy cmd**).

Implementation Information

- FTOS supports SNMP version 1 as defined by RFC 1155, 1157, and 1212, SNMP version 2c as defined by RFC 1901, and SNMP version 3 as defined by RFC 2571.
- FTOS supports up to 16 trap receivers.
- The FTOS implementation of the sFlow MIB supports sFlow configuration via SNMP sets.
- SNMP traps for STP and MSTP state changes are based on BRIDGE MIB (RFC 1483) for STP and IEEE 802.1 *draft ruzin-mstp-mib-02* for MSTP.
- All objects in the LLDP-EXT-DOT1-DCBX-MIB and IEEE8021-PFC-MIB tables are read-only.
- In the LLDP-EXT-DOT1-DCBX-MIB, lldpXdot1dcbxRemoteData tables are not supported.

Configure Simple Network Management Protocol



Note: The configurations in this chapter use a UNIX environment with net-snmp version 5.4. This is only one of many RFC-compliant SNMP utilities you can use to manage your Dell Force10 system using SNMP. Also, these configurations use SNMP version 2c.

Configuring SNMP version 1 or version 2 requires only a single step:

1. Create a community. See [page 517](#).

Configuring SNMP version 3 requires you to configure SNMP users in one of three methods. See [Setting Up User-based Security \(SNMPv3\)](#).

Related Configuration Tasks

The following list contains configuration tasks for SNMP:

- [Setting up SNMP](#)
- [Setting Up User-based Security \(SNMPv3\)](#)
- [Read Managed Object Values](#)
- [Write Managed Object Values](#)
- [Configure Contact and Location Information Using SNMP](#)
- [Subscribe to Managed Object Value Updates using SNMP](#)
- [Copy Configuration Files Using SNMP](#)
- [Manage VLANs Using SNMP](#)
- [Enable and Disable a Port Using SNMP](#)
- [Fetch Dynamic MAC Entries Using SNMP](#)
- [Deriving Interface Indices](#)
- [Monitor Port-channels](#)
- [Troubleshooting SNMP Operations](#)

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, increase the timeout value to greater than 3 seconds, and increase the retry value to greater than 2 on your SNMP server.
- User ACLs override group ACLs.

Setting up SNMP

As previously stated, FTOS supports SNMP version 1 and version 2 which are community-based security models. The primary difference between the two versions is that version 2 supports two additional protocol operations (informs operation and **snmpgetbulk** query) and one additional object (counter64 object).

SNMP version 3 (SNMPv3) is a user-based security model that provides password authentication for user security and encryption for data security and privacy. Three sets of configurations are available for SNMP read/write operations: no password or privacy, password privileges, password and privacy privileges

A maximum of 16 users can be configured even if they are in different groups.

Create a Community

For SNMPv1 and SNMPv2, you must create a community to enable the community-based security in FTOS. The management station generates requests to either retrieve or alter the value of a management object and is called the *SNMP manager*. A network element that processes SNMP requests is called an *SNMP agent*. An *SNMP community* is a group of SNMP agents and managers that are allowed to interact. Communities are necessary to secure communication between SNMP managers and agents; SNMP agents do not respond to requests from management stations that are not part of the community.

FTOS enables SNMP automatically when you create an SNMP community and displays [Message 1](#). You must specify whether members of the community may only retrieve values (read), or retrieve and alter values (read-write).

To create an SNMP community:

Task	Command	Command Mode
Choose a name for the community.	<code>snmp-server community name {ro rw}</code>	CONFIGURATION

Message 1 SNMP Enabled

```
22:31:23: %STKUNIT0-M:CP %SNMP-6-SNMP_WARM_START: Agent Initialized - SNMP WARM_START.
```

View your SNMP configuration, using the command `show running-config snmp` from EXEC Privilege mode, as shown in [Figure 30-1](#).

Figure 30-1. Creating an SNMP Community

```
FTOS(conf)#snmp-server community my-snmp-community ro
22:31:23: %STKUNIT0-M:CP %SNMP-6-SNMP_WARM_START: Agent Initialized - SNMP WARM_START.
FTOS#do show running-config snmp
!
snmp-server community mycommunity ro
FTOS#
```

Setting Up User-based Security (SNMPv3)

When setting up SNMPv3, you can set users up with one of the following three types of configuration for SNMP read/write operations. Users are typically associated to an SNMP group with permissions provided, such as OID view.

- **noauth:** no password or privacy. Select this option to set a user up with no password or privacy privileges. This is the basic configuration. Users must have a group and profile that do not require password privileges.
- **auth:** password privileges. Select this option to set up an user with password authentication
- **priv:** password and privacy privileges. Select this option to set up a user with password and privacy privileges.

Figure 30-2. Select a User-based Security Type

```

FTOS(conf)#snmp-server host 1.1.1.1 traps {oid tree} version 3 ?
auth                Use the SNMPv3 authNoPriv Security Level
noauth              Use the SNMPv3 noAuthNoPriv Security Level
priv                Use the SNMPv3 authPriv Security Level
FTOS(conf)#snmp-server host 1.1.1.1 traps version 3 noauth ?
WORD                SNMPv3 user name
  
```

To set up a user with view privileges only (no password or privacy privileges):

Task	Command	Command Mode
Configure the user.	snmp-server user <i>name group-name 3 noauth</i>	CONFIGURATION
Configure an SNMP group.	snmp-server group <i>group-name 3 noauth auth read name write name</i>	CONFIGURATION
Configure an SNMPv3 view.	snmp-server view <i>view-name oid-tree {included excluded}</i> Note: To give a user read and write view privileges, repeat this step for each privilege type.	CONFIGURATION

To set up a user with password privileges only, use the following commands:

Task	Command	Command Mode
Configure the user with an authorization password	snmp-server user <i>name group-name 3 noauth auth md5 auth-password</i>	CONFIGURATION
Configure an SNMP group.	snmp-server group <i>groupname {oid-tree} auth read name write name</i>	CONFIGURATION
Configure an SNMPv3 view.	snmp-server view <i>view-name 3 noauth {included excluded}</i> Note: To give a user read and write privileges, repeat this step for each privilege type.	CONFIGURATION

To set up a user with password or privacy privileges:

Task	Command	Command Mode
Configure an SNMP group.	snmp-server group <i>group-name {oid-tree} priv read name write name</i>	CONFIGURATION
Configure the user with a secure authorization password and privacy password.	snmp-server user <i>name group-name {oid-tree} auth md5 auth-password priv des56 priv password</i>	CONFIGURATION
Configure an SNMPv3 view.	snmp-server view <i>view-name oid-tree {included excluded}</i>	CONFIGURATION

Read Managed Object Values

You may only retrieve (read) managed object values if your management station is a member of the same community as the SNMP agent.

Dell Force10 supports RFC 4001, *Textual Conventions for Internet Work Addresses* that defines values representing a type of internet address. These values display for ipAddressTable objects using the **snmpwalk** command.

In the following figure, the value “4” displays in the OID before the IP address for IPv4.

```
>snmpwalk -v 2c -c public 10.11.195.63 1.3.6.1.2.1.4.34
IP-MIB::ip.34.1.3.1.4.1.1.1.1 = INTEGER: 1107787778
IP-MIB::ip.34.1.3.1.4.2.1.1.1 = INTEGER: 1107787779
IP-MIB::ip.34.1.3.2.16.254.128.0.0.0.0.0.2.1.232.255.254.139.5.8 = INTEGER: 1107787778
IP-MIB::ip.34.1.4.1.4.1.1.1.1 = INTEGER: 1
IP-MIB::ip.34.1.4.1.4.2.1.1.1 = INTEGER: 1
IP-MIB::ip.34.1.4.2.16.0.1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1 = INTEGER: 1
```

There are several UNIX SNMP commands that read data:

Task	Command
Read the value of a single managed object as shown in Figure 30-3 .	snmpget -v version -c community agent-ip {identifier.instance descriptor.instance}
Figure 30-3. Reading the Value of a Managed Object	
<pre>> snmpget -v 2c -c mycommunity 10.11.131.161 sysUpTime.0 DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (32852616) 3 days, 19:15:26.16 > snmpget -v 2c -c mycommunity 10.11.131.161 .1.3.6.1.2.1.1.3.0 DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (32856932) 3 days, 19:16:09.32</pre>	
Read the value of the managed object directly below the specified object as shown in Figure 30-4 .	snmpgetnext -v version -c community agent-ip {identifier.instance descriptor.instance}
Figure 30-4. Reading the Value of the Next Managed Object in the MIB	
<pre>> snmpgetnext -v 2c -c mycommunity 10.11.131.161 .1.3.6.1.2.1.1.3.0 SNMPv2-MIB::sysContact.0 = STRING: > snmpgetnext -v 2c -c mycommunity 10.11.131.161 sysContact.0 SNMPv2-MIB::sysName.0 = STRING:</pre>	
Read the value of many objects at once as shown in Figure 30-5 .	snmpwalk -v version -c community agent-ip {identifier.instance descriptor.instance}

Task	Command
------	---------

Figure 30-5. Reading the Value of Many Managed Objects at Once

```

> snmpwalk -v 2c -c mycommunity 10.11.209.217 .1.3.6.1.2.1.1
SNMPv2-MIB::sysDescr.0 = STRING: Dell Force10 OS
Operating System Version: 1.0
Application Software Version: E8-3-16-0
Series: MXL-10/40GbE
Copyright (c) 1999-2012 by Dell Inc. All Rights Reserved.
Build Time: Tue May 22 22:40:56 PDT 2012
SNMPv2-MIB::sysObjectID.0 = OID: SNMPv2-SMI::enterprises.6027.1.4.1
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (256676) 0:42:46.76
SNMPv2-MIB::sysContact.0 = STRING:
SNMPv2-MIB::sysName.0 = STRING: FTOS
SNMPv2-MIB::sysLocation.0 = STRING:
SNMPv2-MIB::sysServices.0 = INTEGER: 4

```

Write Managed Object Values

You may only alter (write) a managed object value if your management station is a member of the same community as the SNMP agent, and the object is writable.

To write or write-over the value of a managed object:

Task	Command
To write or write-over the value of a managed object, as shown in Figure 30-6 .	snmpset -v version -c community agent-ip {identifier.instance descriptor.instance} syntax value

Figure 30-6. Writing over the Current Value of a Managed Object

```

> snmpset -v 2c -c mycommunity 10.11.131.161 sysName.0 s "R5"
SNMPv2-MIB::sysName.0 = STRING: R5

```

Configure Contact and Location Information Using SNMP

You may configure system contact and location information from the Dell Force10 system or from the management station using SNMP.

To configure system contact and location information from the Dell Force10 system:

Task	Command	Command Mode
Identify the system manager along with this person's contact information (e.g., E-mail address or phone number). You may use up to 55 characters. Default: None	snmp-server contact <i>text</i>	CONFIGURATION
Identify the physical location of the system. For example, San Jose, 350 Holger Way, 1st floor lab, rack A1-1. You may use up to 55 characters. Default: None	snmp-server location <i>text</i>	CONFIGURATION

To configure the system from the management station using SNMP:

Task	Command	Command Mode
Identify the system manager along with this person's contact information (e.g., E-mail address or phone number). You may use up to 55 characters. Default: None	snmpset -v version -c community agent-ip sysContact.0 s "contact-info"	CONFIGURATION
Identify the physical location of the system. For example, San Jose, 350 Holger Way, 1st floor lab, rack A1-1. You may use up to 55 characters. Default: None	snmpset -v version -c community agent-ip sysLocation.0 s "location-info"	CONFIGURATION

Subscribe to Managed Object Value Updates using SNMP

By default, the Dell Force10 system displays some unsolicited SNMP messages (traps) upon certain events and conditions. You can also configure the system to send the traps to a management station. Traps cannot be saved on the system.

FTOS supports the following three sets of traps:

- **RFC 1157-defined traps:** coldStart, warmStart, linkDown, linkUp, authenticationFailure, egpNeighborLoss
- **Dell Force10 enterpriseSpecific environment traps:** temperature
- **Dell Force10 enterpriseSpecific protocol traps:** ecfm, entity, ets, fips, lacp, pfc, stp, xstp

To configure the system to send SNMP notifications, follow these steps:

Step	Task	Command	Command Mode
1	Configure the Dell Force10 system to send notifications to an SNMP server. <ul style="list-style-type: none"> Enter the keyword traps to send trap messages. Enter the keyword informs to send informational messages. Enter the keyword version to send the SNMP version to use for notification messages. Enter the name of the <i>community-string</i> to identify the SNMPv1 community string. 	snmp-server host <i>ip-address</i> [traps informs] [version 1 2c 3] [<i>community-string</i>]	CONFIGURATION
2	Specify which traps the Dell Force10 system sends to the trap receiver. <ul style="list-style-type: none"> Enable all Dell Force10 enterpriseSpecific and RFC-defined traps using the command snmp-server enable traps from CONFIGURATION mode. Enable all of the RFC-defined traps using the command snmp-server enable traps snmp from CONFIGURATION mode. 	snmp-server enable traps	CONFIGURATION
3	Specify the interfaces out of which FTOS sends SNMP traps.	snmp-server trap-source	CONFIGURATION

Table 30-1 lists the traps the RFC-defined SNMP traps and the command used to enable each. Note that the coldStart and warmStart traps are enabled using a single command.

Table 30-1. RFC 1157 Defined SNMP Traps on FTOS

Command Option	Trap
snmp authentication	SNMP_AUTH_FAIL:SNMP Authentication failed.Request with invalid community string.
snmp coldstart	SNMP_COLD_START: Agent Initialized - SNMP COLD_START. SNMP_WARM_START: Agent Initialized - SNMP WARM_START.
snmp linkdown	PORT_LINKDN:changed interface state to down:%d
snmp linkup	PORT_LINKUP:changed interface state to up:%d

Enable a subset of Dell Force10 enterprise specific SNMP traps using one of the listed command options Table 30-2 with the command **snmp-server enable traps**. Note that the **envmon** option enables all environment traps including those that are enabled with the **envmon temperature** option.

Table 30-2. Dell Force10 Enterprise-specific SNMP Traps

Command Option	Trap
envmon temperature	<p>MINOR_TEMP: Minor alarm: chassis temperature</p> <p>MINOR_TEMP_CLR: Minor alarm cleared: chassis temperature normal (%s %d temperature is within threshold of %dC)</p> <p>MAJOR_TEMP: Major alarm: chassis temperature high (%s temperature reaches or exceeds threshold of %dC)</p> <p>MAJOR_TEMP_CLR: Major alarm cleared: chassis temperature lower (%s %d temperature is within threshold of %dC)</p>
xstp	<p>%SPANMGR-5-STP_NEW_ROOT: New Spanning Tree Root, Bridge ID Priority 32768, Address 0001.e801.fc35.</p> <p>%SPANMGR-5-STP_TOPOLOGY_CHANGE: Bridge port GigabitEthernet 11/38 transitioned from Forwarding to Blocking state.</p> <p>%SPANMGR-5-MSTP_NEW_ROOT_BRIDGE: Elected root bridge for instance 0.</p>
entity	<p>Enable entity change traps</p> <p>Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (1487406) 4:07:54.06, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::mib-2.47.2.0.1, SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 4</p> <p>Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (1488564) 4:08:05.64, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::mib-2.47.2.0.1, SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 5</p> <p>Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (1489064) 4:08:10.64, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::mib-2.47.2.0.1, SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 6</p> <p>Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (1489568) 4:08:15.68, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::mib-2.47.2.0.1, SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 7</p>
<cr>	<p>SNMP Copy Config Command Completed</p> <p>%RPM0-P:CP %SNMP-4-RMON_RISING_THRESHOLD: RMON rising threshold alarm from SNMP OID <oid></p> <p>%RPM0-P:CP %SNMP-4-RMON_FALLING_THRESHOLD: RMON falling threshold alarm from SNMP OID <oid></p> <p>%RPM0-P:CP %SNMP-4-RMON_HC_RISING_THRESHOLD: RMON high-capacity rising threshold alarm from SNMP OID <oid></p>
coldstart	<p>10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (6796) 0:01:07.96, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-MIB::coldStart, SNMPv2-SMI::enterprises.6027.3.1.1.4.1.1 = INTEGER: 6, SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "SNMP_COLD_START: Agent Initialized - SNMP COLD_START.", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 1</p>
warmstart linkdown linkup	<p>10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (6756) 0:01:07.56, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-MIB::warmStart, SNMPv2-SMI::enterprises.6027.3.1.1.4.1.1 = INTEGER: 6, SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "SNMP_WARM_START: Agent Initialized - SNMP WARM_START.", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 1</p> <p>10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (625882) 1:44:18.82, SNMPv2-MIB::snmpTrapOID.0 = OID: IF-MIB::linkUp, IF-MIB::ifIndex.45158657 = INTEGER: 45158657, SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_UP: Changed interface state to up: Te 0/43", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 14</p>

Table 30-2. Dell Force10 Enterprise-specific SNMP Traps

Command Option	Trap
ets	<pre>10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (645746) 1:47:37.46, SNMPv2-MIB::snmpTrapOID.0 = OID: IF-MIB::linkDown, IF-MIB::ifIndex.45420801 = INTEGER: 45420801, SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_DN: Changed interface state to down: Te 0/44", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 22</pre>
	<pre>ETS peer state enabled</pre>
	<pre>10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (625916) 1:44:19.16, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.6027.3.15.4.0.3, SNMPv2-SMI::enterprises.6027.3.15.4.1.1.0 = INTEGER: 45158657, SNMPv2-SMI::enterprises.6027.3.15.4.1.2.0 = INTEGER: 1, SNMPv2-SMI::enterprises.6027.3.15.4.0 = STRING: "ETS_TRAP_TYPE_PEER_STATE_CHANGE: ETS Peer state changed to enabled for port Te 0/43", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 17</pre>
	<pre>ETS peer state disabled</pre>
	<pre>10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (645772) 1:47:37.72, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.6027.3.15.4.0.3, SNMPv2-SMI::enterprises.6027.3.15.4.1.1.0 = INTEGER: 45420801, SNMPv2-SMI::enterprises.6027.3.15.4.1.2.0 = INTEGER: 2, SNMPv2-SMI::enterprises.6027.3.15.4.0 = STRING: "ETS_TRAP_TYPE_PEER_STATE_CHANGE: ETS Peer state changed to disabled for port Te 0/44", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 23</pre>
pfc	<pre>pfc peer state enabled</pre>
	<pre>10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (626100) 1:44:21.00, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.6027.3.15.4.0.7, SNMPv2-SMI::enterprises.6027.3.15.4.1.1.0 = INTEGER: 45420801, SNMPv2-SMI::enterprises.6027.3.15.4.1.2.0 = INTEGER: 1, SNMPv2-SMI::enterprises.6027.3.15.4.0 = STRING: "PFC_TRAP_TYPE_PEER_STATE_CHANGE: PFC Peer state changed to enabled for port Te 0/44", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 21</pre>
	<pre>pfc peer state disbled</pre>
	<pre>10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (645794) 1:47:37.94, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.6027.3.15.4.0.7, SNMPv2-SMI::enterprises.6027.3.15.4.1.1.0 = INTEGER: 45420801, SNMPv2-SMI::enterprises.6027.3.15.4.1.2.0 = INTEGER: 2, SNMPv2-SMI::enterprises.6027.3.15.4.0 = STRING: "PFC_TRAP_TYPE_PEER_STATE_CHANGE: PFC Peer state changed to disabled for port Te 0/44", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 24</pre>

Copy Configuration Files Using SNMP

Use SNMP from a remote client to:

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

- copy configuration files from a server to the Dell Force10 system

You can perform all of these tasks using IPv4 addresses.

The relevant MIBs for these functions are:


Table 30-3. MIB Objects for Copying Configuration Files Using SNMP

MIB Object	OID	Object Values	Description
copySrcFileType	.1.3.6.1.4.1.6027.3.5.1.1.1.1.2	1 = FTOS file 2 = running-config 3 = startup-config	Specifies the type of file to copy from. Valid values are: <ul style="list-style-type: none"> • 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	.1.3.6.1.4.1.6027.3.5.1.1.1.1.3	1 = flash 2 = n/a 3 = tftp 4 = ftp 5 = scp 6 = usbflash	Specifies the location of source file. <ul style="list-style-type: none"> • If the copySrcFileLocation is FTP or SCP, copyServerAddress, copyUserName, and copyUserPassword must be specified.
copySrcFileName	.1.3.6.1.4.1.6027.3.5.1.1.1.1.4	Path (if file is not in current directory) and filename.	Specifies name of the file. <ul style="list-style-type: none"> • If copySourceFileType is set to running-config or startup-config, copySrcFileName is not required.
copyDestFileType	.1.3.6.1.4.1.6027.3.5.1.1.1.1.5	1 = FTOS file 2 = running-config 3 = startup-config	Specifies the type of file to copy to. <ul style="list-style-type: none"> • 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	.1.3.6.1.4.1.6027.3.5.1.1.1.1.6	1 = flash 2 = n/a 3 = tftp 4 = ftp 5 = scp	Specifies the location of destination file. <ul style="list-style-type: none"> • If the copyDestFileLocation is FTP or SCP, copyServerAddress, copyUserName, and copyUserPassword must be specified.
copyDestFileName	.1.3.6.1.4.1.6027.3.5.1.1.1.1.7	Path (if file is not in default directory) and filename.	Specifies the name of destination file.
copyServerAddress	.1.3.6.1.4.1.6027.3.5.1.1.1.1.8	IP Address of the server	The IP address of the server. <ul style="list-style-type: none"> • If the copyServerAddress is specified so must copyUserName, and copyUserPassword.

Table 30-3. MIB Objects for Copying Configuration Files Using 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 the FTP, TFTP, or SCP server. <ul style="list-style-type: none"> 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 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 <i>community-name rw</i>	CONFIGURATION
2	Copy the <i>f10-copy-config.mib</i> MIB from the Dell 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...		
	<ul style="list-style-type: none"> Every specified object must have an object value, which must be preceded by the keyword <i>i</i>. See Table 30-3 for valid values. <i>index</i> must be unique to all previously executed snmpset commands. If an index value has been used previously, a message like the one in Message 2 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 30-3 to complete the command. See Table 30-4 for examples. 		
	Note: You can use the entire OID rather than the object name. Use the form: <i>OID.index i object-value</i> as shown in Figure 30-8 .		

Message 2 snmpset Index Value Error

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

[Table 30-4](#) 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* in the directory from which you are executing the **snmpset** command or in the **snmpset** tool path.



Note: Use the following options in the **snmpset** command to view additional information:

- c: View the community, either public or private
- m: View the MIB files for the SNMP command
- r: Number of retries using the option
- t: View the timeout
- v: View the SNMP version (either 1, 2, 2d, or 3)

Table 30-4. Copying Configuration Files via SNMP

Task

Copy the running-config to the startup-config using the following command from the UNIX machine:

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

Figure 30-7 shows the command syntax using MIB object names. Figure 30-8 shows the same command using the object OIDs. In both cases, the object is followed by a unique index number.

Figure 30-7. Copying Configuration Files via SNMP using Object-Name Syntax

```
> snmpset -v 2c -r 0 -t 60 -c public -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 30-8. Copying Configuration Files via SNMP using OID Syntax

```
> snmpset -v 2c -c public -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)
```

Table 30-4. Copying Configuration Files via SNMP

Task

Copy the startup-config to the running-config using the following command from a UNIX machine:

```
snmpset -c private -v 2c force10system-ip-address copySrcFileType.index i 3 copyDestFileType.index i 2
```

Figure 30-9. Copying Configuration Files via SNMP using Object-Name Syntax

```
> snmpset -c public -v 2c -m ./f10-copy-config.mib 10.11.131.162 copySrcFileType.7 i 3
copyDestFileType.7 i 2
FORCE10-COPY-CONFIG-MIB::copySrcFileType.7 = INTEGER: runningConfig(3)
FORCE10-COPY-CONFIG-MIB::copyDestFileType.7 = INTEGER: startupConfig(2)
```

Figure 30-10. Copying Configuration Files via SNMP using OID Syntax

```
>snmpset -c public -v 2c 10.11.131.162 .1.3.6.1.4.1.6027.3.5.1.1.1.1.2.8 i 3
.1.3.6.1.4.1.6027.3.5.1.1.1.1.5.8 i 2
SNMPv2-SMI::enterprises.6027.3.5.1.1.1.1.2.8 = INTEGER: 3
SNMPv2-SMI::enterprises.6027.3.5.1.1.1.1.5.8 = INTEGER: 2
```

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

```
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 30-11. 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 using TFTP using the following command from the UNIX machine:

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

```
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
```

Table 30-4. Copying Configuration Files via SNMP**Task****Figure 30-12. 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 Dell Force10 system using FTP using the following command:

```
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 30-13. 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
```

Dell Force10 provides additional MIB objects to view copy statistics. These are provided in [Table 30-5](#)

Table 30-5. 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.
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 30-5](#), follow this step:

Step	Task
1	<p>Get a copy-config MIB object value.</p> <pre>snmpset -v 2c -c public -m /f10-copy-config.mib force10system-ip-address [OID.index mib-object.index]</pre> <ul style="list-style-type: none">index is the index value used in the snmpset command used to complete the copy operation. <p> Note: You can use the entire OID rather than the object name. Use the form: <i>OID.index</i> as shown in Figure 30-15.</p>

[Figure 30-14](#) and [Figure 30-15](#) are examples of using the snmpget command 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
- the file *f10-copy-config.mib* is in the current directory

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

[Figure 30-14](#) shows the command syntax using MIB object names, and [Figure 30-15](#) 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 30-14. Obtaining MIB Object Values for a Copy Operation using Object-name Syntax

```
> snmpget -v 2c -c private -m ./f10-copy-config.mib 10.11.131.140 copyTimeCompleted.110
FORCE10-COPY-CONFIG-MIB::copyTimeCompleted.110 = Timeticks: (1179831) 3:16:38.31
```

Figure 30-15. Obtaining MIB Object Values for a Copy Operation using OID Syntax

```
> snmpget -v 2c -c private 10.11.131.140 .1.3.6.1.4.1.6027.3.5.1.1.1.1.13.110
SNMPv2-SMI::enterprises.6027.3.5.1.1.1.1.13.110 = Timeticks: (1179831) 3:16:38.31
```

Manage VLANs Using SNMP

The qBridgeMIB managed objects in the Q-BRIDGE-MIB, defined in RFC 2674, enable you to use SNMP to manage VLANs.

Create a VLAN

Use the dot1qVlanStaticRowStatus object to create a VLAN. The snmpset operation in [Figure 30-16](#) creates VLAN 10 by specifying a value of 4 for instance 10 of the dot1qVlanStaticRowStatus object.

Figure 30-16. Creating a VLAN Using SNMP

```
> snmpset -v2c -c mycommunity 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
```

Assign a VLAN Alias

Write a character string to the dot1qVlanStaticName object to assign a name to a VLAN as shown in [Figure 30-17](#).

Figure 30-17. Assign a VLAN Alias Using SNMP

[Unix system output]

```
> snmpset -v2c -c mycommunity 10.11.131.185 .1.3.6.1.2.1.17.7.1.4.3.1.1.1107787786 s "My
VLAN"
SNMPv2-SMI::mib-2.17.7.1.4.3.1.1.1107787786 = STRING: "My VLAN"
```

[FTOS system output]

```
FTOS#show int vlan 10
Vlan 10 is down, line protocol is down
Vlan alias name is: My VLAN
Address is 00:01:e8:cc:cc:ce, Current address is 00:01:e8:cc:cc:ce
Interface index is 1107787786
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 01:01:00
Queueing strategy: fifo
Time since last interface status change: 01:01:00
```

Display the Ports in a VLAN

FTOS identifies VLAN interfaces using an interface index number that is displayed in the output of the command `show interface vlan`, as shown in [Figure 30-18](#).

Figure 30-18. Identifying the VLAN Interface Index Number

```
FTOS(conf)#do show interface vlan 10
Vlan 10 is down, line protocol is down
Address is 00:01:e8:cc:cc:ce, Current address is 00:01:e8:cc:cc:ce
Interface index is 1107787786
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 00:12:42
Queueing strategy: fifo
Time since last interface status change: 00:12:42
```

To display the ports in a VLAN, send an `snmpget` request for the object `dot1qStaticEgressPorts` using the interface index as the instance number, as shown in [Figure 30-19](#).

Figure 30-19. Display the Ports in a VLAN in SNMP

```
> snmpget -v2c -c mycommunity 10.11.131.185 .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 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00
```

The table that the Dell Force10 system sends in response to the `snmpget` request is a table that contains hexadecimal (hex) pairs, each pair representing a group of eight ports.

- 7 hex pairs represents a stack unit. Seven pairs accommodates the greatest number of ports available on an MXL Switch, 56 ports. The last stack unit is assigned 8 pairs; the eighth pair is unused.

The first hex pair, `00` in [Figure 30-19](#), represents ports 1-7 in Stack Unit 0. The next pair to the right represents ports 8-15. To resolve the hex pair into a representation of the individual ports, convert the hex pair to binary. Consider the first hex pair `00`, which resolves to `0000 0000` in binary:

- Each position in the 8-character string is for one port, starting with Port 1 at the left end of the string, and ending with Port 8 at the right end. A 0 indicates that the port is not a member of the VLAN; a 1 indicates VLAN membership.

Figure 30-19 shows the output for an MXL Switch. All hex pairs are 00, indicating that no ports are assigned to VLAN 10. In Figure 30-20, Port 0/2 is added to VLAN 10 as untagged. And the first hex pair changes from 00 to 04.

Figure 30-20. Displaying Ports in a VLAN using SNMP

```
[Dell Force10 system output]

FTOS(conf)#do show vlan id 10

Codes: * - Default VLAN, G - GVRP VLANs
Q: U - Untagged, T - Tagged
  x - Dot1x untagged, X - Dot1x tagged
  G - GVRP tagged, M - Vlan-stack

      NUM      Status      Description                               Q Ports
      10      Inactive

[Unix system output]

> snmpget -v2c -c mycommunity 10.11.131.185 .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: 40 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00
```

The value 40 is in the first set of 7 hex pairs, indicating that these ports are in Stack Unit 0. The hex value 40 is 0100 0000 in binary. As described above, the left-most position in the string represents Port 1. The next position from the left represents Port 2 and has a value of 1, indicating that Port 0/2 is in VLAN 10. The remaining positions are 0, so those ports are not in the VLAN.

Note that the table contains none of the other information provided by the show vlan command, such as port speed or whether the ports are tagged or untagged.

Add Tagged and Untagged Ports to a VLAN

The value dot1qVlanStaticEgressPorts object is an array of all VLAN members.

The dot1qVlanStaticUntaggedPorts object is an array of only untagged VLAN members. All VLAN members that are not in dot1qVlanStaticUntaggedPorts are tagged.

- To add a tagged port to a VLAN, write the port to the dot1qVlanStaticEgressPorts object, as shown in Figure 30-21.
- To add an untagged port to a VLAN, write the port to the dot1qVlanStaticEgressPorts and dot1qVlanStaticUntaggedPorts objects, as shown in Figure 30-22.



Note: Whether adding a tagged or untagged port, you must specify values for both dot1qVlanStaticEgressPorts and dot1qVlanStaticUntaggedPorts.

In Figure 30-21, Port 0/2 is added as an untagged member of VLAN 10.

Figure 30-21. Adding Untagged Ports to a VLAN using SNMP

```
>snmpset -v2c -c mycommunity 10.11.131.185 .1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786 x "40 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
.1.3.6.1.2.1.17.7.1.4.3.1.4.1107787786 x "40 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 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: 40 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 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-2.17.7.1.4.3.1.4.1107787786 = Hex-STRING: 40 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

In Figure 30-22, Port 0/2 is added as a tagged member of VLAN 10.

Figure 30-22. Adding Tagged Ports to a VLAN using SNMP

```
>snmpset -v2c -c mycommunity 10.11.131.185 .1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786 x "40 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
.1.3.6.1.2.1.17.7.1.4.3.1.4.1107787786 x "00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 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: 40 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 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-2.17.7.1.4.3.1.4.1107787786 = 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
```

Enable and Disable a Port Using SNMP

Step	Task	Command Syntax	Command Mode
1	Create an SNMP community on the Dell Force10 system.	snmp-server community	CONFIGURATION
2	From the Dell Force10 system, identify the interface index of the port for which you want to change the admin status. Or, from the management system, use the snmpwalk command to identify the interface index.	show interface	EXEC Privilege
3	Enter the command snmpset to change the admin status using either the object descriptor or the OID. Choose integer 1 to change the admin status to Up, or 2 to change the admin status to Down. snmpset with descriptor: <code>snmpset -v version -c community agent-ip ifAdminStatus .ifindex i {1 2}</code> snmpset with OID: <code>snmpset -v version -c community agent-ip .1.3.6.1.2.1.2.2.1.7.ifindex i {1 2}</code>		

Fetch Dynamic MAC Entries Using SNMP

Dell Force10 supports the RFC 1493 dot1d table for the default VLAN and the dot1q table for all other VLANs.



Note: The 802.1q Q-BRIDGE MIB defines VLANs with regard to 802.1d, as 802.1d itself does not define them. As a switchport must belong to a VLAN (the default VLAN or a configured VLAN), all MAC address learned on a switchport are associated with a VLAN. For this reason, the Q-Bridge MIB is used for MAC address query. Moreover, specific to MAC address query, dot1dTpFdbTable is indexed by MAC address only for a single forwarding database, while dot1qTpFdbTable has two indices —VLAN ID and MAC address—to allow for multiple forwarding databases and considering that the same MAC address is learned on multiple VLANs. The VLAN ID is added as the first index so that MAC addresses can be read by VLAN, sorted lexicographically. The MAC address is part of the OID instance, so in this case, lexicographic order is according to the most significant octet.

Table 30-6. MIB Objects for Fetching Dynamic MAC Entries in the Forwarding Database

MIB Object	OID	Description	MIB
dot1dTpFdbTable	.1.3.6.1.2.1.17.4.3	List the learned unicast MAC addresses on the default VLAN.	Q-BRIDGE MIB
dot1qTpFdbTable	.1.3.6.1.2.1.17.7.1.2. 2	List the learned unicast MAC addresses on non-default VLANs.	
dot3aCurAggFdb Table	.1.3.6.1.4.1.6027.3.2. 1.1.5	List the learned MAC addresses of aggregated links (LAG).	F10-LINK-AGGREGATION -MIB

In Figure 30-23, R1 has one dynamic MAC address, learned off of port TenGigabitEthernet 1/21, which is a member of the default VLAN, VLAN 1. The SNMP walk returns the values for dot1dTpFdbAddress, dot1dTpFdbPort, and dot1dTpFdbStatus.

Each object is comprised an OID concatenated with an instance number. In the case of these objects, the instance number is the decimal equivalent of the MAC address; derive the instance number by converting each hex pair to its decimal equivalent. For example, the decimal equivalent of E8 is 232, and so the instance number for MAC address 00:01:e8:06:95:ac is .0.1.232.6.149.172.

The value of dot1dTpFdbPort is the port number of the port off which the system learns the MAC address. In this case, of TenGigabitEthernet 1/21, the manager returns the integer 118.

Figure 30-23. Fetching Dynamic MAC Addresses on the Default VLAN

```

-----MAC Addresses on Dell Force10
System-----
FTOS#show mac-address-table
VlanId      Mac Address      Type      Interface      State
  1          00:01:e8:06:95:ac  Dynamic  Tengig 1/21    Active
-----Query from Management Station-----
>snmpwalk -v 2c -c techpubs 10.11.131.162 .1.3.6.1.2.1.17.4.3.1
SNMPv2-SMI::mib-2.17.4.3.1.1.0.1.232.6.149.172 = Hex-STRING: 00 01 E8 06 95 AC
SNMPv2-SMI::mib-2.17.4.3.1.2.0.1.232.6.149.172 = INTEGER: 118
SNMPv2-SMI::mib-2.17.4.3.1.3.0.1.232.6.149.172 = INTEGER: 3

```

In Figure 30-24, TenGigabitEthernet 1/21 is moved to VLAN 1000, a non-default VLAN. Use the objects dot1qTpFdbTable to fetch the MAC addresses learned on non-default VLANs. The instance number is the VLAN number concatenated with the decimal conversion of the MAC address.

Figure 30-24. Fetching Dynamic MAC Addresses on Non-default VLANs

```

-----MAC Addresses on Dell Force10
System-----
FTOS#show mac-address-table
VlanId      Mac Address      Type      Interface      State
  1000       00:01:e8:06:95:ac  Dynamic  Tengig 1/21    Active
-----Query from Management Station-----
>snmpwalk -v 2c -c techpubs 10.11.131.162 .1.3.6.1.2.1.17.7.1.2.2.1
SNMPv2-SMI::mib-2.17.7.1.2.2.1.2.1000.0.1.232.6.149.172 = INTEGER: 118
SNMPv2-SMI::mib-2.17.7.1.2.2.1.3.1000.0.1.232.6.149.172 = INTEGER: 3

```

Use dot3aCurAggFdbTable to fetch the learned MAC address of a port-channel. The instance number is the decimal conversion of the MAC address concatenated with the port-channel number.

Figure 30-25. Fetching Dynamic MAC Addresses on the Default VLAN

```

-----MAC Addresses on Dell Force10
System-----
FTOS(conf)#do show mac-address-table
VlanId      Mac Address      Type      Interface      State
  1000       00:01:e8:06:95:ac  Dynamic  Po 1           Active
-----Query from Management Station-----
>snmpwalk -v 2c -c techpubs 10.11.131.162 .1.3.6.1.4.1.6027.3.2.1.1.5
SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.1.1000.0.1.232.6.149.172.1 = INTEGER: 1000
SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.2.1000.0.1.232.6.149.172.1 = Hex-STRING: 00 01 E8 06 95 AC
SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.3.1000.0.1.232.6.149.172.1 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.4.1000.0.1.232.6.149.172.1 = INTEGER: 1

```

Deriving Interface Indices


FTOS assigns an interface number to each (configured or unconfigured) physical and logical interface. Display the interface index number using the command **show interface** from EXEC Privilege mode, as shown in [Figure 30-26](#).

Figure 30-26. Display the Interface Index Number

```
FTOS#show interface tengig 1/21
TenGigabitEthernet 1/21 is up, line protocol is up
Hardware is Dell Force10Eth, address is 00:01:e8:0d:b7:4e
  Current address is 00:01:e8:0d:b7:4e
Interface index is 72925242
[output omitted]
```

The interface index is a binary number with bits that indicate the slot number, port number, interface type, and card type of the interface. FTOS converts this binary index number to decimal, and displays it in the output of the **show interface** command.

Figure 30-27. Interface Index Binary Calculations




Starting from the least significant bit (LSB):

- the first 14 bits represent the card type
- the next 4 bits represent the interface type
- the next 7 bits represent the port number
- the next 5 bits represent the slot number
- the next 1 bit is 0 for a physical interface and 1 for a logical interface
- the next 1 bit is unused

For example, the index 44634369 is 10101010010001000100000001 in binary. The binary interface index for TenGigabitEthernet 0/41 of an MXL Switch is shown in [Figure 30-28](#). Notice that the physical/logical bit and the final, unused bit are not given. The interface is physical, so this must be represented by a 0 bit, and the unused bit is always 0. These two bits are not given because they are the most significant bits, and leading zeros are often omitted.

Figure 30-28. Binary Representation of Interface Index



For interface indexing, slot and port numbering begins with binary one. If the Dell Force10 system begins slot and port numbering from 0, then binary 1 represents slot and port 0. In S4810, the first interface is 0/0, but in the MXL Switch the first interface is 0/1. Hence, in the MXL Switch 0/0s ifindex is unused and Ifindex creation logic is not changed. Since Zero is reserved for logical interfaces, it starts from 1. For the first interface, port number is set to 1. Adding it causes an increment by 1 for the next interfaces, so it only starts from 2. Therefore, the port number is set to 42 for 0/41.

Note that the interface index does not change if the interface reloads or fails over. If the unit is renumbered (for any reason) the interface index will change during a reload.

Monitor Port-channels

To check the status of a Layer 2 port-channel, use `f10LinkAggMib (.1.3.6.1.4.1.6027.3.2)`. Below, Po 1 is a switchport and Po 2 is in Layer 3 mode.

```
[senthilnathan@lithium ~]$ snmpwalk -v 2c -c public 10.11.1.1 .1.3.6.1.4.1.6027.3.2.1.1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.1 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.2 = INTEGER: 2
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.2.1 = Hex-STRING: 00 01 E8 13 A5 C7
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.2.2 = Hex-STRING: 00 01 E8 13 A5 C8
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.3.1 = INTEGER: 1107755009
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.3.2 = INTEGER: 1107755010
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.4.1 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.4.2 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.5.1 = Hex-STRING: 00 00
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.5.2 = Hex-STRING: 00 00
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.6.1 = STRING: "Tengig 5/84 " << Channel member for Po1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.6.2 = STRING: "Tengig 5/85 " << Channel member for Po2
dot3aCommonAggFdbIndex
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.1.1107755009.1 = INTEGER: 1107755009
dot3aCommonAggFdbVlanId
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.2.1107755009.1 = INTEGER: 1
dot3aCommonAggFdbTagConfig
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.3.1107755009.1 = INTEGER: 2 (Tagged 1 or Untagged 2)
dot3aCommonAggFdbStatus
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.4.1107755009.1 = INTEGER: 1 << Status active, 2 - status inactive
```

If we learn mac address for the LAG, status is shown for those as well.

```
dot3aCurAggVlanId
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.1.1.0.0.0.0.0.1.1 = INTEGER: 1
dot3aCurAggMacAddr
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.2.1.0.0.0.0.0.1.1 = Hex-STRING: 00 00 00 00 00 01
dot3aCurAggIndex
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.3.1.0.0.0.0.0.1.1 = INTEGER: 1
dot3aCurAggStatus
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.4.1.0.0.0.0.0.1.1 = INTEGER: 1 << Status active, 2 - status
inactive
```

For L3 lag we do not have this support. SNMP trap works fine for L2 / L3 / default mode LAG.

```
SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500842) 23:36:48.42
SNMPv2-MIB::snmpTrapOID.0 = OID: IF-MIB::linkDown
IF-MIB::ifIndex.33865785 = INTEGER: 33865785
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_DN: Changed interface state to down: Tengig 0/0"
2010-02-10 14:22:39 10.16.130.4 [10.16.130.4]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500842) 23:36:48.42
SNMPv2-MIB::snmpTrapOID.0 = OID: IF-MIB::linkDown
IF-MIB::ifIndex.1107755009 = INTEGER: 1107755009
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_DN: Changed interface state to down: Po 1"
2010-02-10 14:22:40 10.16.130.4 [10.16.130.4]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500932) 23:36:49.32      SNMPv2-MIB::snmpTrapOID.0 = OID:
IF-MIB::linkUp IF-MIB::ifIndex.33865785 = INTEGER: 33865785      SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 =
STRING: "OSTATE_UP: Changed interface state to up: Tengig 0/0"
2010-02-10 14:22:40 10.16.130.4 [10.16.130.4]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500934) 23:36:49.34      SNMPv2-MIB::snmpTrapOID.0 = OID:
IF-MIB::linkUp IF-MIB::ifIndex.1107755009 = INTEGER: 1107755009
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_UP: Changed interface state to up: Po 1"
```

BMP functionality using SNMP SET

You can enable/disable BMP functionality via SNMP SET.:

Table 30-7. List of Jumpstart MIBs that have both read/write access

MIB Object	OID	Description
f10JumpStartMib	.1.3.6.1.4.1.6027.3.23	NODE
f10JumpStart	.1.3.6.1.4.1.6027.3.23.1	NODE
jsReloadType	.1.3.6.1.4.1.6027.3.23.1.1	LEAF INTEGER
jsAutoSave	.1.3.6.1.4.1.6027.3.23.1.2	LEAF INTEGER
jsConfigDownload	.1.3.6.1.4.1.6027.3.23.1.3	LEAF INTEGER
jsDhcpTimeout e	.1.3.6.1.4.1.6027.3.23.1.4	LEAF INTEGER
jsRetryCount	.1.3.6.1.4.1.6027.3.23.1.5	LEAF INTEGER

Entity MIBS

The Entity MIB provides a mechanism for presenting hierarchies of physical entities using SNMP tables. The Entity MIB contains the following groups, which describe the physical elements and logical elements of a managed system. The following tables are implemented for the MXL Switch Platform.

Physical Entity

A physical entity or physical component represents an identifiable physical resource within a managed system. Zero or more logical entities may utilize a physical resource at any given time. Determining which physical components are represented by an agent in the `EntPhysicalTable` is an implementation-specific matter. Typically, physical resources (e.g., communications ports, backplanes, sensors, daughter-cards, power supplies, the overall chassis), which can be managed via functions associated with one or more logical entities, are included in the MIB.

Containment Tree

Each physical component may be modeled as contained within another physical component. A containment-tree is the conceptual sequence of *entPhysicalIndex* values that uniquely specifies the exact physical location of a physical component within the managed system. It is generated by following and recording each *entPhysicalContainedIn* instance up the tree towards the root, until a value of zero indicating no further containment is found.

Figure 30-29. Sample Entity MIBS outputs

```
FTOS#show inventory optional-module
Unit Slot Expected Inserted Next Boot Power
-----
0      0    QSFP+      QSFP+      AUTO Good
0      1   10GBASE-T  10GBASE-T  AUTO Good
1      0    QSFP+      QSFP+      AUTO Good
1      1   10GBASE-T  10GBASE-T  AUTO Good
2      0    QSFP+      QSFP+      AUTO Good
2      1    SFP+       SFP+       AUTO Good
```

The status for the MIBS is as follows:

```
vijayakrishnan@tapti[3:42pm] : /tftpboot > snmpwalk -c public -v 2c 10.16.130.135
1.3.6.1.2.1.47.1.1.1.1.2
SNMPv2-SMI::mib-2.47.1.1.1.1.2.1 = ""
SNMPv2-SMI::mib-2.47.1.1.1.1.2.2 = STRING: "PowerConnect MXL 10/40GbE"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.3 = STRING: "Module 0"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.4 = STRING: "Unit: 0 Port 1 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.5 = STRING: "Unit: 0 Port 2 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.6 = STRING: "Unit: 0 Port 3 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.7 = STRING: "Unit: 0 Port 4 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.8 = STRING: "Unit: 0 Port 5 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.9 = STRING: "Unit: 0 Port 6 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.10 = STRING: "Unit: 0 Port 7 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.11 = STRING: "Unit: 0 Port 8 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.12 = STRING: "Unit: 0 Port 9 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.13 = STRING: "Unit: 0 Port 10 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.14 = STRING: "Unit: 0 Port 11 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.15 = STRING: "Unit: 0 Port 12 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.16 = STRING: "Unit: 0 Port 13 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.17 = STRING: "Unit: 0 Port 14 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.18 = STRING: "Unit: 0 Port 15 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.19 = STRING: "Unit: 0 Port 16 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.20 = STRING: "Unit: 0 Port 17 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.21 = STRING: "Unit: 0 Port 18 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.22 = STRING: "Unit: 0 Port 19 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.23 = STRING: "Unit: 0 Port 20 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.24 = STRING: "Unit: 0 Port 21 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.25 = STRING: "Unit: 0 Port 22 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.26 = STRING: "Unit: 0 Port 23 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.27 = STRING: "Unit: 0 Port 24 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.28 = STRING: "Unit: 0 Port 25 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.29 = STRING: "Unit: 0 Port 26 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.30 = STRING: "Unit: 0 Port 27 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.31 = STRING: "Unit: 0 Port 28 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.32 = STRING: "Unit: 0 Port 29 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.33 = STRING: "Unit: 0 Port 30 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.34 = STRING: "Unit: 0 Port 31 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.35 = STRING: "Unit: 0 Port 32 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.36 = STRING: "40G QSFP+ port"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.37 = STRING: "Unit: 0 Port 33 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.41 = STRING: "40G QSFP+ port"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.42 = STRING: "Unit: 0 Port 37 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.46 = STRING: "Optional module 0"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.47 = STRING: "2-port 40G QSFP (XL)"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.48 = STRING: "Unit: 0 Port 41 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.52 = STRING: "Unit: 0 Port 45 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.56 = STRING: "Optional module 1"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.57 = STRING: "4-port 10GE BASE-T (XL) "
SNMPv2-SMI::mib-2.47.1.1.1.1.2.58 = STRING: "Unit: 0 Port 49 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.59 = STRING: "Unit: 0 Port 50 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.60 = STRING: "Unit: 0 Port 51 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.61 = STRING: "Unit: 0 Port 52 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.66 = STRING: "PowerConnect MXL 10/40GbE"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.67 = STRING: "Module 0"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.68 = STRING: "Unit: 1 Port 1 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.69 = STRING: "Unit: 1 Port 2 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.70 = STRING: "Unit: 1 Port 3 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.71 = STRING: "Unit: 1 Port 4 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.72 = STRING: "Unit: 1 Port 5 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.73 = STRING: "Unit: 1 Port 6 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.74 = STRING: "Unit: 1 Port 7 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.75 = STRING: "Unit: 1 Port 8 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.76 = STRING: "Unit: 1 Port 9 10G Level"
```


SNMPv2-SMI::mib-2.47.1.1.1.1.2.77 = STRING: "Unit: 1 Port 10 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.78 = STRING: "Unit: 1 Port 11 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.79 = STRING: "Unit: 1 Port 12 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.80 = STRING: "Unit: 1 Port 13 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.81 = STRING: "Unit: 1 Port 14 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.82 = STRING: "Unit: 1 Port 15 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.83 = STRING: "Unit: 1 Port 16 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.84 = STRING: "Unit: 1 Port 17 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.85 = STRING: "Unit: 1 Port 18 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.86 = STRING: "Unit: 1 Port 19 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.87 = STRING: "Unit: 1 Port 20 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.88 = STRING: "Unit: 1 Port 21 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.89 = STRING: "Unit: 1 Port 22 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.90 = STRING: "Unit: 1 Port 23 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.91 = STRING: "Unit: 1 Port 24 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.92 = STRING: "Unit: 1 Port 25 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.93 = STRING: "Unit: 1 Port 26 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.94 = STRING: "Unit: 1 Port 27 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.95 = STRING: "Unit: 1 Port 28 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.96 = STRING: "Unit: 1 Port 29 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.97 = STRING: "Unit: 1 Port 30 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.98 = STRING: "Unit: 1 Port 31 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.99 = STRING: "Unit: 1 Port 32 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.100 = STRING: "40G QSFP+ port"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.101 = STRING: "Unit: 1 Port 33 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.105 = STRING: "40G QSFP+ port"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.106 = STRING: "Unit: 1 Port 37 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.110 = STRING: "Optional module 0"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.111 = STRING: "2-port 40G QSFP (XL)"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.112 = STRING: "Unit: 1 Port 41 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.116 = STRING: "Unit: 1 Port 45 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.120 = STRING: "Optional module 1"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.121 = STRING: "4-port 10GE BASE-T (XL) "
SNMPv2-SMI::mib-2.47.1.1.1.1.2.122 = STRING: "Unit: 1 Port 49 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.123 = STRING: "Unit: 1 Port 50 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.124 = STRING: "Unit: 1 Port 51 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.125 = STRING: "Unit: 1 Port 52 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.130 = STRING: "PowerConnect MXL 10/40GbE"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.131 = STRING: "Module 0"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.132 = STRING: "Unit: 2 Port 1 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.133 = STRING: "Unit: 2 Port 2 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.134 = STRING: "Unit: 2 Port 3 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.135 = STRING: "Unit: 2 Port 4 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.136 = STRING: "Unit: 2 Port 5 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.137 = STRING: "Unit: 2 Port 6 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.138 = STRING: "Unit: 2 Port 7 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.139 = STRING: "Unit: 2 Port 8 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.140 = STRING: "Unit: 2 Port 9 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.141 = STRING: "Unit: 2 Port 10 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.142 = STRING: "Unit: 2 Port 11 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.143 = STRING: "Unit: 2 Port 12 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.144 = STRING: "Unit: 2 Port 13 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.145 = STRING: "Unit: 2 Port 14 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.146 = STRING: "Unit: 2 Port 15 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.147 = STRING: "Unit: 2 Port 16 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.148 = STRING: "Unit: 2 Port 17 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.149 = STRING: "Unit: 2 Port 18 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.150 = STRING: "Unit: 2 Port 19 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.151 = STRING: "Unit: 2 Port 20 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.152 = STRING: "Unit: 2 Port 21 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.153 = STRING: "Unit: 2 Port 22 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.154 = STRING: "Unit: 2 Port 23 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.155 = STRING: "Unit: 2 Port 24 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.156 = STRING: "Unit: 2 Port 25 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.157 = STRING: "Unit: 2 Port 26 10G Level"

```

SNMPv2-SMI::mib-2.47.1.1.1.1.2.158 = STRING: "Unit: 2 Port 27 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.159 = STRING: "Unit: 2 Port 28 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.160 = STRING: "Unit: 2 Port 29 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.161 = STRING: "Unit: 2 Port 30 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.162 = STRING: "Unit: 2 Port 31 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.163 = STRING: "Unit: 2 Port 32 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.164 = STRING: "40G QSFP+ port"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.165 = STRING: "Unit: 2 Port 33 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.169 = STRING: "40G QSFP+ port"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.170 = STRING: "Unit: 2 Port 37 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.174 = STRING: "Optional module 0"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.175 = STRING: "2-port 40G QSFP (XL)"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.176 = STRING: "Unit: 2 Port 41 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.180 = STRING: "Unit: 2 Port 45 40G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.184 = STRING: "Optional module 1"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.185 = STRING: "4-port 10GE SFP+ (XL) "
SNMPv2-SMI::mib-2.47.1.1.1.1.2.186 = STRING: "Unit: 2 Port 49 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.187 = STRING: "Unit: 2 Port 50 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.188 = STRING: "Unit: 2 Port 51 10G Level"
o SNMPv2-SMI::mib-2.47.1.1.1.1.2.189 = STRING: "Unit: 2 Port 52 10G Level"

```

Troubleshooting SNMP Operations

When you use SNMP to retrieve management data from an SNMP agent on a Dell Force10 router, take into account the following behavior:

- When you query an IPv4 icmpMsgStatsInPkts object in the ICMP table by using the snmpwalk command, the output for echo replies may be incorrectly displayed. To correctly display this information under ICMP statistics, use the show ip traffic command.
- When you query an icmpStatsInErrors object in the icmpStats table by using the snmpget or snmpwalk commands, the output for IPv4 addresses may be incorrectly displayed. To correctly display this information under IP and ICMP statistics, use the show ip traffic command.
- When you query an IPv4 icmpMsgStatsInPkts object in the ICMP table by using the snmpwalk command, the echo response output may not be displayed. To correctly display ICMP statistics, such as echo response, use the show ip traffic command.

Stacking

Overview

Stacking is supported on a MXL 10/40GbE Switch on the 40GbE ports (for the base module) or a 2-Port 40GbE QSFP+ module. You can connect up to six MXL 10/40GbE Switches in a single stack. Stacking provides a single point of management and network interface controller (NIC) teaming for high availability and higher throughput.

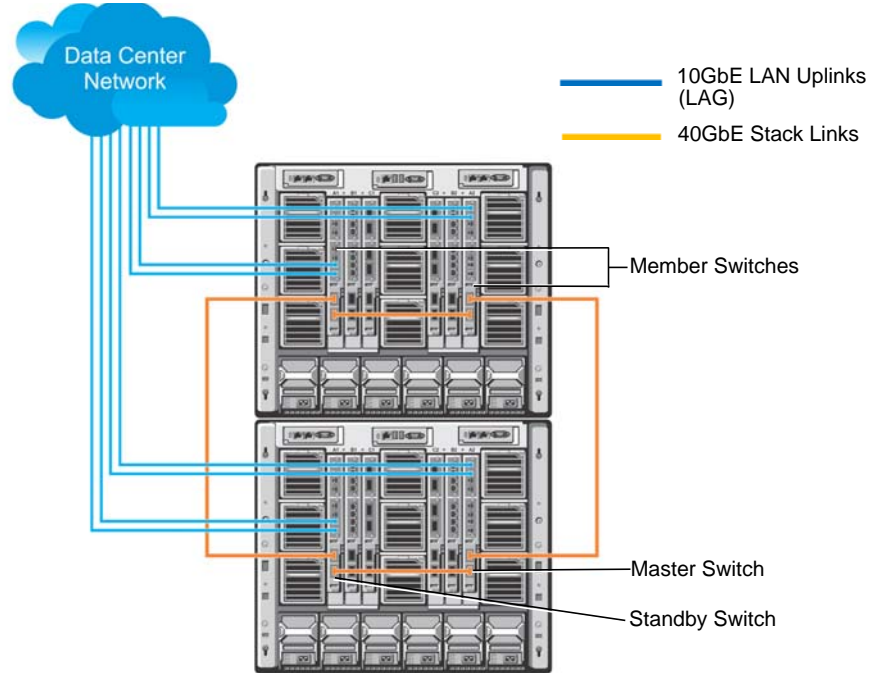
This chapter contains the following sections:

- [Stacking MXL 10/40GbE Switches](#)
- [Stack Group/Port Numbers](#)
- [Configuring a Switch Stack](#)
- [Verifying a Stack Configuration](#)
- [Troubleshooting a Switch Stack](#)
- [Upgrading a Switch Stack](#)
- [Upgrading a Single Stack Unit](#)

Stacking MXL 10/40GbE Switches

A stack of MXL 10/40GbE Switches operates as a virtual chassis with management units (primary and standby) and member units. The Dell Force10 operating software (FTOS) elects a primary (master) and secondary (standby) management unit; all other units are member units. The forwarding database resides on the master switch; all other stack units maintain a synchronized local copy. Each unit in the stack makes forwarding decisions based on their local copy.

[Figure 31-1](#) shows an example of how you can stack four MXL 10/40GbE Switches and the role played by each switch in the stack. The MXL 10/40GbE Switches are connected to operate as a single stack in a ring topology using only the 40GbE ports on the base modules. You can use the 40GbE ports on the base module and FlexIO modules to create a stack in either a ring or daisy-chain topology.

Figure 31-1. Four Stacked MXL 10/40GbE Switches

Stack Management Roles

The stack elects the management units for the stack management:

- Stack master: primary management unit
- Standby: secondary management unit

The master holds the control plane and the other units maintain a local copy of the forwarding databases. From Stack master you can configure:

- System-level features that apply to all stack members
- Interface-level features for each stack member

The master synchronizes the following information with the standby unit:

- Stack unit topology
- Stack running Configuration (which includes ACL, LACP, STP, SPAN, etc.)
- Logs

The master switch maintains stack operation with minimal impact in the event of:

- Switch failure
- Inter-switch stacking link failure
- Switch insertion
- Switch removal

If the master switch goes off line, the standby replaces it as the new master and the switch with the next highest priority or MAC address becomes standby.



Note: For the MXL Switch, the entire stack has only one management IP address.

Stack Master Election

The stack elects a master and standby unit at bootup time based on two criteria:

- Unit priority: This is user-configurable. Valid values are from 1 to 14. A higher value means a higher priority. The default is 0. To remove the stack-unit priority and set the priority back to the default value of zero, use the `no stack-unit priority` command.
- MAC address (in case of priority tie): The unit with the higher MAC value becomes master.

To view which switch is the stack master, use the `show system` command. [Figure 31-2](#) shows sample output from an established stack.

A change in the stack master occurs when:

- You power down the stack master or bring the master switch offline.
- A failover of the master switch occurs.
- You disconnect the master switch from the stack.



Note: When a stack reloads and all the units come up at the same time; for example, when all units boot up from flash, all units participate in the election and the master and standby are chosen based on the priority or MAC address. When the units do not boot up at the same time; for example, some units are powered down just after reloading and powered up later to join the stack, they do not participate in the election process, even though the units that boot up late may have a higher priority configured. This happens because the master and standby have already been elected; therefore, the unit that boots up late joins only as a member. Also, when an up and running standalone unit or stack is merged with another stack, based on election, the losing stack reloads and the master unit of the winning stack becomes the master of the merged stack.

For more information, refer to [“Adding a Stack Unit”](#) and [“Merging Two Stacks”](#). To ensure a fully synchronised bootup, it is possible to reset individual units to force them to give up the management role; or reload the whole stack from the command line interface (CLI).

Figure 31-2. Displaying the Stack Master

```

FTOS# show system brief

Stack MAC : 00:1e:c9:f1:00:7b

Reload Type : jump-start [Next boot : normal-reload]

-- Stack Info --
Unit  UnitType      Status   ReqTyp          CurTyp          Version        Ports
-----
  0    Management    online   MXL-10/40GbE   MXL-10/40GbE   9-1-0-853     56
  1    Standby        online   MXL-10/40GbE   MXL-10/40GbE   9-1-0-853     56
  2    Member         online   MXL-10/40GbE   MXL-10/40GbE   9-1-0-853     56
  3    Member         online   MXL-10/40GbE   MXL-10/40GbE   9-1-0-853     56
  4    Member         online   MXL-10/40GbE   MXL-10/40GbE   9-1-0-853     56
  5    Member         online   MXL-10/40GbE   MXL-10/40GbE   9-1-0-853     56
FTOS#

```

Failover Roles

If the stack master fails (for example, powered off), it is removed from the stack topology. The standby unit detects the loss of peering communication and takes ownership of the stack management, switching from standby to master. The lack of a standby unit triggers an election within the remaining units for a standby role.

After the former master switch recovers, despite having a higher priority or MAC address, it does not recover its master role but instead take the next available role.

MAC Addressing

All port interfaces in the stack use the MAC address of the management interface on the master switch. The MAC address of the chassis in which the master MXL Switch is installed is used as the stack MAC address.

The stack continues to use the master's chassis MAC address even after a failover. The MAC address is not refreshed until the stack is reloaded and a different unit becomes the stack master.

Stacking LAG

When you use multiple links between stack units, FTOS automatically bundles them in a stacking link aggregation group (LAG) to provide aggregated throughput and redundancy. The stacking LAG is established automatically and transparently by FTOS (without user configuration) after peering is detected and behaves as follows:

- The stacking LAG dynamically aggregates; it can lose link members or gain new links.
- Shortest path selection inside the stack: if multiple paths exist between two units in the stack, the shortest path is used.

Supported Stacking Topologies

Stacking is supported on the MXL 10/40GbE Switch in ring and daisy-chain topologies.


Example 1: Dual-Ring Stack Across Multiple Chassis

Using two separate stacks in a dual-ring stacking topology provides redundancy and increased high availability in case of stack failure. Also, stacking upgrades are simplified when you have to take one stack offline (Figure 31-3).



Note: A ring topology is recommended under normal operation because it provides increased resiliency when compared with a daisy chain topology. In daisy chain topology, any change in a non-edge stack unit causes a split stack.


Figure 31-3. Dual-Ring Stacking Topology for MXL 10/40GbE Switches



Example 2: Dual Daisy-Chain Stack Across Multiple Chassis

Using two separate, daisy-chained stacks in a stacking topology provides redundancy and increased high availability in case of stack failure. Also, stacking upgrades are simplified when you have to take one stack offline (Figure 31-4).


Figure 31-4. Dual Daisy-Chain Stacking Topology for MXL 10/40GbE Switches



Stack Group/Port Numbers

By default, each switch in Standalone mode is numbered stack-unit 0. Stack-unit numbers are assigned to member switches when the stack comes up. [Figure 31-5](#) shows the stack-group numbers of 40GbE ports on an MXL 10/40GbE Switch.

Figure 31-5. Stack Groups on an MXL 10/40GbE Switch



Configuring a Switch Stack

To configure and bring up a switch stack, follow these steps:

1. Connect the switches to be stacked with 40G direct attach or QSFP fibre cables.
2. Configure the stacking ports on each switch.
3. All switches must be booted together.
4. (Optional) Configure management priorities, unit numbers, or logical provisioning for stack units.

Stacking Prerequisites

Before you cable and configure a stack of MXL 10/40GbE Switches, review the following prerequisites:

- All MXL 10/40GbE Switches in the stack must be powered up with the initial or startup configuration before you attach the cables.
- All stacked MXL 10/40GbE Switches must run the same FTOS version. The minimum FTOS version required is 8.3.16.0. To check the FTOS version that a switch is running, use the show version command. To download an FTOS version, go to <http://support.dell.com>.
- Stacking is supported only with other MXL 10/40GbE Switches. A maximum of six MXL 10/40GbE Switches is supported in a single stack. You cannot stack the MXL 10/40GbE Switch with the M IO Aggregator or another type of switch.
- A maximum of four stack groups (40GbE ports) is supported on a stacked MXL 10/40GbE switch.
- Interconnect the stack units by following the instructions in [Cabling Stacked Switches](#).

Cabling Stacked Switches

Before you configure MXL Switches in a stack, connect the 40G direct attach or QSFP cables and transceivers to connect 40GbE ports on switches in the same or different chassis.

Cabling Restrictions

The following restrictions apply when setting up a stack of MXL 10/40GbE Switches:

- Only daisy-chain or ring topologies are supported; star and full mesh topologies are not supported.
- Stacking is supported only on 40GbE links by connecting 40GbE ports on the base module or a 2-Port QSFP+ module. Stacking is not supported on 10GbE ports or 4x10GbE ports.
To convert the 40GbE ports on the 2-Port QSFP+ module from their default 4x10GbE mode of operation to 40GbE mode, refer to [Converting 4x10GbE Ports to 40GbE for Stacking](#).
- Use only QSFP transceivers and QSFP or direct attach cables (purchased separately) to connect stacking ports.

Cabling Redundancy

Connect the units in a stack with two or more stacking cables to avoid a stacking port or cable failure. Removing one of the stacked cables between two stacked units does not trigger a reset.

Cabling Procedure

The following cabling procedure uses the stacking topology in [Figure 31-1](#). Follow the same steps to cable switches in any of the stacking topologies shown in [Supported Stacking Topologies](#). To connect the cabling, follow these steps:

1. Connect a 40GbE port on the first switch to a 40GbE port on the second switch.
2. Connect another 40GbE port on the second switch to a 40GbE port on the third switch.
3. Connect another 40GbE port on the third switch to a 40GbE port on the fourth switch.
4. Connect another 40GbE port on the fourth switch to a 40GbE port on the first switch.

The resulting ring topology allows the entire stack to function as a single switch with resilient fail-over capabilities. If you do not connect the last switch to the first switch (Step 4), the stack operates in a daisy chain topology with less resiliency. Any failure in a non-edge stack unit causes a split stack.


Accessing the CLI

To configure a stack, you must access the stack master in one of the following ways:

- For remote out-of-band management (OOB), enter the OOB management interface IP address into a Telnet or secure shell (SSH) client and log in to the switch using the user ID and password to access the CLI.
- For local management, use the attached console connection to the master switch to log in to the CLI. Console access to the stack CLI is available on the master only.
- For remote in-band management from a network management station, enter the virtual local area network (VLAN) IP address of the management port and log in to the switch to access the CLI.

Configuring and Bringing Up a Stack

After you attach the 40G QSFP or direct attach cables in a stack of MXL 10/40GbE Switches, to bring up the stack, follow these steps.

 **Note:** The procedure uses command examples for the stacking topology in [Figure 31-1](#).

Step	Task	Command Syntax	Command Mode
1	Set up a connection to the CLI on an MXL 10/40GbE Switch as described in Accessing the CLI .		
2	Log on to the CLI and enter Global Configuration mode.	Login: username Password: ***** FTOS> enable FTOS# configure	---

Step	Task	Command Syntax	Command Mode
3	Configure a 40GbE port for stacking mode, where: stack-unit <unit-number> is the unit-number of the member stack unit. Valid values: 0 to 5. Default value: 0. stack-group group-number is the number of stacked port on unit. Valid values: 0 to 1 (Figure 31-5).	stack-unit <i>unit-number</i> stack-group <i>group-number</i>	CONFIGURATION
4	Save the stacking configuration on the 40GbE ports.	write memory	EXEC PRIVILEGE
5	Repeat Steps 1 to 4 on each MXL 10/40GbE Switch in the stack.		
6	Log on to the CLI and reboot each switch, one after another, in as short a time as possible.	reload	EXEC PRIVILEGE



Note: If the stacked switches all reboot at approximately the same time, the switch with the highest MAC address is automatically elected as the master switch. The switch with the next highest MAC address is elected as standby. As each switch joins the stack, it is assigned the lowest available stack-unit number from 0 to 5. The default configuration of each stacked switch is stored in the running configuration of the stack. The stack-unit ID numbers are retained after future stack reloads.

To verify the stack-unit number assigned to each switch in the stack, use the show system brief command (Figure 31-8).

To configure stacked switches so that stacking roles are determined by preset priorities, use the stack-unit priority command (refer to [Assigning a Priority to Stacked Switches](#)).

Assigning a Priority to Stacked Switches

To configure the stack so that the roles are assigned according to pre-determined priorities instead of using the highest MAC addresses, use the stack-unit priority command in Global Configuration mode on each stacked switch. The switch with the highest priority number is elected master. The switch with the next highest priority number is elected standby and takes over stack management if the master switch fails.

Task	Command Syntax	Command Mode
Configure the priority of stacked switches to determine stack mastership, where stack-unit <i>unit-number</i> identifies the switch in the stack. priority <i>priority-number</i> specifies the management priority. Range: 1-14. Default: 0.	stack-unit <i>unit-number</i> priority <i>number</i>	CONFIGURATION

To revert the management priority of a stack unit to the default value of 0, use the no form of the stack-unit *unit-number* priority *number* command.

After you reconfigure the priorities of stacked switches, reload the stack so that a new master and standby election is performed.

Renumbering a Stack Unit

To renumber a stack unit to reset the unit numbering for a master, standby or member unit, enter the stack-unit renumber command in EXEC Privilege mode and reload the switch.

Task	Command Syntax	Command Mode
Assign a stack-number to a unit.	stack-unit <i>unit-number</i> renumber <i>new-number</i>	EXEC Privilege

- If you renumber the master switch, you are prompted to reload the entire stack.
- If you renumber the standby switch, only the switch reloads and is replaced by a member switch that is elected as the new standby.
- If you renumber a member switch, only the member switch reloads.
- If you renumber a switch to a number already assigned to another stack unit, the following error message is displayed:

Figure 31-6. Example

```
FTOS#stack-unit 5 renumber 0
% ERROR: stack unit 0 already exists.
```

Provisioning a Stack Unit

You can logically provision a stack-unit number to accept only an MXL 10/40GbE Switch. Provisioning is a type of pre-configuration that is stored on the master switch and applied when a stacked unit is assigned the unit number.

When you provision a unit number for an MXL 10/40GbE Switch:

- The base-module ports on the switch (ports 33 and 37/stack groups 0 and 1) are pre-configured for 40GbE operation.
- The 40GbE ports on FlexIO modules (ports 41 and 45 in slot 0; ports 49 and 53 in slot 1) are pre-configured for 4x10GbE (quad mode) operation.

To provision a stack unit, use the stack-unit provision command in Global Configuration mode, save the provisioning configuration, and reload the stack.

To provision a stack unit, use the following command:

Task	Command Syntax	Command Mode
Create a virtual stack unit by logically provisioning a switch.	stack-unit <i>unit-number</i> provision MXL-10/40GbE	CONFIGURATION



FTOS Behavior: Stacking configuration is handled as follows on an MXL 10/40GbE Switch:

- If a stack unit goes down and is removed from the stack, the logical provisioning configured for the stack-unit number is saved on the master and standby switches.
- When you add a new unit to the stack and the stack already has an existing member unit with the same stack-unit number, the new unit is assigned the smallest available unit number (0 to 5). A configuration mismatch between the newly added unit and a logically provisioned unit occurs in the following situations:
 - The logical provisioning for the unit number configures FlexIO module ports for 4x10GbE operation and the added unit has FlexIO Module ports operating in 40GbE mode.
 - The logical provisioning for the unit number and the added unit have different stack groups configured.

When a configuration mismatch occurs, the newly added switch enters into a Card-Problem state and is disabled. A syslog error message is generated. To restore a stacked switch in a Card-Problem state, refer to [Stack Unit in Card-Problem State Due to Configuration Mismatch](#).

- A stack unit can also enter a Card-Problem state after a split-stack reload in which a unit that was previously neither the master nor standby is elected as the new master and has logical stack-unit provisioning configured for a stack-unit number that creates a mismatch with the stack-unit numbering on other units.

Converting 4x10GbE Ports to 40GbE for Stacking

Stacking is supported only on 40GbE links by connecting 40GbE ports on the base module or a 2-Port QSFP+ module. However, on a 2-Port 40GbE QSFP+ module, the ports operate by default in 4x10GbE (quad) mode with breakout cables as eight 10GbE ports.

To change a port from 4x10GbE to 40GbE mode of operation for stacking, use the `no stack-unit port portmode quad` command (Figure 31-7). After you convert the 4x10GbE ports to 40GbE, you must save the configuration and reload the stack for the change to take effect.

Figure 31-7. `no stack-unit port portmode quad` Command Example

```
FTOS(conf)# no stack-unit unit-number port port-number portmode quad
```

Where:

`stack-unit unit-number` is the unit ID number in the stack unit. Valid values: 0 to 5. To display the stack-unit number, use the `show system brief` command.

`port port-number` specifies the port number of the QSFP+ port to be converted to 40GbE mode. Valid values are: base-module ports: 33 or 37; slot 0: 41 or 45; slot 1: 49 or 53 (refer to Figure 31-5).

`portmode quad` identifies the port as a split 10GbE SFP+ port.

Removing a Port from the Stacking Mode

To remove a 40GbE port from the stack, use the `no` form of the `stack-unit unit-number stack-group number` command. After entering the command, save the configuration and reload the stack for the change to take effect.

To remove a stack port, use the following command:

Task	Command Syntax	Command Mode
Remove a stacked port from a stack.	no stack-unit <i>unit-number</i> stack-group <i>group</i> end write memory reload	CONFIGURATION

When the reload completes, the port comes up in 40GbE mode if it is on the base module and in 4x10GbE (quad) mode if the port is on a FlexIO module, such as a 2-Port 40GbE QSFP+ module.

Removing a Switch from a Stack

After you remove all 40GbE ports from a stack ([Removing a Port from the Stacking Mode](#)), the switch functions in standalone mode but retains the running and startup configuration that was last synchronized by the master switch while it operated as a stack unit.

To remove a switch from a stack, disconnect the stacking cables from the unit either when the unit is powered on or off and is online or offline.

If you remove a unit from the middle of a stack, the stack is split into multiple parts. Each split stack forms a new stack according to MAC addresses or assigned priorities, as described in [Configuring and Bringing Up a Stack](#) and [Assigning a Priority to Stacked Switches](#).

Adding a Stack Unit

You can add a new unit to an existing stack both when the unit has no stacking ports (stack groups) configured and when the unit already has stacking ports configured. If the units to be added to the stack have been previously used, they are assigned the smallest available unit ID in the stack.

If a standalone switch has no stack groups configured, you can add it to a stack. To add a standalone switch to a stack, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Power on the switch.		
2	Attach QSFP or direct attach cables to connect 40GbE ports on the switch to one or more switches in the stack.		
3	Log on to the CLI and enter global configuration mode.	Login: username Password: ***** FTOS> enable FTOS# configure	---

Step	Task	Command Syntax	Command Mode
4	Configure a 40GbE port for stacking, where: stack-unit 0 defines the default ID unit-number in the initial configuration of a switch. stack-group <i>group-number</i> configures a 40GbE port for stacking. Base-module ports are stack groups 0 and 1; 40GbE ports on a FlexIO module in slot 0 are stack groups 2 and 3 and in slot 1 are stack groups 4 and 5 (Figure 31-5).	stack-unit 0 stack-group <i>group-number</i>	CONFIGURATION
5	Save the stacking configuration on the 40GbE ports.	write memory	EXEC Privilege
6	Reload the switch. FTOS automatically assigns a number to the new unit and adds it as member switch in the stack. The new unit synchronizes its running and startup configurations with the stack.	reload	EXEC Privilege

If a standalone switch already has stack groups configured, continue with these steps:

Step	Task	Command Syntax	Command Mode
7	Attach QSFP or direct attach cables to connect the 40GbE ports already configured as stack groups on the switch to one or more switches in the stack.		
8	Power on the switch. FTOS automatically assigns a number to the new unit and adds it as member switch in the stack. The new unit synchronizes its running and startup configurations with the stack.		



FTOS Behavior: When you add a new switch to a stack:

- If the new unit has been configured with a stack number that is already assigned to a stack member, the stack avoids a numbering conflict by assigning the new switch the first available stack number.
- If the stack has been provisioned for the stack number that is assigned to the new unit, the pre-configured provisioning must match the switch type. If there is a conflict between the provisioned switch type and the new unit, a mismatch error message is displayed.

Merging Two Stacks

You can merge two MXL 10/40GbE Switch stacks while they are powered and online. To merge two stacks, connect one stack to the other with 40G QSFP or direct attach cables. After you connect the stacking cables, a merge of the two stacks is performed:

- FTOS selects a master switch for the merged stack from the existing masters in the two stacks. To ensure that one of the two master switches wins the master election in the merged stack, use the stack-unit priority command to configure the highest priority for the unit (refer to [Assigning a Priority to Stacked Switches](#)).
- All the units in the losing stack go for a reboot and then merge with the winning stack that has the stack master.

- If there is no unit numbering conflict, the stack members retain their previous unit numbers. Otherwise, the stack master assigns new unit numbers, based on the order in which they come online.
- The new stack master uses its own startup and running configurations to synchronize the configurations on the new stack members.



Note: Adding a new unit that is powered on and has stack groups configured is the same as merging two stacks (refer to [Adding a Stack Unit](#)). If the new unit has been configured with a higher priority than the current stack master, it becomes the new stack master and the stack reloads. If the new unit does not have a higher priority than the master switch, it is added as a member switch.

Splitting a Stack

To split an MXL 10/40GbE Switch stack, unplug the stacking cables between member units at any time: while the stack is powered on or off and when the units are online or offline. Each portion of the split stack retains the startup and running configuration of the original stack.

For a stack that is split into two smaller stacks, each with multiple units:

- If one of the new stacks receives the master and standby units, it is unaffected by the split.
- If one of the new stacks receives only the master unit, the master switch retains its role and a new standby is elected.
- If one of the new stacks receives only the standby unit, it becomes the master in the new stack and FTOS elects a new standby.
- If one of the new stacks receives neither the master nor the standby unit, the stack is reset so that a new election takes place.

Managing Redundant Stack Management

To manage the redundancy behavior in a stack, use the following redundancy commands.

Task	Command Syntax	Command Mode
Reset the current stack master and make the standby unit the new master. A new standby is elected. When the previous stack master comes back online, it becomes a member unit.	redundancy force-failover stack-unit	EXEC Privilege
Prevent the stack master from rebooting after a failover. This command does not affect a forced failover, manual reset, or a stack-link disconnect.	redundancy disable-auto-reboot stack-unit	CONFIGURATION
Display redundancy information.	show redundancy	EXEC Privilege

Reset a Unit on a Stack

Use the following **reset** commands to reload any of the member units or the standby in a stack. If you try to reset the stack master, an error message is displayed: `Reset of master unit is not allowed.`

Task	Command Syntax	Command Mode
Reload a stack unit from the master switch	reset stack-unit <i>unit-number</i>	EXEC Privilege
Reload a member unit from the unit itself.	reset-self	EXEC Privilege
Reset a stack-unit when the unit is in a problem state.	reset stack-unit <i>unit-number</i> hard	EXEC Privilege

Verifying a Stack Configuration

Using LEDs

Table 31-1 lists the status of a stacked switch according to the color of the System Status light emitting diodes (LEDs) on its front panel.

Table 31-1. System Status LED on a Stacked Switch

Color	Meaning
Blue	The switch is operating as the stack master or as a standalone unit.
Off	The switch is a member or standby unit.
Amber	The switch is booting or a failure condition has occurred.

Using Show Commands

To display information on the stack configuration, use the show commands in Table 31-2 on the master switch.

Table 31-2. Displaying Stack Configurations

Command	Output
<code>show system [brief]</code> (Figure 31-8 and Figure 31-9)	Displays stacking roles (master, standby, and member units) and the stack MAC address.
<code>show inventory optional-module</code> (Figure 31-10)	Displays the FlexIO modules currently installed in expansion slots 0 and 1 on a switch and the expected module logically provisioned for the slot.
<code>show system stack-unit</code> <i>unit-number</i> <code>stack-group</code> configured (Figure 31-11)	Displays the stack groups allocated on a stacked switch. Valid stack-unit numbers: 0 to 5.
<code>show system stack-unit</code> <i>unit-number</i> <code>stack-group</code> (Figure 31-12)	Displays the port numbers that correspond to the stack groups on a switch. Valid stack-unit numbers: 0 to 5.

Table 31-2. Displaying Stack Configurations

Command	Output
<code>show system stack-ports [status topology]</code> (Figure 31-13)	Displays the type of stack topology (ring or daisy chain) with a list of all stacked ports, port status, link speed, and peer stack-unit connection.

Figure 31-8. show system brief Command Example

```
FTOS# show system brief

Stack MAC : 00:1e:c9:f1:00:7b

Reload Type : jump-start [Next boot : normal-reload]

-- Stack Info --
Unit  UnitType  Status  ReqTyp      CurTyp      Version     Ports
-----
  0   Management  online  MXL-10/40GbE  MXL-10/40GbE  9-1-0-853   56
  1   Standby     online  MXL-10/40GbE  MXL-10/40GbE  9-1-0-853   56
  2   Member      online  MXL-10/40GbE  MXL-10/40GbE  9-1-0-853   56
  3   Member      online  MXL-10/40GbE  MXL-10/40GbE  9-1-0-853   56
  4   Member      online  MXL-10/40GbE  MXL-10/40GbE  9-1-0-853   56
  5   Member      online  MXL-10/40GbE  MXL-10/40GbE  9-1-0-853   56
```

Figure 31-9. show system Command Example

```

FTOS#show system

Stack MAC : 00:1e:c9:f1:00:e3

Reload Type : normal-reload [Next boot : normal-reload]

-- Unit 0 --
Unit Type      : Member Unit
Status         : not present
Required Type  : MXL-10/40GbE - 34-port GE/TE/FG (XL)

-- Unit 1 --
Unit Type      : Management Unit
Status         : online
Next Boot      : online
Required Type  : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Current Type   : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Master priority : 14
Hardware Rev   : 2.0
Num Ports     : 56
Up Time       : 19 hr, 30 min
FTOS Version   : 9-1-0-1010
Jumbo Capable : yes
POE Capable   : no
Burned In MAC : 00:1e:c9:f1:00:e3
No Of MACs    : 3

-- Unit 2 --
Unit Type      : Member Unit
Status         : online
Next Boot      : online
Required Type  : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Current Type   : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Master priority : 12
Hardware Rev   : 2.0
Num Ports     : 56
Up Time       : 19 hr, 30 min
FTOS Version   : 9-1-0-1010
Jumbo Capable : yes
POE Capable   : no
Burned In MAC : 00:1e:c9:f1:00:c7
No Of MACs    : 3

-- Unit 3 --
Unit Type      : Member Unit
Status         : not present
Required Type  : MXL-10/40GbE - 34-port GE/TE/FG (XL)

-- Unit 4 --
Unit Type      : Standby Unit
Status         : online
Next Boot      : online
Required Type  : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Current Type   : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Master priority : 13
Hardware Rev   : 3.0
Num Ports     : 56
Up Time       : 19 hr, 30 min
FTOS Version   : 9-1-0-1010
Jumbo Capable : yes
POE Capable   : no

```

Figure 31-10. show inventory optional-module Command Example

```
FTOS# show inventory optional-module
```

Unit	Slot	Expected	Inserted	Next Boot	Power
0	0	SFP+	SFP+	AUTO	Good
0	1	QSFP+	QSFP+	AUTO	Good

* - Mismatch

Figure 31-11. show system stack-unit stack-group configured Command Example

```
FTOS# show system stack-unit 1 stack-group configured
Configured stack groups in stack-unit 1
-----
0
1
4
5
```

Figure 31-12. show system stack-unit stack-group Command Example

```
FTOS#show system stack-unit 1 stack-group
Stack group          Ports
-----
0                    0/33
1                    0/37
2                    0/41
3                    0/45
4                    0/49
5                    0/53
FTOS#
```

Figure 31-13. show system stack-ports (ring) Command Example

```
FTOS# show system stack-ports
Topology: Ring
```

Interface	Connection	Link Speed (Gb/s)	Admin Status	Link Status	Trunk Group
0/33	1/37	40	up	up	
0/37	2/33	40	up	up	
0/41	1/49	40	up	up	
0/45	2/53	40	up	up	
1/33	2/37	40	up	up	
1/37	0/33	40	up	up	
1/49	0/41	40	up	up	
1/53	2/49	40	up	up	
2/33	0/37	40	up	up	
2/37	1/33	40	up	up	
2/49	1/53	40	up	up	
2/53	0/45	40	up	up	

Figure 31-14. show system stack-ports (daisy chain) Command Example

```
FTOS# show system stack-ports
Topology: Daisy Chain
```

Interface	Connection	Link Speed (Gb/s)	Admin Status	Link Status	Trunk Group
0/33	1/37	40	up	up	
0/41	1/49	40	up	up	
1/33	2/37	40	up	up	
1/37	0/33	40	up	up	
1/49	0/41	40	up	up	
1/53	2/49	40	up	up	
2/37	1/33	40	up	up	
2/49	1/53	40	up	up	

Troubleshooting a Switch Stack

Troubleshooting Commands

To perform troubleshooting operations on a switch stack, use the commands in [Table 31-3](#) on the master switch.

Table 31-3. Troubleshooting Stack Commands

Command	Output
show system stack-ports (Figure 31-15)	Displays the status of stacked ports on stack units.
show redundancy (Figure 31-16)	Displays the master standby unit status, failover configuration, and result of the last master-standby synchronization; allows you to verify the readiness for a stack failover.
show hardware stack-unit <i>unit-number</i> stack-port <i>port-number</i> (Figure 31-15)	Displays input and output flow statistics on a stacked port.
clear hardware stack-unit <i>unit-number</i> counters	Clears statistics on the specified stack unit. Valid stack-unit numbers: 0 to 5.

Figure 31-15. show system stack-ports Command Example

```
FTOS# show system stack-ports
Topology: Ring

  Interface    Connection    Link Speed    Admin    Link    Trunk
              (Gb/s)        Status        Status   Group
-----
0/41          2/45          40            up       up
0/45          1/41          40            up       up
1/41          0/45          40            up       up
1/45          2/41          40            up       up
2/41          1/45          40            up       up
2/45          0/41          40            up       up
```

Figure 31-16. show redundancy Command Example

```

FTOS#show redundancy

-- Stack-unit Status --
-----
Mgmt ID:                0
Stack-unit ID:          0
Stack-unit Redundancy Role: Primary
Stack-unit State:       Active ← Indicates Master Unit.
Stack-unit SW Version:  E8-3-16-79
Link to Peer:           Up

-- PEER Stack-unit Status --
-----
Stack-unit State:       Standby ← Indicates Standby Unit
Peer stack-unit ID:    2
Stack-unit SW Version:  E8-3-16-79

-- Stack-unit Redundancy Configuration --
-----
Primary Stack-unit:    mgmt-id    0
Auto Data Sync:        Full
Failover Type:         Hot Failover ← Failover type with redundancy
force-failover
Auto reboot Stack-unit: Disabled
Auto failover limit:   3 times in 60 minutes

-- Stack-unit Failover Record --
-----
Failover Count:        0
Last failover timestamp: None
Last failover Reason:  None
Last failover type:    None
-- Last Data Block Sync Record: --
-----
Stack Unit Config:     succeeded Mar 24 2012 20:07:39
Start-up Config:       succeeded Mar 24 2012 20:07:39 ← Latest sync of config
Runtime Event Log:     succeeded Mar 24 2012 20:07:39
Running Config:        succeeded Mar 24 2012 20:07:39
ACL Mgr:               succeeded Mar 24 2012 20:07:39
LACP:                  no block sync done
STP:                   no block sync done

```


Figure 31-17. show hardware stack-unit stack-port Command Example

```
FTOS# show hardware stack-unit 1 stack-port 53

Input Statistics:
  7934 packets, 1049269 bytes
  0 64-byte pkts, 7793 over 64-byte pkts, 100 over 127-byte pkts
  0 over 255-byte pkts, 7 over 511-byte pkts, 34 over 1023-byte pkts
  70 Multicasts, 0 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  438 packets, 270449 bytes, 0 underruns
  0 64-byte pkts, 57 over 64-byte pkts, 181 over 127-byte pkts
  54 over 255-byte pkts, 0 over 511-byte pkts, 146 over 1023-byte pkts
  72 Multicasts, 0 Broadcasts, 221 Unicasts
  0 throttles, 0 discarded, 0 collisions, 0 wredDrops
Rate info (interval 45 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
```

Failure Scenarios

The following sections describe some of the common fault conditions that can happen in a switch stack and how they are resolved.

Stack Member Fails

Problem: A unit that is not the stack master fails in an operational stack.

Resolution: If a stack member fails in a daisy chain topology, a split stack occurs. If a member unit fails in a ring topology, traffic is re-routed over existing stack links.

The following syslog messages are generated when a member unit fails:

```
FTOS#May 31 01:46:17: %STKUNIT3-M:CP %IPC-2-STATUS: target stack unit 4 not responding

May 31 01:46:17: %STKUNIT3-M:CP %CHMGR-2-STACKUNIT_DOWN: Major alarm: Stack unit 4 down - IPC
timeout

FTOS#May 31 01:46:17: %STKUNIT3-M:CP %IFMGR-1-DEL_PORT: Removed port: Te 4/1-32,41-48, Fo 4/
49,53

FTOS#May 31 01:46:18: %STKUNIT5-S:CP %IFMGR-1-DEL_PORT: Removed port: Te 4/1-32,41-48, Fo 4/
49,53
```

Unplugged Stacking Cable

Problem: A stacking cable is unplugged from a member switch. The stack loses half of its bandwidth from the disconnected switch.

Resolution: Intra-stack traffic is re-routed on a another link using the redundant stacking port on the switch. A recalculation of control plane and data plane connections is performed.

Master Switch Fails

Problem: The master switch fails due to a hardware fault, software crash, or power loss.

Resolution: A failover procedure begins:

1. Keep-alive messages from the MXL 10/40GbE master switch time out after 60 seconds and the switch is removed from the stack.
2. The standby switch takes the master role. Data traffic on the new master switch is uninterrupted. Protocol traffic is managed by the control plane.
3. A member switch is elected as the new standby. Data traffic on the new standby is uninterrupted. The control plane prepares for operation in Warm Standby mode.

Stack-Link Flapping Error

Problem/Resolution: Stacked MXL 10/40GbE Switches monitor their own stack ports and disable any stack port that flaps five times within 10 seconds. If the stacking ports that flap are on the master or standby, KERN-2-INT error messages note the units ([Figure 31-18](#)).

To re-enable a downed stacking port, power cycle the stacked switch on which the port is installed.

Figure 31-18. Recovering from a Stack-Link Flapping Error

```
-----MANAGEMENT UNIT-----
Error: Stack Port 49 has flapped 5 times within 10 seconds.Shutting down this stack port now.
Error: Please check the stack cable/module and power-cycle the stack.
10:55:20: %STKUNIT1-M:CP %KERN-2-INT: Error: Stack Port 50 has flapped 5 times within 10
seconds.Shutting down this stack port now.
10:55:20: %STKUNIT1-M:CP %KERN-2-INT: Error: Please check the stack cable/module and
power-cycle the stack.
-----STANDBY UNIT-----
10:55:18: %STKUNIT1-M:CP %KERN-2-INT: Error: Stack Port 50 has flapped 5 times within 10
seconds.Shutting down this stack port now.
10:55:18: %STKUNIT1-M:CP %KERN-2-INT: Error: Please check the stack cable/module
and power-cycle the stack.
-----MEMBER 2-----
Error: Stack Port 51 has flapped 5 times within 10 seconds.Shutting down this stack port now.
Error: Please check the stack cable/module and power-cycle the stack.
```

Master Switch Recovers from Failure

Problem: The master switch recovers from a failure after a reboot and rejoins the stack:

- As a member unit if there is already a standby
- As a standby if there is no standby in the stack

Protocol and control plane recovery requires time before the switch is fully online.

Resolution: When the entire stack is reloaded, the recovered master switch becomes the master unit of the stack.

Stack Unit in Card-Problem State Due to Incorrect FTOS Version

Problem: A stack unit enters a Card-Problem state because the switch has a different FTOS version than the master unit (Figure 31-19). The switch does not come online as a stack unit.

Resolution: To restore a stack unit with an incorrect FTOS version as a member unit, disconnect the stacking cables on the switch and install the correct FTOS version. Then add the switch to the stack as described in [Adding a Stack Unit](#). To verify that the problem has been resolved and the stacked switch is back online, use the show system brief command (Figure 31-20).

Figure 31-19. Card Problem Error - Different FTOS Versions

```
FTOS#show system brief

Stack MAC : 00:1e:c9:f1:01:57

Reload Type : normal-reload [Next boot : normal-reload]

-- Stack Info --
Unit  UnitType  Status      ReqTyp      CurTyp      Version     Ports
-----
  0  Management  online      MXL-10/40GbE  MXL-10/40GbE  8-3-16-79  56
  1  Member      card problem MXL-10/40GbE  unknown      56
  2  Standby     online      MXL-10/40GbE  MXL-10/40GbE  8-3-16-79  56
  3  Member      not present
  4  Member      not present
  5  Member      not present
```

Figure 31-20. Card Problem Error - Different FTOS Versions: Resolved

```

FTOS#show system brief

Stack MAC : 00:1e:c9:f1:01:57

Reload Type : normal-reload [Next boot : normal-reload]

-- Stack Info --
Unit  UnitType  Status      ReqTyp      CurTyp      Version     Ports
-----
 0  Management  online      MXL-10/40GbE  MXL-10/40GbE  8-3-16-79  56
 1  Member      online     MXL-10/40GbE  MXL-10/40GbE  8-3-16-79  56
 2  Standby     online      MXL-10/40GbE  MXL-10/40GbE  8-3-16-79  56
 3  Member      not present
 4  Member      not present
 5  Member      not present

```

Stack Unit in Card-Problem State Due to Configuration Mismatch

Problem: A stack unit enters a Card-Problem state because there is a configuration mismatch between the logical provisioning stored for the stack-unit number on the master switch and the newly added unit with the same number.

Resolution: The resolution is to reload the stack. When the stack is up, the card problem will be solved

Step	Task	Command Syntax	Command Mode
1	From the Master Switch, reload the entire stack	reload	EXEC Privilege

Upgrading a Switch Stack

To upgrade all switches in a stack with the same FTOS version, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Copy the new FTOS image to a network server.		
2	Download the FTOS image by accessing an interactive CLI that requests the server IP address and image filename, and prompts you to upgrade all member stack units. Specify the system partition on the master switch into which you want to copy the FTOS image; valid values are a: and b: . As shown in Figure 31-21 , the system then prompts you to upgrade all member units with the new FTOS version.	<i>upgrade system { flash: ftp: scp: tftp: usbflash: } partition</i>	EXEC Privilege
3	Reboot all stack units to load the FTOS image from the same partition on all switches in the stack.	<i>boot system stack-unit all primary system partition</i>	CONFIGURATION
4	Save the configuration.	<i>write memory</i>	EXEC Privilege
5	Reload the stack unit to activate the new FTOS version.	<i>reload</i>	CONFIGURATION

[Figure 31-21](#) shows an example of how to upgrade all switches in a stack, including the master switch.

Figure 31-21. Upgrading all Stacked Switches Example

```

FTOS# upgrade system ftp: A:
Address or name of remote host []: 10.11.200.241
Source file name []: $V-9-1-0/NAVASOTA-DEV-9-1-0-887/FTOS-XL-9-1-0-887.bin
User name to login remote host: ftp
Password to login remote host:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Erasing IOM Primary Image, please wait
.!.....
.....Writing.....
.....
31972272 bytes successfully copied
System image upgrade completed successfully.
Upgrade system image for all stack-units [yes/no]: yes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Image upgraded to all
FTOS# configure
FTOS(conf)# boot system stack-unit all primary system: A:
FTOS(conf)# end
FTOS# write memory
Jan 3 14:01:48: %STKUNIT0-M:CP %FILEMGR-5-FILESAVED: Copied running-config to startup-config
in flash by default
Synchronizing data to peer Stack-unit
!!!!
FTOS# reload
Proceed with reload [confirm yes/no]: yes

```


Storm Control

Overview

The storm control feature allows you to control unknown-unicast and broadcast traffic on Layer 2, Layer 3, and multicast physical interfaces.



FTOS Behavior: The Dell Force10 operating software (FTOS) supports broadcast control (storm-control broadcast command) for Layer 2 *and* Layer 3 traffic.



FTOS Behavior: The minimum number of packets per second (PPS) that storm control can limit is two.

Configure Storm Control

Storm control is supported in INTERFACE mode and CONFIGURATION mode

Configure Storm Control from INTERFACE Mode

To configure storm control, use the storm control command from INTERFACE mode.

You can only configure storm control for ingress traffic in INTERFACE mode. If you configure storm control from both INTERFACE and CONFIGURATION mode, the INTERFACE mode configurations override the CONFIGURATION mode configurations.

Configure Storm Control from CONFIGURATION Mode

To configure storm control, use the storm control command from CONFIGURATION mode.

You can configure storm control for ingress traffic in CONFIGURATION mode. Do not apply per-virtual local area network (per-VLAN) quality of service (QoS) on an interface that you have enabled storm-control (either on an interface or globally)

Spanning Tree Protocol (STP)

Overview

The spanning tree protocol (STP) is a Layer 2 protocol—specified by IEEE 802.1d—that eliminates loops in a bridged topology by enabling only a single path through the network. By eliminating loops, the protocol improves scalability in a large network and allows you to implement redundant paths, which can be activated after the failure of active paths. Layer 2 loops, which can occur in a network due to poor network design and without enabling protocols like xSTP, can cause unnecessarily high switch CPU utilization and memory consumption.

Table 33-1 lists the variations of STP that FTOS supports.

Table 33-1. FTOS Supported Spanning Tree Protocols

Dell Force10 Term	IEEE Specification
Spanning Tree Protocol (STP)	802.1d
Rapid Spanning Tree Protocol (RSTP)	802.1w
Multiple Spanning Tree Protocol (MSTP)	802.1s
Per-VLAN Spanning Tree Plus (PVST+)	Third Party

This chapter contains the following sections:

- [Configuring Spanning Tree](#)
- [Configuring Interfaces for Layer 2 Mode](#)
- [Enabling Spanning Tree Protocol Globally](#)
- [Adding an Interface to the Spanning Tree Group](#)
- [Removing an Interface from the Spanning Tree Group](#)
- [Modifying Global Parameters](#)
- [Modifying Interface STP Parameters](#)
- [Enabling PortFast](#)
- [BPDU Filtering](#)
- [STP Root Selection](#)
- [STP Root Guard](#)
- [SNMP Traps for Root Elections and Topology Changes](#)
- [Displaying STP Guard Configuration](#)

Configuring Spanning Tree

Configuring STP is a two-step process:

1. Configure interfaces for Layer 2.
2. Enable STP.

Related Configuration Tasks

- [Adding an Interface to the Spanning Tree Group](#)
- [Removing an Interface from the Spanning Tree Group](#)
- [Modifying Global Parameters](#)
- [Modifying Interface STP Parameters](#)
- [Enabling PortFast](#)
- [Preventing Network Disruptions with BPDU Guard](#)
- [STP Root Selection](#)
- [SNMP Traps for Root Elections and Topology Changes](#)

Important Points to Remember

- STP is disabled by default.
- FTOS supports only one spanning tree instance (0). For multiple instances, you must enable MSTP or PVST+. You can only enable one flavor of STP at any one time.
- All ports in virtual local area networks (VLANs) and all enabled interfaces in Layer 2 mode are automatically added to the STP topology at the time you enable the protocol.
- After you enable STP, to add interfaces to the STP topology, enable the port and configure it for Layer 2, then use the switchport command.
- 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.

Configuring Interfaces for Layer 2 Mode


All interfaces on all switches that participate in STP must be in Layer 2 mode and enabled.

Figure 33-1. Example of Configuring Interfaces for Layer 2 Mode

```

FT05(conf)#interface range tengigabitethernet 1/1 - 4
FT05(conf-if-range-te-1/1-4)#switchport
FT05(conf-if-range-te-1/1-4)#no shutdown
FT05(conf-if-range-te-1/1-4)#show config
!
interface TenGigabitEthernet 1/1
no ip address
switchport
no shutdown
!
interface TenGigabitEthernet 1/2
no ip address
switchport
no shutdown
!
interface TenGigabitEthernet 1/3
no ip address
switchport
no shutdown
!
interface TenGigabitEthernet 1/4
no ip address
switchport
no shutdown
FT05(conf-if-range-te-1/1-4)#

```



To configure the interfaces for Layer 2 and then enable them, follow these steps:

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

To verify that an interface is in Layer 2 mode and enabled, use the show config command from INTERFACE mode (Figure 33-2).

Figure 33-2. show config Command Example

```
FTOS(conf-if-te-1/1)#show config
!
interface TenGigabitEthernet 1/1
  no ip address
  switchport ← Indicates that the interface is in Layer 2 mode
  no shutdown
FTOS(conf-if-te-1/1)#
```

Enabling Spanning Tree Protocol Globally

You must enable STP globally; it is not enabled by default.

To enable STP globally for all Layer 2 interfaces, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Enter 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, use the disable command from PROTOCOL SPANNING TREE mode.

To verify that STP is enabled, use the show config command from PROTOCOL SPANNING TREE mode (Figure 33-3).


Figure 33-3. Verifying STP is Enabled

```
FTOS(conf)#protocol spanning-tree 0
FTOS(conf-span)#show config
!
protocol spanning-tree 0
  no disable ← Indicates that Spanning Tree is enabled
FTOS#
```

When you enable STP, all physical, VLAN, and port-channel interfaces that are enabled and in Layer 2 mode are automatically part of the STP topology (Figure 33-4).

- 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.

Figure 33-4. Spanning Tree Enabled Globally



To view the STP configuration and the interfaces that are participating in STP, use 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 (Figure 33-5).

Figure 33-5. `show spanning-tree 0` Command Example

```
FTOS#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
Bpdu filter disabled globally
Current root has priority 32768, address 0001.e80d.2462
Root Port is 289 (Tengigabitethernet 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 Tengigabitethernet 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 (Tengigabitethernet 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 (Tengigabitethernet 2/2) is Blocking
Port path cost 4, Port priority 8, Port Identifier 8.290

FTOS#
```

To confirm that a port is participating in STP, use the `show spanning-tree 0 brief` command from EXEC privilege mode (Figure 33-6).

Figure 33-6. show spanning-tree brief Command Example

```
FTOS#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
Bpdu filter disabled globally

Interface
Name          PortID Prio Cost Sts Cost   Designated
-----
Tengig 1/1    8.496  8    4  DIS  0    32768 0001.e80d.2462 8.496
Tengig 1/2    8.497  8    4  DIS  0    32768 0001.e80d.2462 8.497
Tengig 1/3    8.513  8    4  FWD  0    32768 0001.e80d.2462 8.513
Tengig 1/4    8.514  8    4  FWD  0    32768 0001.e80d.2462 8.514
FTOS#
```

Adding an Interface to the Spanning Tree Group

To add a Layer 2 interface to the STP topology, use the following command:

Task	Command Syntax	Command Mode
Enable spanning tree on a Layer 2 interface.	<code>spanning-tree 0</code>	INTERFACE

Removing an Interface from the Spanning Tree Group

To remove a Layer 2 interface from the STP topology, use the following command:

Task	Command Syntax	Command Mode
Disable spanning tree on a Layer 2 interface.	<code>no spanning-tree 0</code>	INTERFACE

Modifying Global Parameters

You can modify the STP 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 STP.



Note: Dell Force10 recommends that only experienced network administrators change the STP parameters. Poorly planned modification of the STP parameters can negatively impact network performance.

Table 33-2. Default Values for STP

STP Parameter		Default Value
Forward Delay		15 seconds
Hello Time		2 seconds
Max Age		20 seconds
Port Cost	40-Gigabit Ethernet interfaces	1
	10-Gigabit Ethernet interfaces	2
	Port Channel with 40-Gigabit Ethernet interfaces	1
	Port Channel with 10-Gigabit Ethernet interfaces	1
Port Priority		8

To change the STP global parameters, use the following commands:

Task	Command Syntax	Command Mode
Change the forward-delay parameter (the wait time before the interface enters the Forwarding state). <ul style="list-style-type: none">• Range: 4 to 30• Default: 15 seconds	<code>forward-delay seconds</code>	PROTOCOL SPANNING TREE
Change the hello-time parameter (the BPDU transmission interval). Note: With large configurations (especially those with more ports) Dell Force10 recommends increasing the hello-time. Range: 1 to 10 Default: 2 seconds	<code>hello-time seconds</code>	PROTOCOL SPANNING TREE
Change the max-age parameter (the refresh interval for configuration information that is generated by recomputing the STP topology). Range: 6 to 40 Default: 20 seconds	<code>max-age seconds</code>	PROTOCOL SPANNING TREE

To view the current values for global parameters, use the `show spanning-tree 0` command from EXEC privilege mode (Figure 33-5).

Modifying Interface STP 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 is selected to be a forwarding port.
- **Port priority** influences the likelihood that a port is selected to be a forwarding port in case that several ports have the same port cost.

The default values are listed in [Table 33-2](#).

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 33-2 .	<code>spanning-tree 0 cost cost</code>	INTERFACE
Change the port priority of an interface. Range: 0 to 15 Default: 8	<code>spanning-tree 0 priority priority-value</code>	INTERFACE

To view the current values for interface parameters, use the `show spanning-tree 0` command from EXEC privilege mode ([Figure 33-5](#)).

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. The `bpduguard shutdown-on-violation` option causes the interface hardware to be shutdown when it receives a BPDU. When you only implement `bpduguard`, although the interface is placed in an Error Disabled state when receiving the BPDU, the physical interface remains up and STP drops packets in the hardware after a BPDU violation. BPDUs are dropped in the software after receiving the BPDU violation.



Caution: Enable PortFast only on links connecting to an end station. PortFast can cause loops if you enable it on an interface connected to a network.

BPDU Filtering enabled on an interface, stops sending and receiving BPDUs on the port fast enabled ports. When BPDU guard and BPDU filter is enabled on the port, then BPDU filter takes the highest precedence. By default bpduguard filtering on an interface is disabled.

To enable PortFast on an interface, use the following command:

Task	Command Syntax	Command Mode
Enable PortFast on an interface.	spanning-tree stp-id portfast [bpduguard [shutdown-on-violation] bpdufilter]	INTERFACE

To verify that PortFast is enabled on a port, use the `show spanning-tree` command from EXEC privilege mode or the `show config` command from INTERFACE mode. Dell Force10 recommends using the `show config` command (Figure 33-7).

Figure 33-7. show config Command Example (PortFast Enabled)

```
FTOS#(conf-if-te-1/1)#show conf
!
interface Tengigabitethernet 1/1
  no ip address
  switchport
  spanning-tree 0 portfast ← Indicates that the interface is in PortFast mode
  no shutdown
FTOS#(conf-if-te-1/1)#
```

Preventing Network Disruptions with BPDU Guard

You must configure the Portfast (and Edgeport, in the case of RSTP, PVST+, and MSTP) feature on ports that connect to end stations. End stations do not generate BPDUs, so ports configured with Portfast/Edgeport (edgeports) do not expect to receive BPDUs. If an edgeport does receive a BPDU, it likely means that it is connected to another part of the network, which can negatively effect the STP topology. The BPDU guard feature blocks an edgeport after receiving a BPDU to prevent network disruptions and FTOS displays [Message 1](#).

To enable BPDU guard, use the `bpduguard` option when enabling PortFast or EdgePort. The `bpduguard shutdown-on-violation` option causes the interface hardware to be shutdown when it receives a BPDU. Otherwise, although the interface is placed in an Error Disabled state when receiving the BPDU, the physical interface remains up and STP only drops *packets* after a BPDU violation.

[Figure 33-8](#) shows a scenario in which an edgeport might unintentionally receive a BPDU. The port on the Dell Force10 system is configured with Portfast. If the switch is connected to the hub, the BPDUs that the switch generates might trigger an undesirable topology change. If you enabled BPDU guard, when the edge port receives the BPDU, the BPDU is dropped, the port is blocked, and a console message is generated.

Message 1 BPDU Guard Error

```
4d23h40m: %STKUNIT3-M:CP %SPANMGR-5-BPDU_GUARD_RX_ERROR: Received Spanning Tree BPDU on BPDU guard port. Disable TenGigabitEthernet 3/20.
```



Note: Note that *unless* you enable the shutdown-on-violation option, STP only *drops packets* after a BPDU violation; the physical interface remains up, as shown below:

```
FTOS#sh spanning-tree 0 brief
  Executing IEEE compatible Spanning Tree Protocol
    Root ID Priority 32768, Address 0001.e88a.fdb3 Cost 1
    Root Port 2 (Port-channel 1)
    Root Bridge hello time 2, max age 20, forward delay 15
    Bridge ID Priority 32768, Address 001e.c9f1.00cf
    Configured hello time 2, max age 20, forward delay 15
    Bpdu filter disabled globally
```

Interface Name	PortID	Prio	Cost	Sts	Cost	Designated Bridge ID	PortID
Po 1	8.2	8	1	FWD	0	32768 0001.e88a.fdb3	8.2
Te 3/20	8.317	8	4	EDS	1	32768 001e.c9f1.00cf	8.317
Te 4/20	8.373	8	4	FWD	1	32768 001e.c9f1.00cf	8.373
Te 4/21	8.374	8	4	FWD	1	32768 001e.c9f1.00cf	8.374

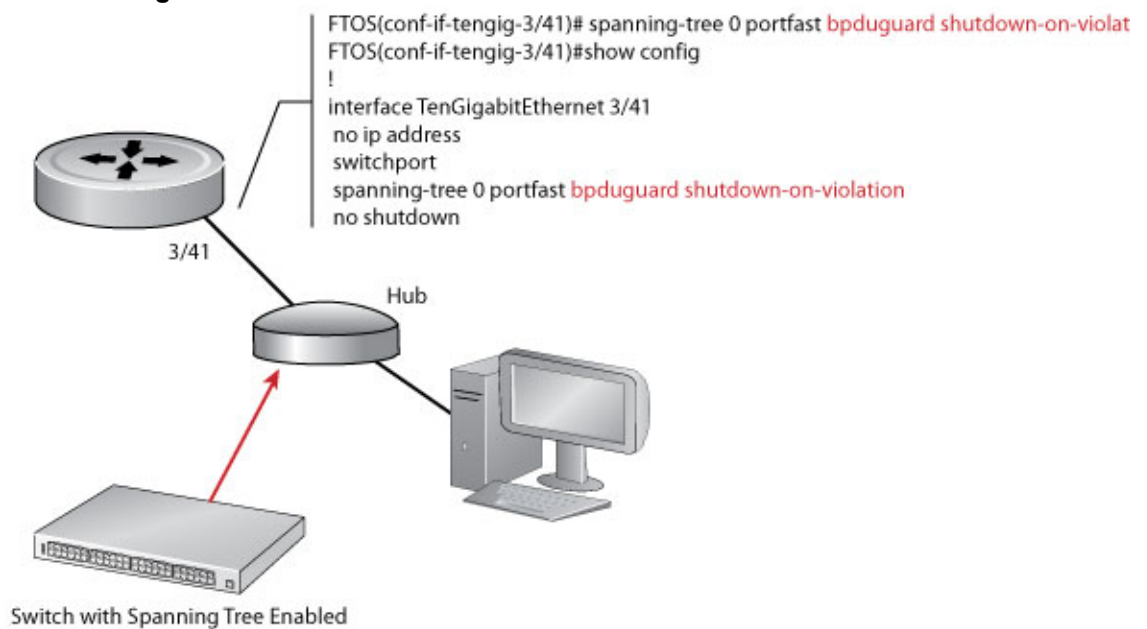
```
FTOS#sh ip int br ten 3/20
Interface          IP-Address      OK Method Status      Protocol
TenGigabitEthernet 3/20  unassigned     YES None  up          up
FTOS#
```



FTOS Behavior: Regarding bpduguard shutdown-on-violation behavior:

- 1 If the interface to be shutdown is a port channel, all the member ports are disabled in the hardware.
- 2 When a physical port is added to a port channel already in an Error Disable state, the new member port is also disabled in the hardware.
- 3 When a physical port is removed from a port channel in an Error Disable state, the Error Disabled state is cleared on this physical port (the physical port is enabled in the hardware).
- 4 You can clear the Error Disabled state with any of the following methods:
 - Use shutdown command on the interface.
 - Disable the shutdown-on-violation command on the interface (no spanning-tree *stp-id* portfast [bpduguard | [shutdown-on-violation]]).
 - Disable STP on the interface (no spanning-tree in INTERFACE mode).
 - Disabling global STP (no spanning-tree in CONFIGURATION mode).

Figure 33-8. Enabling BPDU Guard



FTOS Behavior: BPDU guard blocks BPDUs (refer to [Removing an Interface from the Spanning Tree Group](#)).


- BPDU guard is used on edgeports and blocks all traffic on edgeport if it receives a BPDU.

BPDU Filtering

Global BPDU Filtering

When BPDU Filtering is enabled globally, it should stop transmitting BPDUs on the operational port fast enabled ports by default. When it receives BPDUs, it automatically participates in the spanning tree. By default global bpdud filtering is disabled.

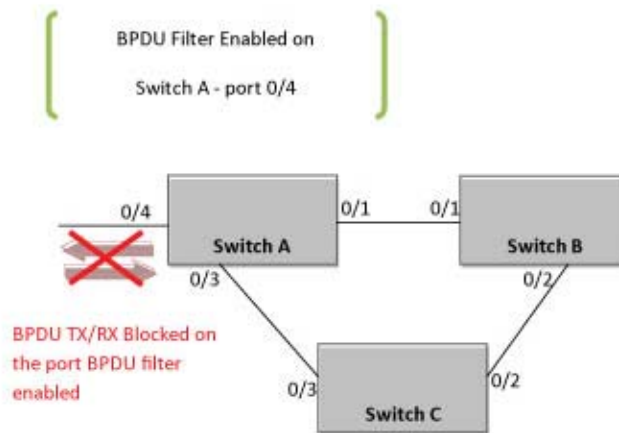
Figure 33-9. BPDU Filtering enabled globally



Interface BPDU Filtering

When BPDU Filtering is enabled on an interface, it should stop sending and receiving BPDUs on the port fast enabled ports. When BPDU guard and BPDU filter is enabled on the port, then BPDU filter takes the highest precedence. By default bpdu filtering on an interface is disabled.

Figure 33-10. BPDU Filtering enabled globally



STP Root Selection

STP determines the root bridge, but you can assign one bridge a lower priority to increase the likelihood that it is 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. <i>priority-value</i> range: 0 to 65535. The lower the number assigned, the more likely this bridge will become the root bridge. The default is 32768. <ul style="list-style-type: none">The primary option specifies a bridge priority of 8192.The secondary option specifies a bridge priority of 16384.	<code>bridge-priority {<i>priority-value</i> primary secondary}</code>	PROTOCOL SPANNING TREE

To view only the root information, use the `show spanning-tree root` command from EXEC privilege mode (Figure 33-11).

Figure 33-11. show spanning-tree root Command Example

```
FTOS#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
FTOS#
```

STP Root Guard

Use the STP root guard feature in a Layer 2 network to avoid bridging loops. In STP, the switch in the network with the lowest priority (as determined by STP or set with the `bridge-priority` command) is selected as the root bridge. If two switches have the same priority, the switch with the lower MAC address is selected as the root. All other switches in the network use the root bridge as the reference to calculate the shortest forwarding path.

Because any switch in an STP network with a lower priority can become the root bridge, the forwarding topology may not be stable. The location of the root bridge can change, resulting in unpredictable network behavior. The STP root guard feature ensures that the position of the root bridge does not change.

Root Guard Scenario

For example, in Figure 33-12 (STP topology 1 upper left), Switch A is the root bridge in the network core. Switch C functions as an access switch connected to an external device. The link between Switch C and Switch B is in a Blocking state. Figure 33-12 shows the flow of STP BPDUs.


In STP topology 2 (Figure 33-12 upper right), STP is enabled on device D on which a software bridge application is started to connect to the network. Because the priority of the bridge in device D is lower than the root bridge in Switch A, device D is elected as root, causing the link between Switches A and B to enter a Blocking state. Network traffic then begins to flow in the directions indicated by the BPDU arrows in the topology. If the links between Switches C and A or Switches C and B cannot handle the increased traffic flow, frames may be dropped.

In STP topology 3 (Figure 33-12 lower middle), if you enabled the root guard feature on the STP port on Switch C that connects to device D, and device D sends a superior BPDU that would trigger the election of device D as the new root bridge, the BPDU is ignored and the port on Switch C transitions from a forwarding to a Root-Inconsistent state (shown by the green X icon). As a result, Switch A becomes the root bridge.

All incoming and outgoing traffic is blocked on an STP port in a Root-Inconsistent state. After the timeout period, the Switch C port automatically transitions to a Forwarding state as soon as device D stops sending BPDUs that advertise a lower priority.

If you enable a root guard on all STP ports on the links where the root bridge should not appear, you can ensure a stable STP network topology and avoid bridging loops.

Figure 33-12. STP Root Guard Prevents Bridging Loops



Root Guard Configuration

You enable STP root guard on a per-port or per-port-channel basis.



FTOS Behavior: The following conditions apply to a port enabled with STP root guard:

- Root guard is supported on any STP-enabled port or port-channel interface except when used as a stacking port.
- Root guard is supported on a port in any Spanning Tree mode:
 - [Spanning Tree Protocol \(STP\) \(STP\)](#)
 - [Rapid Spanning Tree Protocol \(RSTP\)](#)
 - [Multiple Spanning Tree Protocol \(MSTP\)](#)
 - [Per-VLAN Spanning Tree Plus \(PVST+\)](#)
- When enabled on a port, root guard applies to all VLANs configured on the port.
- When used in an MSTP network, if root guard blocks a boundary port in the CIST, the port is also blocked in all other MST instances.

To enable the root guard on an STP-enabled port or port-channel interface in instance 0, use the `spanning-tree 0 rootguard` command.

Task	Command Syntax	Command Mode
Enable root guard on a port or port-channel interface.	<code>spanning-tree {0 mstp rstp pvst} rootguard</code>	INTERFACE
0: Enables root guard on an STP-enabled port assigned to instance 0.		INTERFACE
mstp: Enables root guard on an MSTP-enabled port.		PORT-CHANNEL
rstp: Enables root guard on an RSTP-enabled port.		
pvst: Enables root guard on a PVST-enabled port.		

To disable STP root guard on a port or port-channel interface, use the `no spanning-tree 0 rootguard` command in INTERFACE Configuration mode.

To verify the STP root guard configuration on a port or port-channel interface, use the `show spanning-tree 0 guard [interface interface]` command in Global Configuration mode.

SNMP Traps for Root Elections and Topology Changes

To enable SNMP traps for STP state changes, use the `snmp-server enable traps stp` command.

To enable SNMP traps for RSTP, MSTP, and PVST+ collectively, use the `snmp-server enable traps xstp` command.

Displaying STP Guard Configuration

To verify the STP guard configured on port or port-channel interfaces, use the `show spanning-tree 0 guard [interface interface]` command.

Figure 33-13 shows an example for an STP network (instance 0) in which:

- Root guard is enabled on a port that is in a Root-Inconsistent state.
- BPDU guard is enabled on a port that is shut down (Error Disabled state) after receiving a BPDU.
- Bpdu filter is disabled on the ports.

Figure 33-13. Displaying STP Guard Configuration

```
FTOS#show spanning-tree 0 guard

Interface
Name          Instance Sts          Guard type      Bpdu Filter
-----
Tengig 0/1    0          INCON(Root)    Rootguard       No
Tengig 0/2    0          LIS            Loopguard       No
Tengig 0/3    0          EDS (Shut)    Bpduguard       No
FTOS#
```


System Time and Date

You can set and maintain system times and dates through the network time protocol (NTP). You can also set them through the Dell Force10 operating software (FTOS) command line interfaces (CLIs) and hardware settings.

This chapter includes the following sections:

- [Network Time Protocol](#)
 - [Overview](#)
 - [Implementation Information](#)
 - [Configuring Network Time Protocol](#)
- [FTOS Time and Date](#)
 - [Configuring Time and Date Settings](#)
 - [Set Daylight Savings Time](#)

Network Time Protocol

Network time protocol (NTP) 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.

NTP is a fault-tolerant protocol that automatically selects the best of several available time sources to synchronize to. Multiple candidates can be combined to minimize the accumulated error. Temporarily or permanently insane time sources are detected and avoided.

Dell Force10 recommends configuring NTP for the most accurate time. In FTOS, you can configure other time sources (the hardware and software clocks).

NTP is designed to produce three products: clock offset, roundtrip delay, and dispersion, all of which are relative to a selected reference clock.

- **Clock offset** represents the amount to adjust the local clock to bring it into correspondence with the reference clock.
- **Roundtrip delay** provides the capability to launch a message to arrive at the reference clock at a specified time.
- **Dispersion** represents the maximum error of the local clock relative to the reference clock.

Because most host time servers synchronize using another peer time server, there are two components in each of these three products: those determined by the peer relative to the primary reference source of standard time and those measured by the host relative to the peer.

Each of these components are maintained separately in the protocol in order to facilitate error control and management of the subnet itself. They provide not only precision measurements of offset and delay, but also definitive maximum error bounds, so that the user interface can determine not only the time, but the quality of the time as well.

In what may be the most common client/server model, a client sends an NTP message to one or more servers and processes the replies as received. The server interchanges addresses and ports, overwrites certain fields in the message, recalculates the checksum and returns the message immediately. Information included in the NTP message allows the client to determine the server time with respect to local time and adjust the local clock accordingly. In addition, the message includes information to calculate the expected timekeeping accuracy and reliability, as well as select the best from possibly several servers.

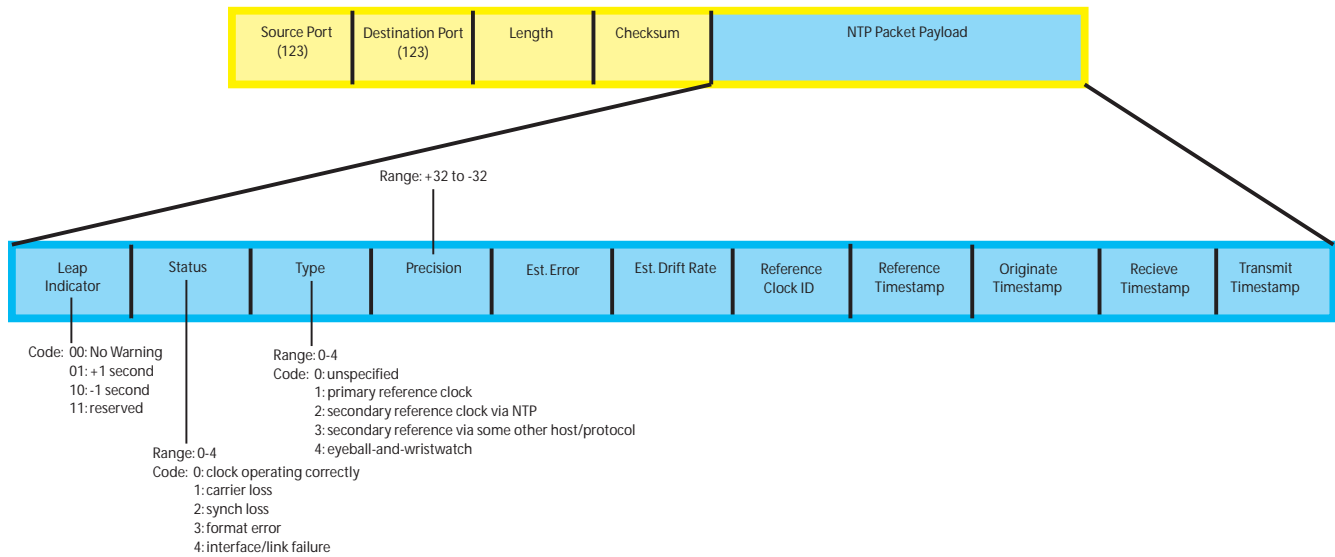
Following conventions established by the telephone industry [BEL86], the accuracy of each server is defined by a number called the stratum, with the topmost level (primary servers) assigned as one and each level downwards (secondary servers) in the hierarchy assigned as one greater than the preceding level.

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.

Overview

NTP message to one or more servers and processes the replies as received. The server interchanges addresses and ports, fills in or overwrites certain fields in the message, recalculates the checksum and returns it immediately (Table 34-1). Information included in the NTP message allows each client/server peer to determine the timekeeping characteristics of its other peers, including the expected accuracies of their clocks. Using this information, each peer is able to select the best time from possibly several other clocks, update the local clock, and estimate its accuracy.

Figure 34-1. NTP Fields



Implementation Information

The MXL Switch can only be an NTP client.

Configuring Network Time Protocol

Configuring NTP is a one-step process:

1. Enable NTP.

Related Configuration Tasks

- [Configure NTP Broadcasts](#)
- [Set the Hardware Clock with the Time Derived from NTP](#)
- [Configure NTP Broadcasts](#)
- [Disable NTP on an Interface](#)
- [Configure a Source IP Address for NTP Packets \(optional\)](#)

Enable NTP

NTP is disabled by default. To enable it, specify an NTP server to which the Dell Force10 system will synchronize. Enter the command multiple times to specify multiple servers. You may specify an unlimited number of servers at the expense of CPU resources.

To specify an NTP server, use the following command.

Task	Command	Command Mode
Specify the NTP server to which the Dell Force10 system will synchronize.	<code>ntp server ip-address</code>	CONFIGURATION

To display the system clock state with respect to NTP, use the `show ntp status` command from EXEC Privilege mode (Figure 34-2).

Figure 34-2. show ntp status Command Example (with respect to NTP)

```
FTOS(conf)#do show ntp status
Clock is synchronized, stratum 2, reference is 192.168.1.1
frequency is -369.623 ppm, stability is 53.319 ppm, precision is 4294967279
reference time is CD63BCC2.0CBBD000 (16:54:26.049 UTC Thu Mar 12 2012)
clock offset is 997.529984 msec, root delay is 0.00098 sec
root dispersion is 10.04271 sec, peer dispersion is 10032.715 msec
peer mode is client
```

To display the calculated NTP synchronization variables received from the server that the system will use to synchronize its clock, use the `show ntp associations` command from EXEC Privilege mode (Figure 34-3).

Figure 34-3. show ntp associations Command Example

```
FTOS(conf)#do show ntp associations
  remote      ref clock      st when poll reach  delay  offset  disp
-----
#192.168.1.1  .LOCL.          1  16  16  76    0.98  -2.470  879.23
* master (syncd), # master (unsyncd), + selected, - candidate
```

Set the Hardware Clock with the Time Derived from NTP

To set the hardware clock with the time value derived from NTP, use the following command:

Task	Command	Command Mode
Periodically update the system hardware clock with the time value derived from NTP (Figure 34-4).	<code>ntp update-calendar</code>	CONFIGURATION

Figure 34-4. Displaying the Calculated NTP Synchronization Variables

```
FTOS(conf)#do show calendar
11:08:48 UTC Tue May 22 2012
FTOS(conf)#ntp update-calendar 1
FTOS(conf)#do show calendar
```


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 command in INTERFACE mode:

Task	Command	Command
Set the interface to receive NTP packets.	ntp broadcast client	INTERFACE

Table 34-1. ntp broadcast client Command Example

```
2w1d11h : NTP: Maximum Slew:-0.000470, Remainder = -0.496884
```

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 INTERFACE mode:

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

To view whether NTP is configured on the interface, use the show config command in INTERFACE mode.

If ntp disable is not listed in the show config command output, NTP is enabled. (The show config command displays only non-default 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 CONFIGURATION mode:

Command Syntax	Command Mode	Purpose
ntp source <i>interface</i>	CONFIGURATION	Enter the following keywords and slot/port or number information: <ul style="list-style-type: none"> 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 128. 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. For a 40-Gigabit Ethernet interface, enter the keyword fortyGigE followed by the slot/port information.

To view the configuration, use the show running-config ntp command in EXEC privilege 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, follow these steps.

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

To view the NTP configuration, use the show running-config ntp command in EXEC privilege mode (Figure 34-5).

Figure 34-5 shows an encrypted authentication key. All keys are encrypted.

Figure 34-5. show running-config ntp Command Example

```

FTOS#show running ntp
!
ntp authenticate
ntp authentication-key 345 md5 5A60910F3D211F02 ← encrypted key
ntp server 11.1.1.1 version 3
ntp trusted-key 345
FTOS#

```

Command Syntax	Command Mode	Purpose
ntp server <i>ip-address</i> [<i>key keyid</i>] [<i>prefer</i>] [<i>version number</i>]	CONFIGURATION	Configure an NTP server. Configure the IP address of a server and the following optional parameters: <ul style="list-style-type: none"> key <i>keyid</i>: 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 <i>number</i>: Enter a number 1 to 4 as the NTP version.

```

FTOS(conf)#1w6d23h : NTP: xmit packet to 192.168.1.1:
 leap 0, mode 3, version 3, stratum 2, ppoll 1024
 rtdel 0219 (8.193970), rtdsp AF928 (10973.266602), refid C0A80101 (192.168.1.1)
 ref CD7F4F63.6BE8F000 (14:51:15.421 UTC Thu Apr 2 2012)
 org CD7F4F63.68000000 (14:51:15.406 UTC Thu Apr 2 2012)
 rec CD7F4F63.6BE8F000 (14:51:15.421 UTC Thu Apr 2 2012)
 xmt CD7F5368.D0535000 (15:8:24.813 UTC Thu Apr 2 2012)
1w6d23h : NTP: rcv packet from 192.168.1.1
 leap 0, mode 4, version 3, stratum 1, ppoll 1024
 rtdel 0000 (0.000000), rtdsp AF587 (10959.090820), refid 4C4F434C (76.79.67.76)
 ref CD7E14FD.43F7CED9 (16:29:49.265 UTC Wed Apr 1 2012)
 org CD7F5368.D0535000 (15:8:24.813 UTC Thu Apr 2 2012)
 rec CD7F5368.D0000000 (15:8:24.812 UTC Thu Apr 2 2012)
 xmt CD7F5368.D0000000 (15:8:24.812 UTC Thu Apr 2 2012)
 inp CD7F5368.D1974000 (15:8:24.818 UTC Thu Apr 2 2012)

rtdel-root delay
rtdsp - round trip dispersion
refid - reference id
org -
rec - (last?) receive timestamp
xmt - transmit timestamp

mode - 3 client, 4 server
stratum - 1 primary reference clock, 2 secondary reference clock (via NTP)
version - NTP version 3
leap -

```

- Leap Indicator (sys.leap, peer.leap, pkt.leap): This is a two-bit code warning of an impending leap second to be inserted in the NTP time scale. The bits are set before 23:59 on the day of insertion and reset after 00:00 on the following day. This causes the number of seconds (rollover interval) in the day of insertion to be increased or decreased by one. In the case of primary servers, the bits are set by operator intervention, while in the case of secondary servers, the bits are set by the protocol. The two bits, bit 0 and bit 1, respectively, are coded as follows:
 - Poll Interval: integer indicating the minimum interval between transmitted messages, in seconds as a power of two. For instance, a value of six indicates a minimum interval of 64 seconds.
 - Precision: integer indicating the precision of the various clocks, in seconds to the nearest power of two. The value must be rounded to the next larger power of two; for instance, a 50-Hz (20 ms) or 60-Hz (16.67ms) power-frequency clock would be assigned the value -5 (31.25 ms), while a 1000-Hz (1 ms) crystal-controlled clock would be assigned the value -9 (1.95 ms).
 - Root Delay (sys.rootdelay, peer.rootdelay, pkt.rootdelay): This is a signed fixed-point number indicating the total roundtrip delay to the primary reference source at the root of the synchronization subnet, in seconds. Note that this variable can take on both positive and negative values, depending on clock precision and skew.
 - Root Dispersion (sys.rootdispersion, peer.rootdispersion, pkt.rootdispersion): This is a signed fixed-point number indicating the maximum error relative to the primary reference source at the root of the synchronization subnet, in seconds. Only positive values greater than zero are possible.
 - Reference Clock Identifier (sys.refid, peer.refid, pkt.refid): This is a 32-bit code identifying the particular reference clock. In the case of stratum 0 (unspecified) or stratum 1 (primary reference source), this is a four-octet, left-justified, zero-padded ASCII string, for example, the case of stratum 2 and greater (secondary reference) this is the four-octet Internet address of the peer selected for synchronization.
 - Reference Timestamp (sys.reftime, peer.reftime, pkt.reftime): This is the local time, in timestamp format, when the local clock was last updated. If the local clock has never been synchronized, the value is zero.
 - Originate Timestamp: The departure time on the server of its last NTP message. If the server becomes unreachable, the value is set to zero.
 - Receive Timestamp: The arrival time on the client of the last NTP message from the server. If the server becomes unreachable, the value is set to zero.
 - Transmit Timestamp: The departure time on the server of the current NTP message from the sender.
 - Filter dispersion: The error in calculating the minimum delay from a set of sample data from a peer.

FTOS Time and Date

You can set the time and date using the FTOS CLI.

Configuring Time and Date Settings

The following list includes the configuration tasks for setting the system time:

- [Set the Time and Date for the Switch Hardware Clock](#)
- [Set the Time and Date for the Switch Software Clock](#)
- [Set the Timezone](#)
- [Set Daylight Savings Time](#)
 - [Set Daylight Saving Time Once](#)
 - [Set Recurring Daylight Saving Time](#)

Set the Time and Date for the Switch Hardware Clock

To set the time and date for the hardware clock, use the following command:

Command Syntax	Command Mode	Purpose
<code>calendar set <i>time month day year</i></code>	EXEC Privilege	Set the hardware clock to the current time and date. <ul style="list-style-type: none">• <i>time</i>: Enter the time in hours:minutes:seconds. For the hour variable, use the 24-hour format, for example, 17:15:00 is 5:15 pm.• <i>month</i>: Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to <i>time day month year</i>.• <i>day</i>: Enter the number of the day. Range: 1 to 31. You can enter the name of a month to change the order of the display to <i>time day month year</i>.• <i>year</i>: Enter a four-digit number as the year. Range: 1993 to 2035.

```
FTOS#calendar set 12:11:00 21 may 2012
FTOS#
```

Set the Time and Date for the Switch Software Clock

You can change the order of the *month* and *day* parameters to enter the time and date as *time day month year*. You cannot delete the software clock.

The software clock runs only when the software is up. The clock restarts, based on the hardware clock, when the switch reboots. To set the time and date for the software clock, use the following command:

Command Syntax	Command Mode	Purpose
clock set <i>time month day year</i>	EXEC Privilege	<p>Set the system software clock to the current time and date.</p> <ul style="list-style-type: none"> <i>time</i>: Enter the time in hours:minutes:seconds. For the hour variable, use the 24-hour format, for example, 17:15:00 is 5:15 pm. <i>month</i>: Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to <i>time day month year</i>. <i>day</i>: Enter the number of the day. Range: 1 to 31. You can enter the name of a month to change the order of the display to <i>time day month year</i>. <i>year</i>: Enter a four-digit number as the year. Range: 1993 to 2035.

```
FTOS#clock set 12:11:00 21 may 2012
FTOS#
```

Set the Timezone

Universal time coordinated (UTC) is the time standard based on the International Atomic Time standard, commonly known as Greenwich Mean time. When determining system time, you must include the differentiator between the UTC and your local timezone. For example, San Jose, CA is the Pacific Timezone with a UTC offset of -8.

To set the timezone, use the following command:

Command Syntax	Command Mode	Purpose
clock timezone <i>timezone-name offset</i>	CONFIGURATION	<p>Set the clock to the appropriate timezone.</p> <p><i>timezone-name</i>: Enter the name of the timezone. Do not use spaces.</p> <p><i>offset</i>: Enter one of the following:</p> <ul style="list-style-type: none"> a number from 1 to 23 as the number of hours in addition to UTC for the timezone. a minus sign (-) followed by a number from 1 to 23 as the number of hours.

Command Syntax	Command Mode	Purpose
<pre>FTOS#conf FTOS(conf)#clock timezone Pacific -8 FTOS#</pre>		

Set Daylight Savings Time

FTOS supports setting the system to daylight savings time once or on a recurring basis every year.

Set Daylight Saving Time Once

Set a date (and time zone) on which to convert the switch to daylight savings time on a one-time basis.

To set daylight saving time once, use the following command:

Command Syntax	Command Mode	Purpose
<pre>clock summer-time <i>time-zone date</i> <i>start-month start-day start-year</i> <i>start-time end-month end-day</i> <i>end-year end-time [offset]</i></pre>	CONFIGURATION	<p>Set the clock to the appropriate timezone and daylight savings time.</p> <ul style="list-style-type: none"> • <i>time-zone</i>: Enter the three-letter name for the time zone. This name is displayed in the show clock output. • <i>start-month</i>: Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to <i>time day month year</i> • <i>start-day</i>: Enter the number of the day. Range: 1 to 31. You can enter the name of a month to change the order of the display to <i>time day month year</i>. • <i>start-year</i>: Enter a four-digit number as the year. Range: 1993 to 2035 • <i>start-time</i>: Enter the time in hours:minutes. For the hour variable, use the 24-hour format, example, 17:15 is 5:15 pm. • <i>end-month</i>: Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to <i>time day month year</i>. • <i>end-day</i>: Enter the number of the day. Range: 1 to 31. You can enter the name of a month to change the order of the display to <i>time day month year</i>. • <i>end-year</i>: Enter a four-digit number as the year. Range: 1993 to 2035. • <i>end-time</i>: Enter the time in hours:minutes. For the hour variable, use the 24-hour format, example, 17:15 is 5:15 pm. • <i>offset</i>: (OPTIONAL) Enter the number of minutes to add during the summer-time period. Range: 1 to 1440. Default: 60 minutes

Command Syntax**Command Mode****Purpose**

```
FTOS(conf)#clock summer-time pacific date Mar 14 2012 00:00 Nov 7 2012 00:00
```

```
FTOS(conf)#
```

Set Recurring Daylight Saving Time

Set a date (and time zone) on which to convert the switch to daylight savings time on a specific day every year.

If you have already set daylight savings for a one-time setting, you can set that date and time as the recurring setting using the `clock summer-time time-zone recurring` command.

To set a recurring daylight saving time, use the following command:

Command Syntax**Command Mode****Purpose**

```
clock summer-time time-zone
recurring start-week start-day
start-month start-time end-week
end-day end-month end-time [offset]
```

CONFIGURATION

Set the clock to the appropriate timezone and adjust to daylight savings time every year.

- **time-zone:** Enter the three-letter name for the time zone. This name is displayed in the `show clock` output.
- **start-week:** (OPTIONAL) Enter one of the following as the week that daylight savings begins and then enter values for **start-day** through **end-time**:
- **week-number:** Enter a number from 1-4 as the number of the week in the month to start daylight savings time.
- **first:** Enter this keyword to start daylight savings time in the first week of the month.
- **last:** Enter this keyword to start daylight savings time in the last week of the month.
- **start-month:** Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to *time day month year*.
- **start-day:** Enter the number of the day. Range: 1 to 31. You can enter the name of the month to change the order of the display to *time day month year*.
- **start-year:** Enter a four-digit number as the year. Range: 1993 to 2035
- **start-time:** Enter the time in hours:minutes. For the hour variable, use the 24-hour format, example, 17:15 is 5:15 pm.

Command Syntax	Command Mode	Purpose
		<ul style="list-style-type: none"> • <i>end-week</i>: If you entered a start-week, enter one of the following as the week that daylight savings ends: • <i>week-number</i>: enter a number from 1 to 4 as the number of the week to end daylight savings time. • <i>first</i>: enter the keyword first to end daylight savings time in the first week of the month. • <i>last</i>: enter the keyword last to end daylight savings time in the last week of the month. • <i>end-month</i>: Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to <i>time day month year</i>. • <i>end-day</i>: Enter the number of the day. Range: 1 to 31. You can enter the name of a day to change the order of the display to <i>time day month year</i>. • <i>end-year</i>: Enter a four-digit number as the year. Range: 1993 to 2035. • <i>end-time</i>: Enter the time in hours:minutes. For the hour variable, use the 24-hour format, example, 17:15 is 5:15 pm. • <i>offset</i>: (OPTIONAL) Enter the number of minutes to add during the summer-time period. Range: 1 to 1440. Default: 60 minutes

```
FTOS(conf)#clock summer-time pacific recurring Mar 14 2012 00:00 Nov 7 2012 00:00
```

```
FTOS(conf)#
```

Note: If you enter <CR> after entering the recurring command parameter, and you have already set a one-time daylight saving time/date, the system uses that time and date as the recurring setting.

```
FTOS(conf)#clock summer-time pacific recurring ?
```

```
<1-4>           Week number to start
```

```
first           Week number to start
```

```
last            Week number to start
```

```
<cr>
```

```
FTOS(conf)#clock summer-time pacific recurring
```

```
FTOS(conf)#
```


Uplink Failure Detection (UFD)

Feature Description


Uplink Failure Detection (UFD) provides detection of the loss of upstream connectivity and, if used with NIC teaming, automatic recovery from a failed link.

A switch provides upstream connectivity for devices, such as servers. If a switch loses its upstream connectivity, downstream devices also lose their connectivity. However, the devices do not receive a direct indication that upstream connectivity is lost since connectivity to the switch is still operational.

UFD allows a switch to associate downstream interfaces with upstream interfaces. When upstream connectivity fails, the switch disables the downstream links. Failures on the downstream links allow downstream devices to recognize the loss of upstream connectivity.

For example, in [Figure 35-1](#) Switches S1 and S2 both have upstream connectivity to Router R1 and downstream connectivity to the server. UFD operation is shown in Steps A through C:

- In Step A, the server configuration uses the connection to S1 as the primary path. Network traffic flows from the server to S1 and then upstream to R1.
- In Step B, the upstream link between S1 and R1 fails. The server continues to use the link to S1 for its network traffic, but the traffic is not successfully switched through S1 because the upstream link is down.
- In Step C, UFD on S1 disables the link to the server. The server then stops using the link to S1 and switches to using its link to S2 to send traffic upstream to R1.


Figure 35-1. Uplink Failure Detection

How Uplink Failure Detection Works

UFD creates an association between upstream and downstream interfaces. The association of uplink and downlink interfaces is called an *uplink-state group*. An interface in an uplink-state group can be a physical interface or a port-channel (LAG) aggregation of physical interfaces.

An enabled uplink-state group tracks the state of all assigned upstream interfaces. Failure on an upstream interface results in the automatic disabling of downstream interfaces in the uplink-state group. As a result, downstream devices can execute the protection or recovery procedures they have in place to establish alternate connectivity paths as shown in [Figure 35-2](#).

Figure 35-2. Uplink Failure Detection Example



If only one of the upstream interfaces in an uplink-state group goes down, a specified number of downstream ports associated with the upstream interface are put into a link-down state. This number is user-configurable and is calculated by the ratio of upstream port bandwidth to downstream port bandwidth in the same uplink-state group. This calculation ensures that there are no traffic drops due to insufficient bandwidth on the upstream links to the routers/switches.

By default, if all upstream interfaces in an uplink-state group go down, all downstream interfaces in the same uplink-state group are put into a link-down state.

Using UFD, you can configure the automatic recovery of downstream ports in an uplink-state group when the link status of an upstream port changes. The tracking of upstream link status does not have a major impact on CPU usage.

UFD and NIC Teaming

Uplink Failure Detection on a switch can be used with network adapter teaming on a server (see [Network Interface Controller \(NIC\) Teaming on page 309](#)) to implement a rapid failover solution. For example, in [Figure 35-2](#) the switch/router with UFD detects the uplink failure and automatically disables the associated downstream link port to the server. The server with NIC teaming detects the disabled link and automatically switches over to the backup link in order to continue to transmit traffic upstream.

Important Points to Remember

When you configure Uplink Failure Detection, the following conditions apply:

- You can configure up to sixteen uplink-state groups. By default, no uplink-state groups are created. An uplink-state group is considered to be operationally *up* if it has at least one upstream interface in the link-up state.
An uplink-state group is considered to be operationally *down* if it has no upstream interfaces in the link-up state. No uplink-state tracking is performed when a group is disabled or in an operationally down state.
- You can assign physical port or port-channel interfaces to an uplink-state group.
You can assign an interface to only one uplink-state group. Each interface assigned to an uplink-state group must be configured as either an upstream or downstream interface, but not both.
You can assign individual member ports of a port channel to the group. An uplink-state group can contain either the member ports of a port channel or the port channel itself, but not both.
If you assign a port channel as an upstream interface, the port channel interface enters a link-down state when the number of port-channel member interfaces in a link-up state drops below the configured Minimum Number of Members parameter.
- If one of the upstream interfaces in an uplink-state group goes down, either a user-configurable set of downstream ports or all the downstream ports in the group are put in an operationally down state with an UFD Disabled error. The order in which downstream ports are disabled is from the lowest numbered port to the highest.
If one of the upstream interfaces in an uplink-state group that was down comes up, the set of UFD-disabled downstream ports (which were previously disabled due to this upstream port going down) are brought up and the UFD Disabled error is cleared.
- If an uplink-state group is disabled, the downstream interfaces are not disabled regardless of the state of the upstream interfaces.
If an uplink-state group has no upstream interfaces assigned, downstream interfaces will not be disabled.
- To enable the debug messages for events related to a specified uplink-state group or all groups, enter the **debug uplink-state-group** [*group-id*] command, where *group-id* is 1 to 16.
To turn off debugging event messages, enter the **no debug uplink-state-group** [*group-id*] command.
For an example of debug log messages, see [Message 1](#).

Configuring Uplink Failure Detection

To configure Uplink Failure Detection, follow these steps:

Step	Command Syntax and Mode	Description
1	uplink-state-group <i>group-id</i> Command Mode: CONFIGURATION	Creates an uplink-state group and enabling the tracking of upstream links on the switch/router. Valid <i>group-id</i> values are 1 to 16. To delete an uplink-state group, enter the no uplink-state-group <i>group-id</i> command.
2	{upstream downstream} <i>interface</i> Command Mode: UPLINK-STATE-GROUP	Assigns a port or port-channel to the uplink-state group as an upstream or downstream interface. For <i>interface</i> , enter one of the following interface types: 10-Gigabit Ethernet: tengigabitethernet { <i>slot/port</i> <i>slot/port-range</i> } 40-Gigabit Ethernet: fortygigabitethernet { <i>slot/port</i> <i>slot/port-range</i> } Port channel: port-channel {1-512 <i>port-channel-range</i> } Where <i>port-range</i> and <i>port-channel-range</i> specify a range of ports separated by a dash (-) and/or individual ports/port channels in any order; for example: upstream tengigabitethernet 1/1-2,5,9,11-12 downstream port-channel 1-3,5 A comma is required to separate each port and port-range entry. To delete an interface from the group, enter the no {upstream downstream} interface command.
3	downstream disable links {number all} Command Mode: UPLINK-STATE-GROUP	(Optional) Configures the number of downstream links in the uplink-state group that will be disabled (Oper Down state) if one upstream link in the group goes down. <i>number</i> specifies the number of downstream links to be brought down. Range: 1 to 1024. all brings down all downstream links in the group. Default: No downstream links are disabled when an upstream link goes down. To revert to the default setting, enter the no downstream disable links command.
4	downstream auto-recover Command Mode: UPLINK-STATE-GROUP	(Optional) Enables auto-recovery so that UFD-disabled downstream ports in the uplink-state group come up when a disabled upstream port in the group comes back up. Default: Auto-recovery of UFD-disabled downstream ports is enabled. To disable auto-recovery, enter the no downstream auto-recover command.

Step	Command Syntax and Mode	Description
5	description <i>text</i> Command Mode: UPLINK-STATE-GROUP	(Optional) Enters a text description of the uplink-state group. Maximum length: 80 alphanumeric characters.
6	no enable Command Mode: UPLINK-STATE-GROUP	(Optional) Disables upstream-link tracking without deleting the uplink-state group. Default: Upstream-link tracking is automatically enabled in an uplink-state group. To re-enable upstream-link tracking, enter the enable command.

Clearing a UFD-Disabled Interface

You can manually bring up a downstream interface in an uplink-state group that has been disabled by UFD and is in a UFD-disabled error state. To re-enable one or more disabled downstream interfaces and clear the UFD-disabled error state, enter the following command:

Command Syntax	Description
clear ufd-disable { interface <i>interface</i> uplink-state-group <i>group-id</i> } Command Mode: EXEC mode	Re-enables a downstream interface on the switch/router that is in a UFD-disabled error state so that it can send and receive traffic. For <i>interface</i> , enter one of the following interface types: 10-Gigabit Ethernet: tengigabitethernet { <i>slot/port</i> <i>slot/port-range</i> } 40-Gigabit Ethernet: fortygigabitethernet { <i>slot/port</i> <i>slot/port-range</i> } Port channel: port-channel {1-512 <i>port-channel-range</i> } Where <i>port-range</i> and <i>port-channel-range</i> specify a range of ports separated by a dash (-) and/or individual ports/port channels in any order; for example: tengigabitethernet 1/1-2,5,9,11-12 port-channel 1-3,5 A comma is required to separate each port and port-range entry. uplink-state-group <i>group-id</i> re-enables all UFD-disabled downstream interfaces in the group. Valid values are 1 to 16.

Message 1 shows the Syslog messages displayed when you clear the UFD-disabled state from all disabled downstream interfaces in an uplink-state group by entering the **clear ufd-disable uplink-state-group group-id** command. All downstream interfaces return to an operationally up state.

Message 1 Syslog Messages before and after entering clear ufd-disable uplink-state-group Command

```
00:10:12: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN: Changed interface Admin state to down: Te 0/1
00:10:12: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN: Changed interface Admin state to down: Te 0/2
00:10:12: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN: Changed interface Admin state to down: Te 0/3
00:10:12: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Te 0/1
00:10:12: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Te 0/2
00:10:12: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Te 0/3
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed uplink state group state to down: Group
3
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Downstream interface set to UFD error-disabled:
Te 0/4
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Downstream interface set to UFD error-disabled:
Te 0/5
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Downstream interface set to UFD error-disabled:
Te 0/6
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Te 0/4
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Te 0/5
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Te 0/6

FTOS(conf-if-range-te-0/1-3)#do clear ufd-disable uplink-state-group 3

00:11:50: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Downstream interface cleared from UFD
error-disabled: Te 0/4
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Downstream interface cleared from UFD
error-disabled: Te 0/5
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Downstream interface cleared from UFD
error-disabled: Te 0/6
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Changed interface state to up: Te 0/4
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Changed interface state to up: Te 0/5
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Changed interface state to up: Te 0/6
```

Displaying Uplink Failure Detection

To display information on the Uplink Failure Detection feature, enter any of the following **show** commands:

Show Command Syntax	Description
show uplink-state-group [<i>group-id</i>] [detail] Command Mode: EXEC	Displays status information on a specified uplink-state group or all groups. Valid <i>group-id</i> values are 1 to 16. detail displays additional status information on the upstream and downstream interfaces in each group (see Figure 35-3).
show interfaces <i>interface</i> Command Mode: EXEC	Displays the current status of a port or port-channel interface assigned to an uplink-state group. <i>interface</i> specifies one of the following interface types: 10-Gigabit Ethernet: Enter tengigabitethernet <i>slot/port</i> . 40-Gigabit Ethernet: Enter fortygigabitethernet <i>slot/port</i> . Port channel: Enter port-channel {1-512}. If a downstream interface in an uplink-state group has been disabled (Oper Down state) by uplink-state tracking because an upstream port went down, the message error-disabled[UFD] is displayed in the output (see Figure 35-4).
show running-config uplink-state-group [<i>group-id</i>] Command Mode: EXEC Or show configuration Command Mode: UPLINK-STATE-GROUP	Displays the current configuration of all uplink-state groups (Figure 35-5) or a specified group (Figure 35-6). Valid <i>group-id</i> values are 1 to 16.

Figure 35-3. show uplink-state-group Command Output

```
FTOS# show uplink-state-group

Uplink State Group: 1   Status: Enabled, Up
Uplink State Group: 3   Status: Enabled, Up
Uplink State Group: 5   Status: Enabled, Down
Uplink State Group: 6   Status: Enabled, Up
Uplink State Group: 7   Status: Enabled, Up
Uplink State Group: 16  Status: Disabled, Up

FTOS# show uplink-state-group 16
Uplink State Group: 16  Status: Disabled, Up

FTOS#show uplink-state-group detail
(Up): Interface up      (Dwn): Interface down  (Dis): Interface disabled

Uplink State Group   : 1           Status: Enabled, Up
Upstream Interfaces  :
Downstream Interfaces:

Uplink State Group   : 3           Status: Enabled, Up
Upstream Interfaces  : Tengig 0/46(Up) Tengig 0/47(Up)
Downstream Interfaces: Te 13/0(Up) Te 13/1(Up) Te 13/3(Up) Te 13/5(Up) Te 13/6(Up)

Uplink State Group   : 5           Status: Enabled, Down
Upstream Interfaces  : Tengig 0/0(Dwn) Tengig 0/3(Dwn) Tengig 0/5(Dwn)
Downstream Interfaces: Te 13/2(Dis) Te 13/4(Dis) Te 13/11(Dis) Te 13/12(Dis) Te 13/13(Dis)
                      Te 13/14(Dis) Te 13/15(Dis)

Uplink State Group   : 6           Status: Enabled, Up
Upstream Interfaces  :
Downstream Interfaces:

Uplink State Group   : 7           Status: Enabled, Up
Upstream Interfaces  :
Downstream Interfaces:

Uplink State Group   : 16          Status: Disabled, Up
Upstream Interfaces  : Tengig 0/41(Dwn) Po 8(Dwn)
Downstream Interfaces: Tengig 0/40(Dwn)
```

Figure 35-4. show interfaces Command: UFD Output

```

FTOS#show interfaces tengigabitethernet 7/45
TenGigabitEthernet 7/45 is up, line protocol is down (error-disabled[UFD])
Hardware is Dell Force10Eth, address is 00:01:e8:32:7a:47
  Current address is 00:01:e8:32:7a:47
Interface index is 280544512
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 1000 Mbit, Mode auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:25:46
Queueing strategy: fifo
Input Statistics:
  0 packets, 0 bytes
  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 runs, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  0 packets, 0 bytes, 0 underruns
  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 Unicasts
  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: 00:01:23

```

Figure 35-5. show running-config uplink-state-group Command: UFD Output

```

FTOS#show running-config uplink-state-group
!
uplink-state-group 1
no enable
downstream TenGigabitEthernet 0/0
upstream TenGigabitEthernet 0/1
FTOS#

```

Figure 35-6. show configuration Command: UFD Output

```

FTOS(conf-uplink-state-group-16)# show configuration
!
uplink-state-group 16
no enable
description test
downstream disable links all
downstream TengigabitEthernet 0/40
upstream TengigabitEthernet 0/41
upstream Port-channel 8

```

Sample Configuration: Uplink Failure Detection

Figure 35-7 shows a sample configuration of Uplink Failure Detection on a switch/router in which you:

- Configure uplink-state group 3.
- Add downstream links TenGigabitEthernet 0/1, 0/2, 0/5, 0/9, 0/11, and 0/12.
- Configure two downstream links to be disabled if an upstream link fails.
- Add upstream links TenGigabitEthernet 0/3 and 0/4.
- Add a text description for the group.
- Verify the configuration with various **show** commands.

Figure 35-7. Configuring Uplink Failure Detection

```

FTOS(conf)#uplink-state-group 3
FTOS(conf-uplink-state-group-3)#

00:23:52: %STKUNIT0-M:CP %IFMGR-5-ASTATE_UP: Changed uplink state group Admin state to up:
Group 3

FTOS(conf-uplink-state-group-3)#downstream tengigabitethernet 0/1-2,5,9,11-12
FTOS(conf-uplink-state-group-3)#downstream disable links 2
FTOS(conf-uplink-state-group-3)#upstream tengigabitethernet 0/3-4
FTOS(conf-uplink-state-group-3)#description Testing UFD feature
FTOS(conf-uplink-state-group-3)#show config
!
uplink-state-group 3
  description Testing UFD feature
  downstream disable links 2
  downstream TenGigabitEthernet 0/1-2,5,9,11-12
  upstream TenGigabitEthernet 0/3-4

FTOS#show running-config uplink-state-group
!
uplink-state-group 3
  description Testing UFD feature
  downstream disable links 2
  downstream TenGigabitEthernet 0/1-2,5,9,11-12
  upstream TenGigabitEthernet 0/3-4

FTOS#show uplink-state-group 3

Uplink State Group: 3   Status: Enabled, Up

FTOS#show uplink-state-group detail

(Up): Interface up   (Dwn): Interface down   (Dis): Interface disabled

Uplink State Group   : 3           Status: Enabled, Up
Upstream Interfaces  : Te 0/3(Up) Te 0/4(Up)
Downstream Interfaces: Te 0/1(Up) Te 0/2(Up) Te 0/5(Up) Te 0/9(Up) Te 0/11(Up)
                    Te 0/12(Up)

< After a single uplink port fails >

FTOS#show uplink-state-group detail

(Up): Interface up   (Dwn): Interface down   (Dis): Interface disabled

Uplink State Group   : 3           Status: Enabled, Up
Upstream Interfaces  : Te 0/3(Dwn) Te 0/4(Up)
Downstream Interfaces: Te 0/1(Dis) Te 0/2(Dis) Te 0/5(Up) Te 0/9(Up) Te 0/11(Up)
                    Te 0/12(Up)

```

Upgrade Procedures

Find the Upgrade Procedures

To see all the requirements to upgrade to the desired Dell Force10 operating software (FTOS) version, go to the *FTOS Release Notes* for your system type. Follow the procedures in the *FTOS Release Notes* for the software version you wish to upgrade to.

Get Help with Upgrades

Direct any questions or concerns about the FTOS Upgrade Procedures to the Dell Force10 Technical Support Center. You can reach Technical Support:

- On the Web: www.force10networks.com/support/
- By email: support@force10networks.com
- By phone: US and Canada: 866.965.5800, International: 408.965.5800

Virtual LANs (VLAN)

This section contains the following subsections:

- [Default VLAN](#)
- [Port-Based VLANs](#)
- [VLANs and Port Tagging](#)
- [Configuration Task List for VLANs](#)
- [Enable Null VLAN as the Default VLAN](#)

Virtual LANs (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. The Dell Force10 operating software (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 more information about VLANs, refer to IEEE Standard 802.1Q *Virtual Bridged Local Area Networks*. In this guide, see also:

- [Bulk Configuration in Interfaces](#)
- [VLAN Stacking](#)

For a complete listing of all commands related to FTOS VLANs, refer to these *FTOS Command Reference Guide* chapters:

- [Interfaces](#)
- [Security](#)
- [GARP VLAN Registration Protocol \(GVRP\)](#)
- [VLAN Stacking](#)
- [Per-VLAN Spanning Tree Plus \(PVST+\)](#)

Table 37-1 lists the defaults for VLANs in FTOS.

Table 37-1. 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

Default VLAN

When you configure interfaces 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 37-1 shows 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, the show vlan command in EXEC privilege mode indicates that the interface is now part of the Default VLAN (VLAN 1).

Figure 37-1. Interfaces and the Default VLAN Example

```

FTOS(conf)#int tengig 3/2
FTOS(conf-if)#no shut
FTOS(conf-if)#switchport
FTOS(conf-if)#show config
!
interface Tengigabitethernet 3/2
  no ip address
  switchport
  no shutdown
FTOS(conf-if)#end
FTOS#show vlan

Codes: * - Default VLAN, G - GVRP VLANs

  NUM   Status   Q Ports
 *   1     Active   U Tengig 3/2
    2     Active   T Po1(So 0/0-1)
                        T Tengig 3/0
FTOS#

```

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)

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



Note: You cannot assign an IP address to the default VLAN. To assign an IP address to a VLAN that is currently the default VLAN, create another VLAN and assign it to be the default VLAN. For more information about assigning IP addresses, refer to [Assign an IP Address to a VLAN](#).

Untagged interfaces must be part of a VLAN. To remove an untagged interface from the default VLAN, you must create another VLAN and place the interface into that VLAN. Alternatively, use the `no switchport` command, and FTOS removes the interface from the default VLAN.

A tagged interface requires an additional step to remove it from Layer 2 mode. Because 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 to the default vlan can you use the `no switchport` command to remove the interface from Layer 2 mode. For more information, refer to [VLANs and Port Tagging](#).

Port-Based VLANs

Port-based VLANs are a broadcast domain defined by different ports or interfaces. In FTOS, a port-based VLAN can contain interfaces from different stack units within the chassis. 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. Because traffic is only broadcast or flooded to the interfaces within a VLAN, the VLAN 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 you enable tagging, a tag header is added to the frame after the destination and source MAC addresses. That information is preserved as the frame moves through the network.

[Figure 37-2](#) shows the structure of a frame with a tag header. The VLAN ID is inserted in the tag header.

Figure 37-2. Tagged Frame Format

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

FN00001B

The tag header contains some key information used by FTOS:

- The VLAN protocol identifier identifies the frame as tagged according to the IEEE 802.1Q specifications (2 bytes).
- Tag control information (TCI) includes the VLAN ID (2 bytes total). The VLAN ID can have 4,096 values, but two 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 system 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

This section contains the following VLAN configuration tasks:

- [Create a Port-Based VLAN](#) (mandatory)
- [Assign Interfaces to a VLAN](#) (optional)
- [Assign an IP Address to a VLAN](#) (optional)
- [Enable Null VLAN as the Default VLAN](#)

Create a Port-Based VLAN

The default VLAN as VLAN 1 is part of the system startup configuration and does not require configuration. To configure a port-based VLAN, you must create the VLAN and then add physical interfaces or port channel (LAG) interfaces to the VLAN.

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

Command Syntax	Command Mode	Purpose
<code>interface vlan <i>vlan-id</i></code>	CONFIGURATION	Configure a port-based VLAN (if the VLAN ID is different from the default VLAN ID) and enter INTERFACE VLAN mode. After you create a VLAN, to activate the VLAN, you must assign interfaces in Layer 2 mode to the VLAN.

To view the configured VLANs, use the `show vlan` command in EXEC privilege mode (Figure 37-3).

Figure 37-3. show vlan Command Example

```
FTOS#show vlan

Codes: * - Default VLAN, G - GVRP VLANs, R - Remote Port Mirroring VLANs, P - Primary, C -
Community, I - Isolated
Q: U - Untagged, T - Tagged
   x - Dot1x untagged, X - Dot1x tagged
   G - GVRP tagged, M - Vlan-stack, H - VSN tagged
   i - Internal untagged, I - Internal tagged, v - VLT untagged, V - VLT tagged

  NUM      Status      Description          Q Ports
  ---      -
  *  1       Inactive    a
  *  20      Active
  *  1002    Active
  *  1002    Active

          U Po32()
          U Te 0/3,5,13,53-56
          T Te 0/3,13,55-56

FTOS#
```

A VLAN is active only if the VLAN contains interfaces and those interfaces are operationally up. In Figure 37-3, VLAN 1 is inactive because it contains the interfaces that are not active. The other VLANs listed in the Figure 37-3 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 effect 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

To assign interfaces in Layer 2 mode to a VLAN, use the `tagged` and `untagged` commands. To place an interface in Layer 2 mode, use the `switchport` command.

These Layer 2 interfaces can further be designated as tagged or untagged. For more information, refer to the [Interfaces](#) and [Configure Layer 2 \(Data Link\) Mode](#). When you place an interface in Layer 2 mode using 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 37-3 shows that six VLANs are configured, 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 about this command, refer to the command statement in the Layer 2 chapter of the *FTOS Command Reference Guide*.

To view just the interfaces that are in Layer 2 mode, use the `show interfaces switchport` command in EXEC privilege mode or 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. To tag interfaces, follow these steps:

Step	Command Syntax	Command Mode	Purpose
1	interface vlan <i>vlan-id</i>	CONFIGURATION	Access INTERFACE VLAN mode of the VLAN to which you want to assign the interface.
2	tagged <i>interface</i>	INTERFACE	Enable an interface to include the IEEE 802.1Q tag header.

Figure 37-4 shows the steps to add a tagged interface (in this case, port channel 1) to VLAN 4.

Figure 37-4. Example of Adding an Interface to Another VLAN

```

FTOS#show vlan

Codes: * - Default VLAN, G - GVRP VLANs

  NUM   Status   Q Ports
*  1     Inactive
  2     Active   T Pol(So 0/0-1)
                   T Tengig 3/0
  3     Active   T Pol(So 0/0-1)
                   T Tengig 3/1

FTOS#config
FTOS(conf)#int vlan 4
FTOS(conf-if-vlan)#tagged po 1
FTOS(conf-if-vlan)#show conf
!
interface Vlan 4
  no ip address
  tagged Port-channel 1
FTOS(conf-if-vlan)#end
FTOS#show vlan

Codes: * - Default VLAN, G - GVRP VLANs

  NUM   Status   Q Ports
*  1     Inactive
  2     Active   T Pol(So 0/0-1)
                   T Tengig 3/0
  3     Active   T Pol(So 0/0-1)
                   T Tengig 3/1
  4     Active   T Pol(So 0/0-1)
FTOS#

```

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.

Except for hybrid ports, only a tagged interface can be a member of multiple VLANs. You can assign hybrid ports to two VLANs if the port is untagged in one VLAN and tagged in all others.

When you remove a tagged interface from a VLAN (using the `no tagged interface` command), it remains tagged only if it is a tagged interface in another VLAN. If you remove the tagged interface from the only VLAN to which it belongs, the interface is placed in the default VLAN as an untagged interface.

To move untagged interfaces from the default VLAN to another VLAN, use the `untagged` command and follow these steps

Step	Command Syntax	Command Mode	Purpose
1	<code>interface vlan <i>vlan-id</i></code>	CONFIGURATION	Access INTERFACE VLAN mode of the VLAN to which you want to assign the interface.
2	<code>untagged <i>interface</i></code>	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 37-5](#) shows the steps and commands to move an untagged interface from the default VLAN to another VLAN.

Figure 37-5. Example of Moving an Untagged Interface to Another VLAN

```

FTOS#show vlan
Codes: * - Default VLAN, G - GVRP VLANs

  NUM   Status   Q Ports
*   1     Active   U Tengig 3/0
  2     Active   T Pol(So 0/0-1)
                        T Tengig 3/0
  3     Active   T Pol(So 0/0-1)
                        T Tengig 3/1
  4     Inactive

FTOS#conf
FTOS(conf)#int vlan 4
FTOS(conf-if-vlan)#untagged tengig 3/2
FTOS(conf-if-vlan)#show config
!
interface Vlan 4
no ip address
untagged Tengigabitethernet 3/2
FTOS(conf-if-vlan)#end
FTOS#show vlan
Codes: * - Default VLAN, G - GVRP VLANs

  NUM   Status   Q Ports
*   1     Inactive
  2     Active   T Pol(So 0/0-1)
                        T Tengig 3/0
  3     Active   T Pol(So 0/0-1)
                        T Tengig 3/1
  4     Active   U Tengig 3/2
FTOS#

```

Use the `show vlan` command to determine interface status. Interface (tengig 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 (tengig 3/2). Since the Default VLAN no longer contains any interfaces, it is listed as inactive.

The only way to remove an interface from the default VLAN is to place the interface in Default mode by using the `no switchport` command in `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.



Note: You cannot assign an IP address to the default VLAN, which, by default, is VLAN 1. To assign another VLAN ID to the default VLAN, use the `default vlan-id vlan-id` command.

To assign an IP address, use the following command in `INTERFACE` mode:

Command Syntax	Command Mode	Purpose
<code>ip address <i>ip-address mask</i> [secondary]</code>	INTERFACE	Configure an IP address and mask on the interface. <ul style="list-style-type: none"> <i>ip-address mask</i> — Enter an address in dotted-decimal format (A.B.C.D) and the mask must be in slash format (/24). <i>secondary</i> — This is the interface’s backup IP address.

In FTOS, you can place VLANs and other logical interfaces in Layer 3 mode to receive and send routed traffic. For more information, refer to [Bulk Configuration](#).

Native VLANs


Traditionally, 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. Such ports are called “hybrid ports”. Physical and port-channel interfaces may be hybrid ports.

Native VLAN is useful in deployments where a Layer 2 port can receive both tagged and untagged traffic on the same physical port. The classic example is connecting a voice over IP (VOIP) phone and a PC to the same port of the switch. The VOIP phone is configured to generate tagged packets (with `VLAN = VOICE VLAN`), and the attached PC generates untagged packets.

To configure a port so that it can be a member of an untagged and tagged VLANs, follow these steps:

Step	Task	Command	Command Mode
1	Remove any Layer 2 or Layer 3 configurations from the interface.		INTERFACE
2	Configure the interface for hybrid mode.	portmode hybrid	INTERFACE
3	Configure the interface for switchport mode.	switchport	INTERFACE
4	Add the interface to a tagged or untagged VLAN.	[tagged untagged]	VLAN INTERFACE

 **Note:** You cannot configure an existing switchport or port channel interface for Native VLAN. Interfaces must have no other Layer 2 or Layer 3 configurations when entering the command portmode hybrid or a message similar to [Message 1](#) is displayed.

Message 1 Native VLAN Error

```
% Error: Port is in Layer-2 mode Tengig 5/6.
```

Enable Null VLAN as the Default VLAN

In a Carrier Ethernet for Metro Service environment, service providers who perform frequent reconfigurations for customers with changing requirements occasionally enable multiple interfaces, each connected to a different customer, before the interfaces are fully configured. This presents a vulnerability because both interfaces are initially placed in the native VLAN (VLAN 1) and for that period customers are able to access each other's networks. FTOS has a null VLAN to eliminate this vulnerability.

When you enable the null VLAN, all ports are placed into by it default, so that even if you activate the physical ports of multiple customers, no traffic is allowed to traverse the links until each port is placed in another VLAN.

To disable the default VLAN, use the following command:

Task	Command Syntax	Command Mode
Disable the default VLAN, so that all ports belong to the Null VLAN until configured as a member of another VLAN.	default-vlan disable Default: the default VLAN is enabled (no default-vlan disable).	CONFIGURATION

Virtual Router Redundancy Protocol (VRRP)

This chapter covers the following information:

- [Overview](#)
- [VRRP Benefits](#)
- [VRRP Implementation](#)
- [VRRP Configuration](#)
- [Sample Configurations](#)

Overview

Virtual router redundancy protocol (VRRP) is designed to eliminate a single point of failure in a statically routed network.

VRRP specifies a MASTER router that owns the next hop IP and MAC address for end stations on a local area network (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 that new MASTER continues routing traffic.

VRRP uses the virtual router identifier (VRID) to identify each virtual router configured. 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.

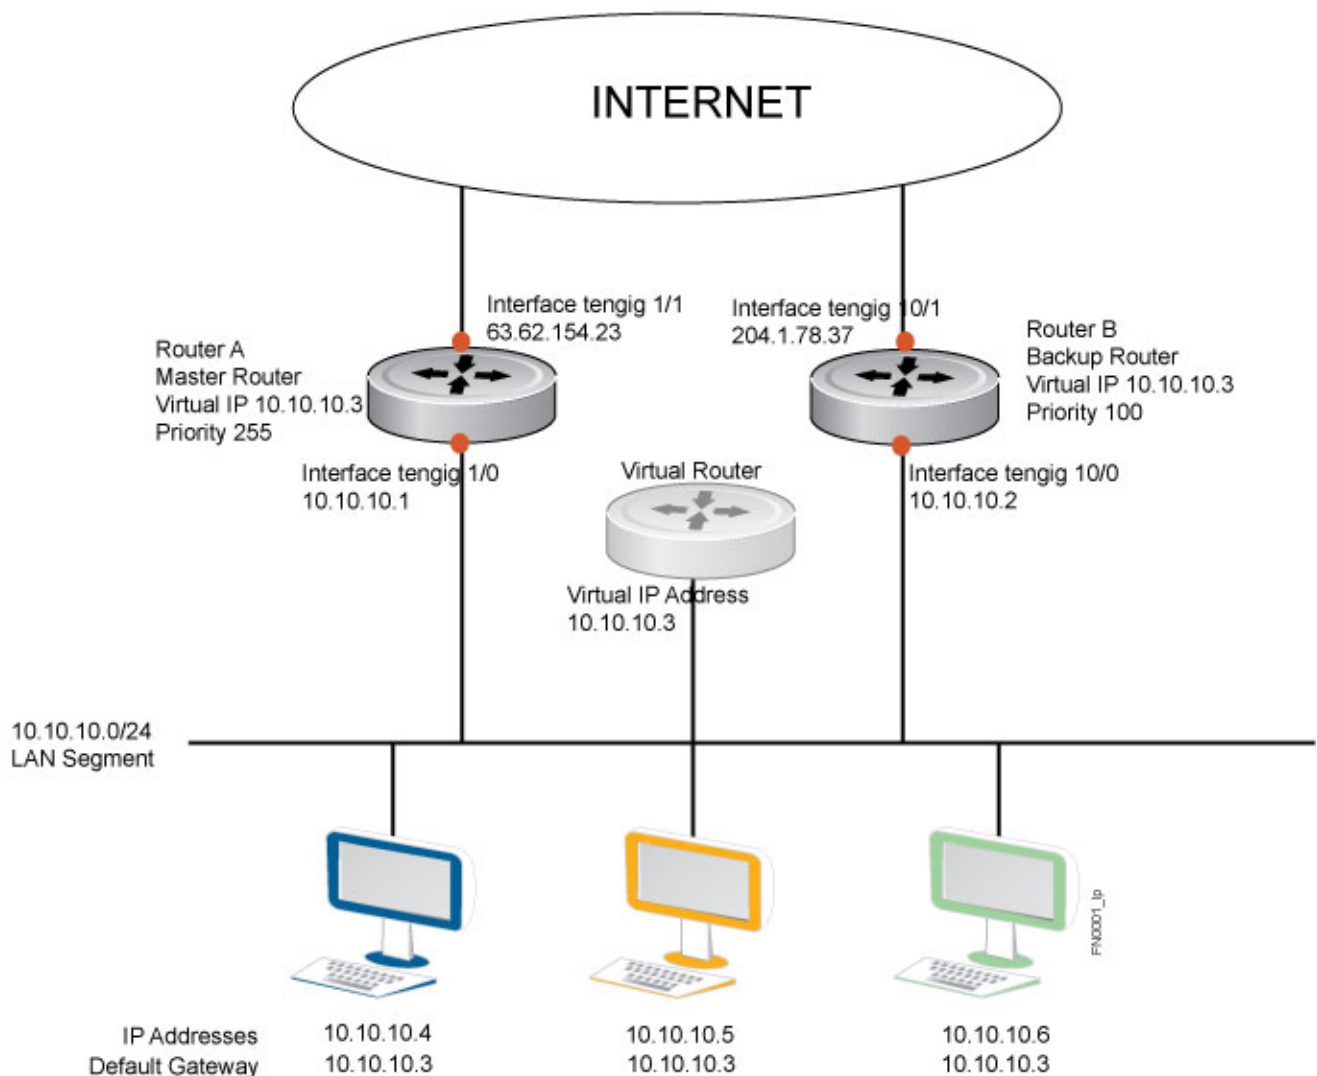
VRRP packets are transmitted with the virtual router MAC address as the source MAC address. The MAC address is in the following format: 00-00-5E-00-01-`{VRID}`. The first three octets are unchangeable. The next two octets (00-01) indicate the address block assigned to the VRRP protocol, and are unchangeable. The final octet changes depending on the VRRP VRID and allows for up to 255 VRRP routers on a network.

[Figure 38-1](#) 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 IP address configured on the virtual router. When any host on the LAN segment wants to access the Internet, it sends packets to the IP address of the virtual router.

In [Figure 38-1](#), 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 TenGigabitEthernet 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 becomes unavailable, VRRP elects a new MASTER Router. 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 Tengigabitethernet 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.

Figure 38-1. Basic VRRP Configuration



For more information about 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 point-of-failure. End-station connections to the network are redundant and they are not dependent on internal gateway protocol (IGP) to converge or update routing tables.

VRRP Implementation

The MXL 10/40GbE Switch supports a total of 2000 VRRP groups on a switch and 255 VRRP groups per interface (Table 38-1).

Within a single VRRP group, up to 12 virtual IP addresses are supported. Virtual IP addresses can belong to 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.

Default VRRP settings may affect the maximum number of groups that can be configured and work efficiently, as a result of hardware throttling VRRP advertisement packets reaching the CP on the MXL Switch. To avoid throttling VRRP advertisement packets, Dell Force10 recommends increasing the VRRP advertisement interval to a value higher than the default value of 1 second. Table 38-1 list the recommendations.

Table 38-1. Recommended VRRP Advertise Intervals

	Recommended Advertise Interval	Groups/Interface
Total VRRP Groups		
Less than 250	1 second	255
Between 250 and 450	2 - 3 seconds	255
Between 450 and 600	3 - 4 seconds	255
Between 600 and 800	4 seconds	255
Between 800 and 1000	5 seconds	255
Between 1000 and 1200	7 seconds	255
Between 1200 and 1500	8 seconds	255

The recommendations in Table 38-1 may vary depending on various factors like address resolution protocol (ARP) broadcasts, IP broadcasts, or spanning tree protocol (STP) before changing the advertisement interval. When the number of packets processed by the CP processor increases or decreases based on the dynamics of the network, the advertisement intervals in 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. Take extra caution when increasing the advertisement interval, as the increased dead interval may cause packets to be dropped during that switch-over time.

VRRP Configuration

By default, VRRP is not configured.

Configuration Task List for VRRP

The following list specifies the configuration tasks for VRRP:

- [Create a Virtual Router](#) (mandatory)
- [Assign Virtual IP addresses](#) (mandatory)
- [Set the VRRP Group \(Virtual Router\) Priority](#) (optional)
- [Configure VRRP Authentication](#) (optional)
- [Disable Preempt](#) (optional)
- [Change the Advertisement Interval](#) (optional)
- [Track an Interface or Object](#) (optional)
- [VRRP Initialization Delay](#)

For a complete listing of all commands related to VRRP, refer to *FTOS Command Line Interface Guide*.

Create a Virtual Router

To enable VRRP, you must create a virtual router. In FTOS, a VRRP group is identified by the virtual router identifier (VRID).

To enable a virtual router, use the following command in INTERFACE mode (Figure 38-2). To delete a VRRP group, use the `no vrrp-group vrid` command in INTERFACE mode.

Task	Command Syntax	Command Mode
Create a virtual router for that interface with a VRID. VRID Range: 1 to 255	<code>vrrp-group vrid</code>	INTERFACE

Note: The interface must already have a Primary IP Address defined and be enabled.

Figure 38-2. vrrp-group Command Example

```
FTOS(conf)#int tengig 1/1
FTOS(conf-if-te-1/1)#vrrp-group 111
FTOS(conf-if-te-1/1-vrid-111)#
```

← Virtual Router ID and VRRP Group identifier

Figure 38-3. show config Command Example

```
FTOS(conf-if-te-1/1)#show conf
!
interface Tengigabitethernet 1/1
 ip address 10.10.10.1/24
!
 vrrp-group 111
 no shutdown
FTOS(conf-if-te-1/1)#
```

Note that the interface has an IP Address and is enabled

Assign Virtual IP addresses

Virtual routers contain virtual IP addresses configured for that VRRP group (VRID). A VRRP group does not transmit VRRP packets until you assign the virtual IP address to the VRRP group.

To activate a VRRP group on an interface (so that VRRP group starts transmitting VRRP packets), configure at least one virtual IP address in a VRRP group. The virtual IP address is the IP address of the virtual router and does not require the IP address mask.

You can configure up to 12 virtual IP addresses on a single VRRP group (VRID).

The following rules apply to virtual IP addresses:

- The virtual IP addresses must be in the same subnet as 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, Dell Force10 recommends configuring virtual IP addresses belonging to the *same* IP subnet for any one VRRP group.
 - For example, 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. The VRRP group (VRID 1) must 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).
- 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 configure multiple VRRP groups on an interface, only one of the VRRP groups can contain the interface primary or secondary IP address.

To configure a virtual IP address, follow these steps:

Step	Task	Command Syntax	Command Mode
1	Configure a VRRP group. VRID Range: 1 to 255	<code>vrrp-group vrrp-id</code>	INTERFACE
2	Configure virtual IP addresses for this VRID. Range: up to 12 addresses	<code>virtual-address ip-address1 [...ip-address12]</code>	INTERFACE -VRID

Figure 38-4. virtual-address Command Example

```
FTOS(conf-if-te-1/1-vrid-111)#virtual-address 10.10.10.1
FTOS(conf-if-te-1/1-vrid-111)#virtual-address 10.10.10.2
FTOS(conf-if-te-1/1-vrid-111)#virtual-address 10.10.10.3
FTOS(conf-if-te-1/1-vrid-111)#
```

Figure 38-5. show config Command Example

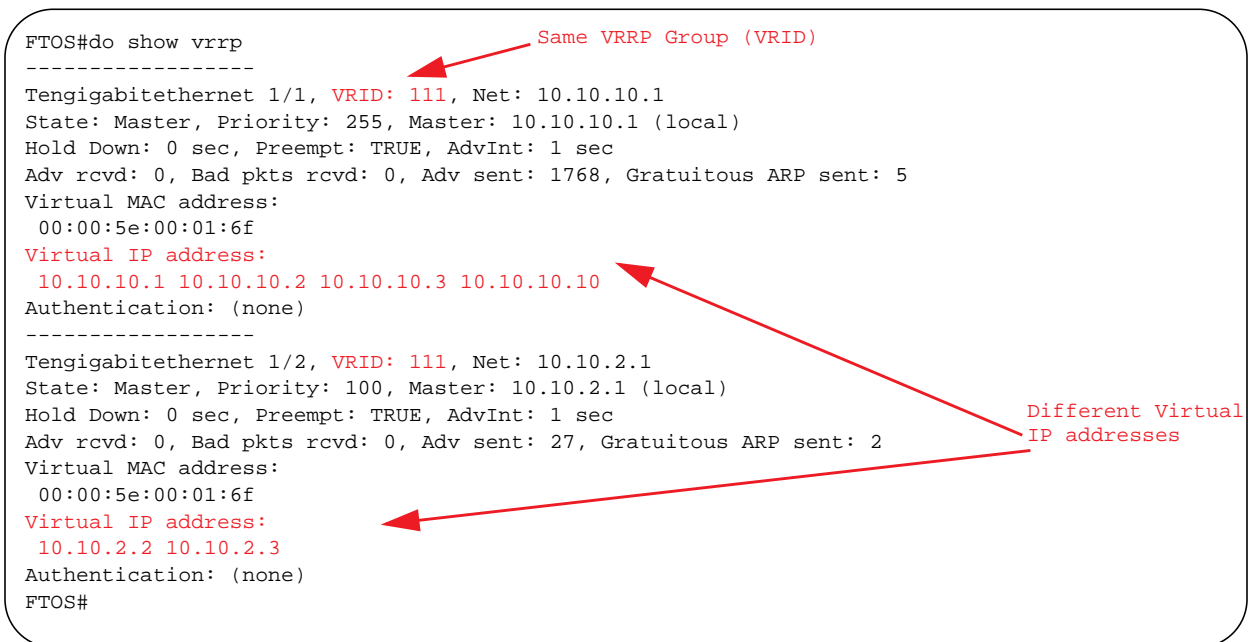
```
FTOS(conf-if-te-1/1)#show conf
!
interface Tengigabitethernet 1/1
 ip address 10.10.10.1/24
!
vrrp-group 111
 priority 255
 virtual-address 10.10.10.1
 virtual-address 10.10.10.2
 virtual-address 10.10.10.3
!
vrrp-group 222
 no shutdown
FTOS(conf-if-te-1/1)#
```

Note that the Primary IP address and the Virtual IP addresses are on the same subnet

Figure 38-6 shows the same VRRP group configured on multiple interfaces on different subnets.

Figure 38-6. show vrrp Command Example

```
FTOS#do show vrrp
-----
Tengigabitethernet 1/1, VRID: 111, Net: 10.10.10.1
State: Master, Priority: 255, Master: 10.10.10.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 1768, Gratuitous ARP sent: 5
Virtual MAC address:
 00:00:5e:00:01:6f
Virtual IP address:
 10.10.10.1 10.10.10.2 10.10.10.3 10.10.10.10
Authentication: (none)
-----
Tengigabitethernet 1/2, VRID: 111, Net: 10.10.2.1
State: Master, Priority: 100, Master: 10.10.2.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 27, Gratuitous ARP sent: 2
Virtual MAC address:
 00:00:5e:00:01:6f
Virtual IP address:
 10.10.2.2 10.10.2.3
Authentication: (none)
FTOS#
```



When the VRRP process completes its initialization, the State field contains either Master or Backup.

Set the VRRP Group (Virtual Router) Priority

Setting a virtual router priority to 255 ensures that router is the OWNER virtual router for the VRRP group. VRRP elects the MASTER router by choosing the router with the highest priority. The default priority for a virtual router is 100. The higher the number, the higher the priority. If the MASTER router fails, VRRP begins the election process to choose a new MASTER router based on the next-highest priority.

If two routers in a VRRP group come up at the same time and have 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 configure the VRRP group's priority, use the following command:

Task	Command Syntax	Command Mode
Configure the priority for the VRRP group. Range: 1 to 255 Default: 100	<code>priority priority</code>	INTERFACE -VRID

Figure 38-7. priority Command Example

```
FTOS(conf-if-te-1/2)#vrrp-group 111
FTOS(conf-if-te-1/2-vrid-111)#priority 125
```

Figure 38-8. show vrrp Command Example

```
FTOS#show vrrp
-----
Tengigabitethernet 1/1, VRID: 111, Net: 10.10.10.1
State: Master, Priority: 255, Master: 10.10.10.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 2343, Gratuitous ARP sent: 5
Virtual MAC address:
  00:00:5e:00:01:6f
Virtual IP address:
  10.10.10.1 10.10.10.2 10.10.10.3 10.10.10.10
Authentication: (none)
-----
Tengigabitethernet 1/2, VRID: 111, Net: 10.10.2.1
State: Master, Priority: 125, Master: 10.10.2.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 601, Gratuitous ARP sent: 2
Virtual MAC address:
  00:00:5e:00:01:6f
Virtual IP address:
  10.10.2.2 10.10.2.3
Authentication: (none)
FTOS(conf)#
```

Configure VRRP Authentication

Simple authentication of VRRP packets ensures that only trusted routers participate in VRRP processes. When you enable authentication, FTOS includes the password in its VRRP transmission, and the receiving router uses that password to verify the transmission.



Note: You must configure all virtual routers in the VRRP group at the same: authentication must be enabled with the same password or authentication is disabled.

To configure simple authentication, use the following command:

Task	Command Syntax	Command Mode
Configure a simple text password. Parameters: <i>encryption-type</i> : 0 indicates unencrypted; 7 indicates encrypted <i>password</i> : plain text	<code>authentication-type simple [<i>encryption-type</i>] <i>password</i></code>	INTERFACE-VRID

Figure 38-9. authentication-type Command Example

```
FTOS(conf-if-te-1/1-vrid-111)#authentication-type ?
FTOS(conf-if-te-1/1-vrid-111)#authentication-type simple 7 force10
```




Figure 38-10. show config Command Example (a Simple Password is Configured)

```
FTOS(conf-if-te-1/1-vrid-111)#show conf
!
vrrp-group 111
 authentication-type simple 7 387a7f2df5969da4 ← Encrypted password
 priority 255
 virtual-address 10.10.10.1
 virtual-address 10.10.10.2
 virtual-address 10.10.10.3
 virtual-address 10.10.10.10
FTOS(conf-if-te-1/1-vrid-111)#
```

Disable Preempt

The preempt command is enabled by default. The command forces the system to change the MASTER router if another router with a higher priority comes online.

To prevent the BACKUP router with the higher priority from becoming the MASTER router, disable the preempt command.



Note: You must configure all virtual routers in the VRRP group at the same: all configured with preempt enabled or configured with preempt disabled.

Because preempt is enabled by default, disable the preempt function with the following command in the VRRP mode. To re-enable preempt, use the preempt command. When preempt is enabled, it does not display in the show commands because it is a default setting.

Task	Command Syntax	Command Mode
Prevent any BACKUP router with a higher priority from becoming the MASTER router.	no preempt	INTERFACE-VRID

Figure 38-11. no preempt Command Example

```
FTOS(conf-if-te-1/1)#vrrp-group 111
FTOS(conf-if-te-1/1-vrid-111)#no preempt
FTOS(conf-if-te-1/1-vrid-111)#show conf
```

Figure 38-12. show config Command Example

```
FTOS(conf-if-te-1/1-vrid-111)#show conf
!
vrrp-group 111
 authentication-type simple 7 387a7f2df5969da4
 no preempt
 priority 255
 virtual-address 10.10.10.1
 virtual-address 10.10.10.2
 virtual-address 10.10.10.3
 virtual-address 10.10.10.10
FTOS(conf-if-te-1/1-vrid-111)#
```

Change the Advertisement Interval

By default, the MASTER router transmits a VRRP advertisement to all members of the VRRP group every 1 second, indicating it is operational and is the MASTER router. If the VRRP group misses three consecutive advertisements, the election process begins and the BACKUP virtual router with the highest priority transitions to MASTER.



Note: Dell Force10 recommends increasing the VRRP advertisement interval to a value higher than the default value of 1 second to avoid throttling VRRP advertisement packets. 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:

Task	Command Syntax	Command Mode
Change the advertisement interval setting. Range: 1-255 seconds Default: 1 second	advertise-interval <i>seconds</i>	INTERFACE-VRID

Figure 38-13. advertise-interval Command Example

```
FTOS(conf-if-te-1/1)#vrrp-group 111
FTOS(conf-if-te-1/1-vrid-111)#advertise-interval 10
FTOS(conf-if-te-1/1-vrid-111)#
```

Figure 38-14. show config Command Example

```
FTOS(conf-if-te-1/1-vrid-111)#show conf
!
vrrp-group 111
  advertise-interval 10
  authentication-type simple 7 387a7f2df5969da4
  no preempt
  priority 255
  virtual-address 10.10.10.1
  virtual-address 10.10.10.2
  virtual-address 10.10.10.3
  virtual-address 10.10.10.10
FTOS(conf-if-te-1/1-vrid-111)#
```

Track an Interface or Object

Set FTOS to monitor the state of any interface according to the virtual group. Each VRRP group can track up to 12 interfaces, which may affect the priority of the VRRP group. If 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.

Each VRRP group can track changes in the status of up to 12 interfaces and up to 20 additional objects, which may affect the priority of the VRRP group. If a tracked interface or object goes down, the VRRP group's priority is decreased by a default value of 10 (also known as cost). If the state of a tracked interface or object 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 must be less than 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.

For a virtual group, you can track the line-protocol state or the routing status of any of the following interfaces with the `interface interface` parameter:

- 40-Gigabit Ethernet: Enter `fortygigabitethernet slot/port` in the `track interface` command (refer to Step 1 below).
- 10-Gigabit Ethernet: Enter `tengigabitethernet slot/port`.
- Port channel: Enter `port-channel number`, where valid port-channel numbers are: 1 to 128.
- VLAN: Enter `vlan vlan-id`, where valid VLAN IDs are from 1 to 4094.

For a virtual group, you can also track the status of a configured object (`track object-id` command) by entering its object number. For more information, refer to [Object Tracking Configuration](#).

You can configure a tracked object for a VRRP group (using the `track object-id` command in INTERFACE-VRID mode) before you actually create the tracked object (using a `track object-id` command in CONFIGURATION mode) (Figure 38-15) and (Figure 38-16). However, no changes in the VRRP group's priority occur until the tracked object is defined and determined to be down.

To track an interface, use the following commands in VRRP mode:

Task	Command Syntax	Command Mode
Monitor an interface and, optionally, set a value to be subtracted from the interface's VRRP group priority. Cost Range: 1-254 Default: 10	<code>track <i>interface</i> [priority-cost <i>cost</i>]</code>	INTERFACE-VRID
(Optional) Display the configuration and UP or DOWN state of tracked objects, including the client (VRRP group) that is tracking an object's state.	<code>show track</code>	EXEC EXEC Privilege
(Optional) Display the configuration and UP or DOWN state of tracked interfaces and objects in VRRP groups, including the time since the last change in an object's state.	<code>show vrrp</code>	EXEC EXEC Privilege
(Optional) Display the configuration of tracked objects in VRRP groups on a specified interface.	<code>show running-config interface <i>interface</i></code>	EXEC EXEC Privilege

The sum of all the costs for all tracked interfaces must be less than the configured priority of the VRRP group.

Figure 38-15. track Command Example

```
FTOS(conf-if-te-1/1)#vrrp-group 111
FTOS(conf-if-te-1/1-vrid-111)#track tengigabitethernet 1/2
FTOS(conf-if-te-1/1-vrid-111)#
```

Figure 38-16. track Command Example (VRID mode)

```
FTOS(conf-if-te-1/1-vrid-111)#show conf
!
vrrp-group 111
  advertise-interval 10
  authentication-type simple 7 387a7f2df5969da4
  no preempt
  priority 255
  track Tengigabitethernet 1/2
  virtual-address 10.10.10.1
  virtual-address 10.10.10.2
  virtual-address 10.10.10.3
  virtual-address 10.10.10.10
FTOS(conf-if-te-1/1-vrid-111)#
```

Figure 38-17. show vrrp Command Example

```
FTOS#show vrrp
-----
TenGigabitEthernet 1/3, IPv4 VRID: 21, Version: 2, Net: 10.1.1.1
State: Master, Priority: 100, Master: 10.1.1.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 72, Gratuitous ARP sent: 1
Virtual MAC address:
  00:00:5e:00:01:15
Virtual IP address:
  10.1.1.2
Authentication: (none)
FTOS#
```

Figure 38-18. show running-config interface Command Example

```
FTOS#show running-config interface tengigabitethernet 1/3
!
interface TenGigabitEthernet 1/3
 ip address 10.1.1.1/24
!
 vrrp-group 21
  virtual-address 10.1.1.2
 no shutdown
FTOS#
```

VRRP Initialization Delay

When configured, VRRP is enabled immediately upon system reload or boot. VRRP initialization can be delayed to allow IGP and EGP protocols to be enabled prior to selecting the VRRP MASTER. This delay ensures that VRRP initializes with no errors or conflicts. You can configure the delay for up to 15 minutes, after which VRRP enables normally.

The delay timer is set on individual interfaces and is supported on all physical interfaces, VLANs, and link aggregation groups (LAGs).

When you configure both CLIs, the later timer rules the VRRP enabling. For example, if you configure `vrrp delay reload 600` and `vrrp delay minimum 300`, the following behavior occurs:

- When the system reloads, VRRP waits 600 seconds (10 minutes) to bring up VRRP on all interfaces that are up and configured for vrrp.
- When an interface comes up and becomes operational, the system waits 300 seconds (5 minutes) to bring up VRRP on that interface.

Task	Command Syntax	Command Mode
Set the delay time for VRRP initialization on an individual interface. This is the gap between an interface coming up and being operational, and VRRP enabling. Seconds range: 0-900 Default: 0	<code>vrrp delay minimum <i>seconds</i></code>	INTERFACE
Set the delay time for VRRP initialization on all the interfaces in the system configured for VRRP. This is the gap between system boot up completion and VRRP enabling. Seconds range: 0-900 Default: 0	<code>vrrp delay reload <i>seconds</i></code>	INTERFACE

Sample Configurations

VRRP for IPv4 Configuration

The configuration in [Figure 38-19](#) shows how to enable IPv4 VRRP. This example does not contain comprehensive directions and is intended to provide guidance for only a typical VRRP configuration. You can copy and paste from the example to your CLI. Be sure you make the necessary changes to support your own IP addresses, interfaces, names, etc. [Figure 38-19](#) shows the VRRP topology created with the CLI configuration in [Figure 38-20](#).

Figure 38-19. VRRP for IPv4 Topology

```
R2#show vrrp
-----
TenGigabitEthernet 2/31, VRID: 99, Net: 10.1.1.1
State: Master, Priority: 100, Master: 10.1.1.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 661, Gratuitous ARP sent: 1
Virtual MAC address:
00:00:5e:00:01:63
Virtual IP address:
10.1.1.3
Authentication: (none)
R2#
```

State Master: R2 was the first interface configured with VRRP

Virtual MAC is automatically assigned and is the same on both Routers

```
R3#show vrrp
-----
TenGigabitEthernet 3/21, VRID: 99, Net: 10.1.1.1
State: Backup, Priority: 100, Master: 10.1.1.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 331, Bad pkts rcvd: 0, Adv sent: 0, Gratuitous ARP sent: 0
Virtual MAC address:
00:00:5e:00:01:63
Virtual IP address:
10.1.1.3
Authentication: (none)
R3#
```

State Backup: R3 was the second interface configured with VRRP




Figure 38-20. Configure VRRP for IPv4 Router

```

R2(conf)#int tengig 2/31
R2(conf-if-te-2/31)#ip address 10.1.1.1/24
R2(conf-if-te-2/31)#vrrp-group 99
R2(conf-if-te-2/31-vrid-99)#priority 200
R2(conf-if-te-2/31-vrid-99)#virtual 10.1.1.3
R2(conf-if-te-2/31-vrid-99)#no shut
R2(conf-if-te-2/31)#show conf
!
interface Tengigabitethernet 2/31
 ip address 10.1.1.1/24
!
 vrrp-group 99
  priority 200
  virtual-address 10.1.1.3
 no shutdown
R2(conf-if-te-2/31)#end

R2#show vrrp
-----
Tengigabitethernet 2/31, VRID: 99, Net: 10.1.1.1
State: Master, Priority: 200, Master: 10.1.1.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 817, Gratuitous ARP sent: 1
Virtual MAC address:
 00:00:5e:00:01:63
Virtual IP address:
 10.1.1.3
Authentication: (none)
R2#
Router 3
R3(conf)#int tengig 3/21
R3(conf-if-te-3/21)#ip address 10.1.1.2/24
R3(conf-if-te-3/21)#vrrp-group 99
R3(conf-if-te-3/21-vrid-99)#virtual 10.1.1.3
R3(conf-if-te-3/21-vrid-99)#no shut
R3(conf-if-te-3/21)#show conf
!
interface Tengigabitethernet 3/21
 ip address 10.1.1.1/24
!
 vrrp-group 99
  virtual-address 10.1.1.3
 no shutdown
R3(conf-if-te-3/21)#end
R3#show vrrp
-----
Tengigabitethernet 3/21, VRID: 99, Net: 10.1.1.2
State: Backup, Priority: 100, Master: 10.1.1.1
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 698, Bad pkts rcvd: 0, Adv sent: 0, Gratuitous ARP sent: 0
Virtual MAC address:
 00:00:5e:00:01:63
Virtual IP address:
 10.1.1.3
Authentication: (none)

```


Debugging and Diagnostics

The chapter contains the following sections:

- [Offline Diagnostics](#)
- [Trace Logs](#)
- [Show Hardware Commands](#)
- [Environmental Monitoring](#)
- [Buffer Tuning](#)
- [Troubleshooting Packet Loss](#)
- [Application Core Dumps](#)
- [Mini Core Dumps](#)
- [TCP Dumps](#)

Offline Diagnostics

The offline diagnostics test suite is useful for isolating faults and debugging hardware. The diagnostics tests are grouped into three levels:

- **Level 0**—Level 0 diagnostics check for the presence of various components and perform essential path verifications. In addition, they verify the identification registers of the components on the board.
- **Level 1**—A smaller set of diagnostic tests. Level 1 diagnostics perform status, self-test, access, and read/write tests for all the components on the board and test their registers for appropriate values. In addition, they perform extensive tests on memory devices (for example, SDRAM, flash, NVRAM, EEPROM) wherever possible.
- **Level 2**—The full set of diagnostic tests. Level 2 diagnostics are used primarily for on-board MAC level, Physical level, and external loopback tests and more extensive component diagnostics. Various components on the board are put into loopback mode, and test packets are transmitted through those components. These diagnostics also perform snake tests using virtual local area network (VLAN) configurations.



Note: Diagnostic is not allowed in Stacking mode, including member stacking. Avoid stacking before executing the diagnostic tests in the chassis.

Important Points to Remember

- You can only perform offline diagnostics on an offline standalone unit. You cannot perform diagnostics if the ports are configured in a stacking group. Remove the port(s) from the stacking group before executing the diagnostic test.
- Diagnostics only test connectivity, not the entire data path.
- Diagnostic results are stored on the flash of the unit on which you performed the diagnostics.
- When offline diagnostics are complete, the unit or stack member reboots automatically.

Running Offline Diagnostics

To run offline diagnostics, follow these steps:

1. Place the unit in the offline state using the offline stack-unit command from EXEC Privilege mode (Figure 39-1).



The system reboots when the off-line diagnostics complete. This is an automatic process. A warning message appears when you implement the offline stack-unit command.

Warning - Diagnostic execution will cause stack-unit to reboot after completion of diags.

Proceed with Offline-Diags [confirm yes/no]:y

Figure 39-1. Taking a Stack Unit Offline

```
FTOS#offline stack-unit 2
Warning - Diagnostic execution will cause stack-unit to reboot after completion of diags.
Proceed with Offline-Diags [confirm yes/no]:y
5w6d12h: %STKUNIT0-M:CP %CHMGR-2-STACKUNIT_DOWN: Stack unit 2 down - stack unit offline
5w6d12h: %STKUNIT0-M:CP %IFMGR-1-DEL_PORT: Removed port: Tengig 2/1-48
FTOS#5w6d12h: %STKUNIT1-S:CP %IFMGR-1-DEL_PORT: Removed port: Tengig 2/1-48
```

2. Use the show system brief command from EXEC Privilege mode to confirm offline status (Figure 39-2).

Figure 39-2. Verifying the Offline/Online Status of a Stack Unit

```
FTOS#show system brief | no-more

Stack MAC : 00:1e:c9:bb:02:c4

Reload Type : normal-reload [Next boot : normal-reload]

-- Stack Info --
Unit  UnitType  Status      ReqTyp      CurTyp      Version     Ports
-----
  0   Management  online      MXL-10/40GbE  MXL-10/40GbE  8-3-16-0    56
  1   Member      not present
  2   Member      not present
  3   Member      not present
  4   Member      not present
  5   Member      not present
FTOS#
```

Trace Logs

In addition to the syslog buffer, the Dell Force10 operating software (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.

Auto Save on Crash or Rollover

Exception information on Master or Standby units is stored in the **flash://TRACE_LOG_DIR** directory. This directory contains files that save trace information when there has been a task crash or timeout.

On a Master unit, you can reach the **TRACE_LOG_DIR** files by file transfer protocol (FTP) or by using the show file command from the **flash://TRACE_LOG_DIR** directory.

On a Standby unit, you can reach the **TRACE_LOG_DIR** files only by using the show file command from the **flash://TRACE_LOG_DIR** directory.



Note: Non-management Member units do not support this functionality.

Figure 39-3. Command Example

```
FTOS#dir flash://TRACE_LOG_DIR
Directory of flash://TRACE_LOG_DIR

 1  drwx      4096   Jan 17 2011 15:02:16 +00:00 .
 2  drwx      4096   Jan 01 1980 00:00:00 +00:00 ..
 3  -rwx     100583  Feb 11 2011 20:41:36 +00:00 failure_trace0_RPM0_CP

flash: 2143281152 bytes total (2069291008 bytes free)
```

Show Hardware Commands

The show hardware command tree consists of EXEC Privilege commands used with the MXL Switch. These commands display information from a hardware sub-component and from hardware-based feature tables.

[Table 39-1](#) lists the show hardware commands available as of the latest FTOS version.



Note: Use the show hardware commands only under the guidance of Dell Force10 Technical Assistance Center.

Table 39-1. show hardware Commands

Command	Description
show hardware stack-unit {0-5} cpu management statistics	View the internal interface status of the stack-unit CPU port which connects to the external management interface.
show hardware stack-unit {0-5} cpu data-plane statistics	View the driver-level statistics for the data-plane port on the CPU for the specified stack-unit. It provides insight into the packet types entering the CPU to see whether CPU-bound traffic is internal (IPC traffic) or network control traffic, which the CPU must process.
show hardware stack-unit {0-5} buffer total-buffer	View the modular packet buffers details per stack unit and the mode of allocation.
show hardware stack-unit {0-5} buffer unit {0-1} total-buffer	View the modular packet buffers details per unit and the mode of allocation.
show hardware stack-unit {0-5} buffer unit {0-1} port {1-64 all} buffer-info	View the forwarding plane statistics containing the packet buffer usage per port per stack unit.
show hardware stack-unit {0-5} buffer unit {0-1} port {1-64} queue {0-14 all} buffer-info	View the forwarding plane statistics containing the packet buffer statistics per COS per port.
show hardware stack-unit {0-5} cpu party-bus statistics	View input and output statistics on the party bus, which carries inter-process communication traffic between CPUs.
show hardware stack-unit {0-5} drops unit {0-0} port {33-56}	View the ingress and egress internal packet-drop counters, MAC counters drop, and FP packet drops for the stack unit on per port basis. It assists in identifying the stack unit/port pipe/port that may experience internal drops.
show hardware stack-unit {0-5} stack-port {33-56}	View the input and output statistics for a stack-port interface.
show hardware stack-unit {0-5} unit {0-0} counters	View the counters in the field processors of the stack unit.
show hardware stack-unit {0-5} unit {0-0} details	View the details of the FP devices and Hi gig ports on the stack-unit.
show hardware stack-unit {0-5} unit {0-0} execute-shell-cmd {command}	Execute a specified bShell commands from the CLI without going into the bShell.
show hardware stack-unit {0-5} unit {0-0} ipmc-replication	View the Multicast IPMC replication table from the bShell.
show hardware stack-unit {0-5} unit {0-0} port-stats [detail]	View the internal statistics for each port-pipe (unit) on per port basis.
show hardware stack-unit {0-5} unit {0-0} register	View the stack-unit internal registers for each port-pipe.
show hardware stack-unit {0-5} unit {0-0} table-dump {table name}	View the tables from the bShell through the CLI without going into the bShell.

Environmental Monitoring

The MXL Switch components use environmental monitoring hardware to detect transmit power readings, receive power readings, and temperature updates. To receive periodic power updates, you must enable the enable optic-info-update interval command. The output in [Figure 39-4](#) shows the environment status.

Figure 39-4. show interfaces transceiver Command Example

```
FTOS#show int ten 0/49 transceiver
SFP is present
SFP 49 Serial Base ID fields
SFP 49 Id = 0x03
SFP 49 Ext Id = 0x04
SFP 49 Connector = 0x07
SFP 49 Transceiver Code = 0x00 0x00 0x00 0x01 0x20 0x40 0x0c 0x01
SFP 49 Encoding = 0x01
SFP 49 BR Nominal = 0x0c
SFP 49 Length(9um) Km = 0x00
SFP 49 Length(9um) 100m = 0x00
SFP 49 Length(50um) 10m = 0x37
SFP 49 Length(62.5um) 10m = 0x1e
SFP 49 Length(Copper) 10m = 0x00
SFP 49 Vendor Rev =
SFP 49 Laser Wavelength = 850 nm
SFP 49 CheckCodeBase = 0x78
SFP 49 Serial Extended ID fields
SFP 49 Options = 0x00 0x12
SFP 49 BR max = 0
SFP 49 BR min = 0
SFP 49 Vendor SN = P11C0B0
SFP 49 Datecode = 020919
SFP 49 CheckCodeExt = 0xb6

SFP 49 Diagnostic Information
=====
SFP 49 Rx Power measurement type = Average
=====
SFP 49 Temp High Alarm threshold = 100.000C
SFP 49 Voltage High Alarm threshold = 5.000V
SFP 49 Bias High Alarm threshold = 100.000mA
SFP 49 TX Power High Alarm threshold = 5.000mW
SFP 49 RX Power High Alarm threshold = 5.000mW
SFP 49 Temp Low Alarm threshold = -50.000C
SFP 49 Voltage Low Alarm threshold = 0.000V
SFP 49 Bias Low Alarm threshold = 0.000mA
SFP 49 TX Power Low Alarm threshold = 0.000mW
SFP 49 RX Power Low Alarm threshold = 0.000mW
=====
SFP 49 Temp High Warning threshold = 100.000C
SFP 49 Voltage High Warning threshold = 5.000V
SFP 49 Bias High Warning threshold = 100.000mA
SFP 49 TX Power High Warning threshold = 5.000mW
SFP 49 RX Power High Warning threshold = 5.000mW
SFP 49 Temp Low Warning threshold = -50.000C
SFP 49 Voltage Low Warning threshold = 0.000V
SFP 49 Bias Low Warning threshold = 0.000mA
SFP 49 TX Power Low Warning threshold = 0.000mW
SFP 49 RX Power Low Warning threshold = 0.000mW
=====
SFP 49 Temperature = 40.844C
SFP 49 Voltage = 3.169V
SFP 49 Tx Bias Current = 0.000mA
SFP 49 Tx Power = 0.000mW
SFP 49 Rx Power = 0.227mW
=====
SFP 49 Data Ready state Bar = False
SFP 49 Rx LOS state = False
SFP 49 Tx Fault state = False
```

Recognize an Over-Temperature Condition

An over-temperature condition occurs for one of two reasons:

- The card genuinely is too hot.
- A sensor has malfunctioned.

Inspect cards adjacent to the one reporting condition to discover the cause.

- If directly adjacent cards are not a normal temperature, suspect a genuine overheating condition.
- If directly adjacent cards are a 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 [Message 1](#).

Message 1 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, use the show alarms threshold command ([Figure 39-5](#)).

Figure 39-5. show alarms threshold Command Example

```
FTOS#show alarms threshold

-- Temperature Limits (deg C) --
-----
Unit0      BelowNormal  Normal  Elevated  Critical  Trip/Shutdown
          <=40         41      71        81        86
FTOS#
```

Troubleshoot an Over-Temperature 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. Ensure the air ducts are clean and that all fans are working correctly.
3. After the software has determined that the temperature levels are within normal limits, the card can be re-powered safely. To bring the stack unit back online, use the power-on command in EXEC mode.

In addition, Dell Force10 requires that you install blanks in all slots without a line card to control airflow for adequate system cooling.

Figure 39-6. show environment Command Example

```
FTOS#show environment

-- Unit Environment Status --
Unit  Status      Temp  Voltage
-----
* 0   online      71C   ok

* Management Unit

-- Thermal Sensor Readings (deg C) --
Unit  Sensor0  Sensor1  Sensor2  Sensor3  Sensor4  Sensor5  Sensor6  Sensor7  Sensor8
Sensor9
-----
0     45       43       66       61       66       62       70       65       67       71
```



Note: Exercise care when removing a card; if it has exceeded the major or shutdown thresholds, the card could be hot to the touch

Recognize an Under-Voltage Condition

If the system detects an under-voltage condition, it sends an alarm. To recognize this condition, look for the system messages in [Message 2](#).

Message 2 Under-Voltage Condition System Messages

```
%CHMGR-1-CARD_SHUTDOWN: Major alarm: Line card 2 down - auto-shutdown due to under voltage
```

[Message 2](#) indicates that the specified card is not receiving enough power. In response, the system first shuts down Power over Ethernet (PoE).

Troubleshoot an Under-Voltage Condition

To troubleshoot an under-voltage condition, check that the correct number of power supplies are installed and their Status light emitting diodes (LEDs) are lit.

The simple network management protocol (SNMP) traps and OIDs in [Table 39-2](#) provide information about environmental monitoring hardware and hardware components.

Table 39-2. SNMP Traps and OIDs

OID String	OID Name	Description
Receiving power		
.1.3.6.1.4.1.6027.3.10.1.2.5.1.6	chSysPortXfpRecvPower	OID to display the receiving power of the connected optics.
Transmitting power		
.1.3.6.1.4.1.6027.3.10.1.2.5.1.8	chSysPortXfpTxPower	OID to display the transmitting power of the connected optics.
Temperature		
.1.3.6.1.4.1.6027.3.10.1.2.5.1.7	chSysPortXfpRecvTemp	OID to display the Temperature of the connected optics. Note: These OIDs are only generated if you enable the CLI enable optic-info-update-interval is enabled command.
Hardware MIB Buffer Statistics		
.1.3.6.1.4.1.6027.3.16.1.1.4	fpPacketBufferTable	View the modular packet buffers details per stack unit and the mode of allocation.
.1.3.6.1.4.1.6027.3.16.1.1.5	fpStatsPerPortTable	View the forwarding plane statistics containing the packet buffer usage per port per stack unit.
.1.3.6.1.4.1.6027.3.16.1.1.6	fpStatsPerCOSTable	View the forwarding plane statistics containing the packet buffer statistics per COS per port.

Buffer Tuning

Buffer tuning allows you to modify the way your switch allocates buffers from its available memory and helps prevent packet drops during a temporary burst of traffic. The application-specific integrated circuit (ASICs) implement the key functions of queuing, feature lookups, and forwarding lookups in the hardware.

- Forwarding processor (FP) ASICs provide Ethernet MAC functions, queueing and buffering, as well as store feature and forwarding tables for hardware-based lookup and forwarding decisions. 10G and 40G interfaces use different FPs.

You can tune buffers at three locations ([Figure 39-7](#)).

- CSF – Output queues going from the CSF.
- FP Uplink—Output queues going from the FP to the CSF IDP links.
- Front-End Link—Output queues going from the FP to the front-end PHY.

All ports support eight queues, four for data traffic and four for control traffic. All eight queues are tunable.


Physical memory is organized into cells of 128 bytes. The cells are organized into two buffer pools—a dedicated buffer and a dynamic buffer.

- **Dedicated buffer** is reserved memory that cannot be used by other interfaces on the same ASIC or by other queues on the same interface. This buffer is always allocated, and no dynamic recarving takes place based on changes in interface status. Dedicated buffers introduce a trade-off. They provide each interface with a guaranteed minimum buffer to prevent an overused and congested interface from starving all other interfaces. However, this minimum guarantee means the buffer manager does not reallocate the buffer to an adjacent congested interface, which means that in some cases, memory is underused.
- **Dynamic buffer** is shared memory that is allocated as needed, up to a configured limit. Using dynamic buffers provides the benefit of statistical buffer sharing. An interface requests dynamic buffers when its dedicated buffer pool is exhausted. The buffer manager grants the request based on three conditions:
 - The number of used and available dynamic buffers.
 - The maximum number of cells that an interface can occupy.
 - Available packet pointers (2k per interface). Each packet is managed in the buffer using a unique packet pointer. Thus, each interface can manage up to 2k packets.

You can configure dynamic buffers per port on both 1G and 10G FPs and per queue on CSFs. By default, the FP dynamic buffer allocation is 10 times oversubscribed. For the 48-port 1G card:

- Dynamic Pool= Total Available Pool(16384 cells) – Total Dedicated Pool = 5904 cells
- Oversubscription ratio = 10
- Dynamic Cell Limit Per port = $59040/29 = 2036$ cells

Figure 39-7. Buffer Tuning Points



Deciding to Tune Buffers

Dell Force10 recommends exercising caution when configuring any non-default buffer settings, as tuning can significantly affect system performance. The default values work for most cases.

As a guideline, consider tuning buffers if traffic is very bursty (and coming from several interfaces). In this case:

- Reduce the dedicated buffer on all queues/interfaces.
- Increase the dynamic buffer on all interfaces.
- Increase the cell pointers on a queue that you are expecting will receive the largest number of packets.

Buffer Tuning Commands

To tune the buffers, use the following commands:

Task	Command	Command Mode
Define a buffer profile for the FP queues.	buffer-profile fp fsqueue	CONFIGURATION
Define a buffer profile for the CSF queues.	buffer-profile csf csqueue	CONFIGURATION
Change the dedicated buffers on a physical 1G interface.	buffer dedicated	BUFFER PROFILE
Change the maximum amount of dynamic buffers an interface can request.	buffer dynamic	BUFFER PROFILE
Change the number of packet-pointers per queue.	buffer packet-pointers	BUFFER PROFILE
Apply the buffer profile to a CSF to FP link.	buffer csf linecard	CONFIGURATION



FTOS Behavior: If you attempt to apply a buffer profile to a non-existent port-pipe, FTOS displays the following message. However, the configuration still appears in the running-config.

```
%DIFFSERV-2-DSA_BUFF_CARVING_INVALID_PORT_SET: Invalid FP port-set 2 for linecard 2. Valid range of port-set is <0-1>
```

Configuration changes take effect immediately and appear in the running configuration. Because under normal conditions all ports do not require the maximum possible allocation, the configured dynamic allocations can exceed the actual amount of available memory; this is called oversubscription. If you choose to oversubscribe the dynamic allocation, a burst of traffic on one interface might prevent other interfaces from receiving the configured dynamic allocation, which causes packet loss.

You cannot allocate more than the available memory for the dedicated buffers. If the system determines that the sum of the configured dedicated buffers allocated to the queues is more than the total available memory, the configuration is rejected, returning a syslog message similar to the following.

Table 39-3. Buffer Allocation Error

```
00:04:20: %S50N:0 %DIFFSERV-2-DSA_DEVICE_BUFFER_UNAVAILABLE: Unable to allocate dedicated buffers for stack-unit 0, port pipe 0, egress port 25 due to unavailability of cells
```



FTOS Behavior: When you remove a buffer-profile using the `no buffer-profile [fp | csf]` command from CONFIGURATION mode, the buffer-profile name still appears in the output of `show buffer-profile [detail | summary]`.

After a stack unit is reset, the buffer profile correctly returns to the default values, but the profile name remains. Remove it from the `show buffer-profile [detail | summary]` command output by using the `no buffer [fp-uplink | csf] stack-unit port-set buffer-policy` command from CONFIGURATION mode and the `no buffer-policy` command from INTERFACE mode.

Display the allocations for any buffer profile using the show commands in [Figure 39-9](#). Display the default buffer profile using the `show buffer-profile {summary | detail}` command from EXEC Privilege mode ([Figure 39-8](#)).

Figure 39-8. Display the Default Buffer Profile

```
FTOS#show buffer-profile detail interface tengigabitethernet 0/1
Interface Tengig 0/1
Buffer-profile -
Dynamic buffer 194.88 (Kilobytes)
Queue#           Dedicated Buffer      Buffer Packets
                  (Kilobytes)
0                 2.50                 256
1                 2.50                 256
2                 2.50                 256
3                 2.50                 256
4                 9.38                 256
5                 9.38                 256
6                 9.38                 256
7                 9.38                 256
```

Figure 39-9. Displaying Buffer Profile Allocations

```

FTOS#show running-config interface tengigabitethernet 2/0 !
interface TenGigabitEthernet 2/0
no ip address
mtu 9252
switchport
no shutdown
buffer-policy myfsbufferprofile

FTOS#show buffer-profile detail int tengig 0/10
Interface Tengig 0/10
Buffer-profile fsqueue-fp
Dynamic buffer 1256.00 (Kilobytes)
Queue#          Dedicated Buffer      Buffer Packets
                (Kilobytes)
0                3.00                 256
1                3.00                 256
2                3.00                 256
3                3.00                 256
4                3.00                 256
5                3.00                 256
6                3.00                 256
7                3.00                 256

FTOS#show buffer-profile detail fp-uplink stack-unit 0 port-set 0
Linecard 0 Port-set 0
Buffer-profile fsqueue-hig
Dynamic Buffer 1256.00 (Kilobytes)
Queue#          Dedicated Buffer      Buffer Packets
                (Kilobytes)
0                3.00                 256
1                3.00                 256
2                3.00                 256
3                3.00                 256
4                3.00                 256
5                3.00                 256
6                3.00                 256
7                3.00                 256

```

Using a Pre-Defined Buffer Profile

FTOS provides two pre-defined buffer profiles, one for single-queue (for example, non-QoS) applications, and one for four-queue (for example, QoS) applications.

Task	Command	Mode
Apply one of two pre-defined buffer profiles for all port pipes in the system.	buffer-profile global [1Q 4Q]	CONFIGURATION

You must reload the system for the global buffer profile to take effect ([Message 3](#)).

Message 3 Reload After Applying Global Buffer Profile

```
% Info: For the global pre-defined buffer profile to take effect, please save the config and reload the system.
```



FTOS Behavior: After you configure buffer-profile global 1Q, [Message 3](#) is displayed during every bootup. Only one reboot is required for the configuration to take effect; afterwards this bootup message may be ignored.



FTOS Behavior: The buffer profile does not returned to the default, 4Q. If you configure 1Q, save the running-config to the startup-config, and then delete the startup-config and reload the chassis. The only way to return to the default buffer profile is to explicitly configure 4Q, and then reload the chassis.

The buffer-profile global command fails if you have already applied a custom buffer profile on an interface.

Message 4 Global Buffer Profile Error

```
% Error: User-defined buffer profile already applied. Failed to apply global pre-defined buffer profile. Please remove all user-defined buffer profiles.
```

Similarly, when you configure buffer-profile global, you cannot not apply a buffer profile on any single interface.

Message 5 Global Buffer Profile Error

```
% Error: Global pre-defined buffer profile already applied. Failed to apply user-defined buffer profile on interface Tengig 0/1. Please remove global pre-defined buffer profile.
```

If the default buffer profile (4Q) is active, FTOS displays an error message instructing you to remove the default configuration using the no buffer-profile global command.

Sample Buffer Profile Configuration

The two general types of network environments are sustained data transfers and voice/data. Dell Force10 recommends a single-queue approach for data transfers ([Figure 39-10](#)).

Figure 39-10. Single Queue Application with Default Packet Pointers

```

!
buffer-profile fp fsqueue-fp
buffer dedicated queue0 3 queue1 3 queue2 3 queue3 3 queue4 3 queue5 3 queue6 3 queue7 3
buffer dynamic 1256
!
buffer-profile fp fsqueue-hig
buffer dedicated queue0 3 queue1 3 queue2 3 queue3 3 queue4 3 queue5 3 queue6 3 queue7 3
buffer dynamic 1256

!
buffer fp-uplink stack-unit 0 port-set 0 buffer-policy fsqueue-hig
buffer fp-uplink stack-unit 0 port-set 1 buffer-policy fsqueue-hig
!
Interface range tengig 0/1 - 48
buffer-policy fsqueue-fp

FTOS#sho run int Tengig 0/10
!
interface TenGigabitEthernet 0/10
no ip address

```

Troubleshooting Packet Loss

The show hardware stack-unit command is intended primarily to troubleshoot packet loss.

- show hardware stack-unit 0-5 cpu data-plane statistics
- show hardware stack-unit 0-5 cpu party-bus statistics
- show hardware stack-unit 0-5 drops unit 0-0 port 1-56
- show hardware stack-unit 0-5 stack-port 33-56
- show hardware stack-unit 0-5 unit 0-0 {counters | details | port-stats [detail] | register | ipmc-replication | table-dump}:
- show hardware {layer2| layer3} {eg acl | in acl} stack-unit 0-5 port-set 0-0
- show hardware layer3 qos stack-unit 0-5 port-set 0-0
- show hardware system-flow layer2 stack-unit 0-5 port-set 0-1 [counters]
- clear hardware stack-unit 0-5 counters
- clear hardware stack-unit 0-5 unit 0-0 counters
- clear hardware stack-unit 0-5 cpu data-plane statistics
- clear hardware stack-unit 0-5 cpu party-bus statistics
- clear hardware stack-unit 0-5 stack-port 33-56

Displaying Drop Counters

The show hardware stack-unit 0–11 drops [unit 0 [port 0–63]] command assists in identifying which stack unit, port pipe, and port is experiencing internal drops (Figure 39-11) and (Figure 39-12).

Figure 39-11. Displaying Drop Counter Statistics

```
FTOS#show hardware stack-unit 0 drops
UNIT No: 0
Total Ingress Drops :0
Total IngMac Drops :0
Total Mmu Drops :0
Total EgMac Drops :0
Total Egress Drops :0
UNIT No: 1
Total Ingress Drops :0
Total IngMac Drops :0
Total Mmu Drops :0
Total EgMac Drops :0
Total Egress Drops :0

FTOS#show hardware stack-unit 0 drops unit 0
Port# :Ingress Drops :IngMac Drops :Total Mmu Drops :EgMac Drops :Egress
Drops
1 0 0 0 0 0
2 0 0 0 0 0
3 0 0 0 0 0
4 0 0 0 0 0
5 0 0 0 0 0
6 0 0 0 0 0
7 0 0 0 0 0
8 0 0 0 0 0
```

Display drop counters with the show hardware stack-unit drops unit port command (Figure 39-12).

Figure 39-12. Displaying Buffer Statistics, Displaying Drop Counters

```

FTOS#show hardware stack-unit 0 drops unit 0 port 1
  --- Ingress Drops      ---
Ingress Drops           : 30
IBP CBP Full Drops     : 0
PortSTPnotFwd Drops    : 0
IPv4 L3 Discards       : 0
Policy Discards        : 0
Packets dropped by FP   : 14
(L2+L3) Drops          : 0
Port bitmap zero Drops : 16
Rx VLAN Drops          : 0

  --- Ingress MAC counters---
Ingress FCSDrops       : 0
Ingress MTUExceeds    : 0

  --- MMU Drops          ---
HOL DROPS              : 0
TxPurge CellErr       : 0
Aged Drops            : 0

  --- Egress MAC counters---
Egress FCS Drops      : 0

  --- Egress FORWARD PROCESSOR Drops ---
IPv4 L3UC Aged & Drops : 0
TTL Threshold Drops   : 0
INVALID VLAN CNTR Drops : 0
L2MC Drops           : 0
PKT Drops of ANY Conditions : 0
Hg MacUnderflow      : 0
TX Err PKT Counter   : 0

```

Dataplane Statistics

The `show hardware stack-unit cpu data-plane statistics` command provides insight into the packet types coming to the CPU. As shown in [Figure 39-13](#), the command output has been augmented, providing detailed RX/TX packet statistics on a per-queue basis. The objective is to see whether CPU-bound traffic is internal (so-called party bus or IPC traffic) or network control traffic, which the CPU must process.

Figure 39-13. Displaying Buffer Statistics, Displaying Dataplane Statistics

```
FTOS#show hardware stack-unit 2 cpu data-plane statistics
```

```
bc pci driver statistics for device:
```

```
rxHandle          :0
noMhdr            :0
noMbuf            :0
noClus            :0
recvd             :0
dropped           :0
recvToNet         :0
rxError           :0
rxDatapathErr    :0
rxPkt(COS0)      :0
rxPkt(COS1)      :0
rxPkt(COS2)      :0
rxPkt(COS3)      :0
rxPkt(COS4)      :0
rxPkt(COS5)      :0
rxPkt(COS6)      :0
rxPkt(COS7)      :0
rxPkt(UNIT0)     :0
rxPkt(UNIT1)     :0
rxPkt(UNIT2)     :0
rxPkt(UNIT3)     :0
transmitted       :0
txRequested       :0
noTxDesc          :0
txError           :0
txReqTooLarge    :0
txInternalError   :0
txDatapathErr    :0
txPkt(COS0)      :0
txPkt(COS1)      :0
txPkt(COS2)      :0
txPkt(COS3)      :0
txPkt(COS4)      :0
txPkt(COS5)      :0
txPkt(COS6)      :0
txPkt(COS7)      :0
txPkt(UNIT0)     :0
```

The `show hardware stack-unit cpu party-bus statistics` command displays input and output statistics on the party bus, which carries inter-process communication traffic between CPUs (Figure 39-14).

Figure 39-14. Displaying Party Bus Statistics

```
FTOS#sh hardware stack-unit 2 cpu party-bus statistics
```

```
Input Statistics:
```

```
  27550 packets, 2559298 bytes
  0 dropped, 0 errors
```

```
Output Statistics:
```

```
 1649566 packets, 1935316203 bytes
 0 errors
```

Displaying Stack Port Statistics

The show hardware stack-unit stack-port command displays input and output statistics for a stack-port interface (Figure 39-15).

Figure 39-15. Displaying Stack Unit Statistics

```

FTOS#show hardware stack-unit 2 stack-port 49
Input Statistics:
  27629 packets, 3411731 bytes
  0 64-byte pkts, 27271 over 64-byte pkts, 207 over 127-byte pkts
  17 over 255-byte pkts, 56 over 511-byte pkts, 78 over 1023-byte pkts
  0 Multicasts, 5 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  1649714 packets, 1948622676 bytes, 0 underruns
  0 64-byte pkts, 27234 over 64-byte pkts, 107970 over 127-byte pkts
  34 over 255-byte pkts, 504838 over 511-byte pkts, 1009638 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts, 1649714 Unicasts
  0 throttles, 0 discarded, 0 collisions
Rate info (interval 45 seconds):
  Input 00.00 Mbits/sec,          2 packets/sec, 0.00% of line-rate
  Output 00.06 Mbits/sec,        8 packets/sec, 0.00% of line-rate
FTOS#

```

Displaying Stack Member Counters

The show hardware stack-unit 0–5 {counters | details | port-stats [detail] | register} command displays internal receive and transmit statistics, based on the selected command option. A sample of the output is shown for the counters option in Figure 39-16.

Figure 39-16. Displaying Stack Unit Counters

RIPC4.ge0	:	1,202	+1,202
RUC.ge0	:	1,224	+1,217
RDBG0.ge0	:	34	+24
RDBG1.ge0	:	366	+235
RDBG5.ge0	:	16	+12
RDBG7.ge0	:	18	+12
GR64.ge0	:	5,176	+24
GR127.ge0	:	1,566	+1,433
GR255.ge0	:	4	+4
GRPKT.ge0	:	1,602	+1,461
GRBYT.ge0	:	117,600	+106,202
GRMCA.ge0	:	366	+235
GRBCA.ge0	:	12	+9
GT64.ge0	:	4	+3
GT127.ge0	:	964	+964
GT255.ge0	:	4	+4
GT511.ge0	:	1	+1
GTPKT.ge0	:	973	+972
GTBCA.ge0	:	1	+1
GTBYT.ge0	:	71,531	+71,467
RUC.cpu0	:	972	+971
TDBG6.cpu0	:	1,584	+1,449=

Application Core Dumps

Application core dumps are disabled by default. A core dump file can be very large. Due to memory requirements, the file can only be sent directly to an FTP server. It is not stored on the local flash. To enable full application core dumps, use the following command:

Task	Command Syntax	Command Mode
Enable RPM core dumps and specify the shutdown mode.	logging coredump server	CONFIGURATION

To undo this command, use the no logging coredump server command.

Mini Core Dumps

FTOS supports mini core dumps for application and kernel crashes. The mini core dump applies to Master, Standby, and Member units.

Application and kernel mini core dumps are always enabled. The mini core dumps contain the stack space and some other very minimal information that you can use to debug a crash. These files are small files and are written into flash until space is exhausted. When the flash is full, the write process is stopped.

A mini core dump contains critical information in the event of a crash.

- Mini core dump files are located in flash:/ (root dir).
- The application mini core file name format is f10StkUnit<Stack_unit_no>.<Application name>.acore.mini.txt.
- The kernel mini core file name format is f10StkUnit<Stack_unit_no>.kcore.mini.txt.

Sample files names are shown in [Figure 39-17](#) and a sample file text is shown in [Figure 39-18](#).

Figure 39-17. Mini application core file naming example

```

FTOS#dir
Directory of flash:

 1 drw-      16384   Jan 01 1980 00:00:00 +00:00 .
 2 drwx      1536   Sep 03 2009 16:51:02 +00:00 ..
 3 drw-       512   Aug 07 2009 13:05:58 +00:00 TRACE_LOG_DIR
 4 d---       512   Aug 07 2009 13:06:00 +00:00 ADMIN_DIR
 5 -rw-      8693   Sep 03 2009 16:50:56 +00:00 startup-config
 6 -rw-      8693   Sep 03 2009 16:44:22 +00:00 startup-config.bak
 7 -rw-       156   Aug 28 2009 16:16:10 +00:00 f10StkUnit0.mrtm.acore.mini.txt
 8 -rw-       156   Aug 28 2009 17:17:24 +00:00 f10StkUnit0.vrrp.acore.mini.txt
 9 -rw-       156   Aug 28 2009 18:25:18 +00:00 f10StkUnit0.sysd.acore.mini.txt
10 -rw-       156   Aug 28 2009 19:07:36 +00:00 f10StkUnit0.frrp.acore.mini.txt
11 -rw-       156   Aug 31 2009 16:18:50 +00:00 f10StkUnit2.sysd.acore.mini.txt
12 -rw-       156   Aug 29 2009 14:28:34 +00:00 f10StkUnit0.ipml.acore.mini.txt
13 -rw-       156   Aug 31 2009 16:14:56 +00:00 f10StkUnit0.acl.acore.mini.txt

flash: 3104256 bytes total (2959872 bytes free)
FTOS#

```

When a member or standby unit crashes, the mini core file gets uploaded to master unit. When the master unit crashes, the mini core file is uploaded to new master. In the MXL Switch, only the master unit has the ability to upload the core dump.

Figure 39-18. Mini core text file example

```

                VALID MAGIC
-----PANIC STRING -----
panic string is :<null>
-----STACK TRACE START-----
0035d60c <f10_save_mmu+0x120>:
00274f8c <panic+0x144>:
0024e2b0 <db_fncall+0x134>:
0024dee8 <db_command+0x258>:
0024d9c4 <db_command_loop+0xc4>:
002522b0 <db_trap+0x158>:
0026a8d0 <mi_switch+0x1b0>:
0026a00c <bpendtsleep>:
-----STACK TRACE END-----

-----FREE MEMORY-----
uvmexp.free = 0x2312

```

The panic string contains key information regarding the crash. Several panic string types exist, and they are displayed in regular English text to allow easier understanding of the crash cause.

TCP Dumps

TCP dump captures CPU bound control plane traffic to improve troubleshooting and system manageability. When enabled, a TCP dump captures all the packets on the local CPU, as specified in the CLI.

You can save the traffic capture files to flash, FTP, SCP, or TFTP. The files saved on the flash are located in the `flash://TCP_DUMP_DIR/Tcpdump_<time_stamp_dir>/` directory, and labeled **tcpdump_*.pcap**. There can be up to 20 `Tcpdump_<time_stamp_dir>` directories. The file after 20 overwrites the oldest saved file. The maximum file size for a TCP dump capture is 1MB. When a file reaches 1MB, a new file is created, up to the specified total number of files.

Maximize the number of packets recorded in a file by specifying the `snap-length` to capture the file headers only.

The `tcpdump` command has a finite run process. When you enable the command, it runs until the `capture-duration` timer and/or the `packet-count` counter threshold is met. If no threshold is set, the system uses a default of five minute `capture-duration` and/or a single 1k file as the stopping point for the dump.

You can use the `capture-duration` timer and the `packet-count` counter at the same time. The TCP dump stops when the first of the thresholds is met. That means that even if the duration timer is 9000 seconds, if the maximum file count parameter is met first, the dumps stop.

Task	Command Syntax	Command Mode
Enable a TCP dump for CPU bound traffic.	<code>tcpdump cp [capture-duration <i>time</i> filter <i>expression</i> max-file-count <i>value</i> packet-count <i>value</i> snap-length <i>value</i> write-to path]</code>	CONFIGURATION

Standards Compliance

This chapter contains the following sections:

- [IEEE Compliance](#)
- [RFC and I-D Compliance](#)
- [MIB Location](#)



Note: Unless noted, when a standard cited here is listed as supported by Dell Force10 operating software (FTOS), FTOS also supports predecessor standards. One way to search for predecessor standards is to use the <http://tools.ietf.org/> website. Click on “**Browse and search IETF documents**”, enter an RFC number, and inspect the top of the resulting document for obsolescence citations to related RFCs.

IEEE Compliance

- 802.1AB — LLDP
- 802.1D — Bridging, STP
- 802.1p — L2 Prioritization
- 802.1Q — VLAN Tagging, Double VLAN Tagging, GVRP
- 802.1s — MSTP
- 802.1w — RSTP
- 802.3ac — Frame Extensions for VLAN Tagging
- 802.3ad — Link Aggregation with LACP
- 802.3ae — 10 Gigabit Ethernet (10GBASE-W, 10GBASE-X)
- 802.3ak — 10 Gigabit Ethernet (10GBASE-CX4)
- 802.3i — Ethernet (10BASE-T)
- 802.3x — Flow Control
- 802.1Qaz — Enhanced Transmission Selection
- 802.1Qbb — Priority-based Flow Control
- ANSI/TIA-1057 — LLDP-MED
- Dell Force10 — FRRP (Force10 Redundant Ring Protocol)
- Dell Force10 — PVST+
- SFF-8431 — SFP+ Direct Attach Cable (10GSFP+Cu)
- MTU — 12,000 bytes

RFC and I-D Compliance

The following standards are supported by FTOS, and are grouped by related protocol. The columns showing support by platform indicate which version of FTOS first supports the standard.

General Internet Protocols

RFC#	Full Name
768	User Datagram Protocol
793	Transmission Control Protocol
854	Telnet Protocol Specification
959	File Transfer Protocol (FTP)
1321	The MD5 Message-Digest Algorithm
1350	The TFTP Protocol (Revision 2)
1661	The Point-to-Point Protocol (PPP)
1989	PPP Link Quality Monitoring
1990	The PPP Multilink Protocol (MP)
1994	PPP Challenge Handshake Authentication Protocol (CHAP)
2474	Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers
2698	A Two Rate Three Color Marker
3164	The BSD syslog Protocol
draft-ietf-bfd-base-03	Bidirectional Forwarding Detection

General IPv4 Protocols

RFC#	Full Name
791	Internet Protocol
792	Internet Control Message Protocol
826	An Ethernet Address Resolution Protocol
1027	Using ARP to Implement Transparent Subnet Gateways
1035	DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION (client)
1042	A Standard for the Transmission of IP Datagrams over IEEE 802 Networks
1191	Path MTU Discovery
1305	Network Time Protocol (Version 3) Specification, Implementation and Analysis
1519	Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy
1542	Clarifications and Extensions for the Bootstrap Protocol
1812	Requirements for IP Version 4 Routers
2131	Dynamic Host Configuration Protocol
2338	Virtual Router Redundancy Protocol (VRRP)
3021	Using 31-Bit Prefixes on IPv4 Point-to-Point Links
3046	DHCP Relay Agent Information Option
3069	VLAN Aggregation for Efficient IP Address Allocation
3128	Protection Against a Variant of the Tiny Fragment Attack

Border Gateway Protocol (BGP)

RFC#	Full Name
1997	BGP Communities Attribute
2385	Protection of BGP Sessions via the TCP MD5 Signature Option
2439	BGP Route Flap Damping
2796	BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)
2842	Capabilities Advertisement with BGP-4
2858	Multiprotocol Extensions for BGP-4
2918	Route Refresh Capability for BGP-4
3065	Autonomous System Confederations for BGP
4360	BGP Extended Communities Attribute
4893	BGP Support for Four-octet AS Number Space
5396	Textual Representation of Autonomous System (AS) Numbers
draft-ietf-idr-bgp4-20	A Border Gateway Protocol 4 (BGP-4)
draft-ietf-idr-restart-06	Graceful Restart Mechanism for BGP

Open Shortest Path First (OSPF)

RFC#	Full Name
1587	The OSPF Not-So-Stubby Area (NSSA) Option
2154	OSPF with Digital Signatures
2328	OSPF Version 2
2370	The OSPF Opaque LSA Option
3623	Graceful OSPF Restart
4222	Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance

Routing Information Protocol (RIP)

RFC#	Full Name
1058	Routing Information Protocol
2453	RIP Version 2

Network Management

RFC#	Full Name
1155	Structure and Identification of Management Information for TCP/IP-based Internets
1156	Management Information Base for Network Management of TCP/IP-based internets
1157	A Simple Network Management Protocol (SNMP)
1212	Concise MIB Definitions
1215	A Convention for Defining Traps for use with the SNMP
1493	Definitions of Managed Objects for Bridges [except for the dot1dTpLearnedEntryDiscards object]
1724	RIP Version 2 MIB Extension
1850	OSPF Version 2 Management Information Base
1901	Introduction to Community-based SNMPv2
2011	SNMPv2 Management Information Base for the Internet Protocol using SMIPv2
2012	SNMPv2 Management Information Base for the Transmission Control Protocol using SMIPv2
2013	SNMPv2 Management Information Base for the User Datagram Protocol using SMIPv2
2024	Definitions of Managed Objects for Data Link Switching using SMIPv2
2096	IP Forwarding Table MIB
2570	Introduction and Applicability Statements for Internet Standard Management Framework
2571	An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks
2572	Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
2574	User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)

Network Management (continued)

RFC#	Full Name
2575	View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)
2576	Coexistence Between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework
2578	Structure of Management Information Version 2 (SMIv2)
2579	Textual Conventions for SMIv2
2580	Conformance Statements for SMIv2
2618	RADIUS Authentication Client MIB, except the following four counters: radiusAuthClientInvalidServerAddresses radiusAuthClientMalformedAccessResponses radiusAuthClientUnknownTypes radiusAuthClientPacketsDropped
2665	Definitions of Managed Objects for the Ethernet-like Interface Types
2674	Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions
2787	Definitions of Managed Objects for the Virtual Router Redundancy Protocol
2819	Remote Network Monitoring Management Information Base: Ethernet Statistics Table, Ethernet History Control Table, Ethernet History Table, Alarm Table, Event Table, Log Table
2863	The Interfaces Group MIB
2865	Remote Authentication Dial In User Service (RADIUS)
3273	Remote Network Monitoring Management Information Base for High Capacity Networks (64 bits): Ethernet Statistics High-Capacity Table, Ethernet History High-Capacity Table
3416	Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)

Network Management (continued)

RFC#	Full Name
3418	Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)
3434	Remote Monitoring MIB Extensions for High Capacity Alarms, High-Capacity Alarm Table (64 bits)
5060	Protocol Independent Multicast MIB
ANSI/TIA-1057	The LLDP Management Information Base extension module for TIA-TR41.4 Media Endpoint Discovery information
draft-grant-taca-cs-02	The TACACS+ Protocol
draft-ietf-idr-bgp4-mib-06	Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIV2
IEEE 802.1AB	Management Information Base module for LLDP configuration, statistics, local system data and remote systems data components.
IEEE 802.1AB	The LLDP Management Information Base extension module for IEEE 802.1 organizationally defined discovery information. (LLDP DOT1 MIB and LLDP DOT3 MIB)
IEEE 802.1AB	The LLDP Management Information Base extension module for IEEE 802.3 organizationally defined discovery information. (LLDP DOT1 MIB and LLDP DOT3 MIB)
ruzin-mstp-mib-02 (Traps)	Definitions of Managed Objects for Bridges with Multiple Spanning Tree Protocol
sFlow.org	sFlow Version 5
sFlow.org	sFlow Version 5 MIB
FORCE10-BGP4-V2-MIB	Force10 BGP MIB (draft-ietf-idr-bgp4-mibv2-05)
FORCE10-IF-EXTENSION-MIB	Force10 Enterprise IF Extension MIB (extends the Interfaces portion of the MIB-2 (RFC 1213) by providing proprietary SNMP OIDs for other counters displayed in the “show interfaces” output)

Network Management (continued)

RFC#	Full Name
FORCE10-LINKAGG-MIB	Force10 Enterprise Link Aggregation MIB
FORCE10-COPY-CONFIG-MIB	Force10 File Copy MIB (supporting SNMP SET operation)
FORCE10-MON-MIB	Force10 Monitoring MIB
FORCE10-PRODUCTS-MIB	Force10 Product Object Identifier MIB
FORCE10-SS-CHASSIS-MIB	Force10 S-Series Enterprise Chassis MIB
FORCE10-SMI	Force10 Structure of Management Information
FORCE10-SYSTEM-COMPONENT-MIB	Force10 System Component MIB (enables the user to view CAM usage information)
FORCE10-TC-MIB	Force10 Textual Convention
FORCE10-TRAP-ALARM-MIB	Force10 Trap Alarm MIB
FORCE10-FIP-SNOOPING-MIB	Force10 FIP Snooping MIB (Based on T11-FCoE-MIB mentioned in FC-BB-5)
FORCE10-DCB-MIB	Force10 DCB MIB
IEEE 802.1Qaz	Management Information Base extension module for IEEE 802.1 organizationally defined discovery information (LDP-EXT-DOT1-DCBX-MIB)
IEEE 802.1Qbb	Priority-based Flow Control module for managing IEEE 802.1Qbb

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, contact Dell Force10 TAC for assistance.

Index

Numerics

- 10/100/1000 Base-T Ethernet line card, auto negotiation 250
- 100/1000 Ethernet interfaces
 - port channels 232
- 802.1AB 673
- 802.1D 673
- 802.1p 673
- 802.1p/Q 673
- 802.1Q 673
- 802.1s 673
- 802.1w 673
- 802.1X 673
- 802.3ac 673
- 802.3ad 673
- 802.3ae 673
- 802.3af 673
- 802.3ak 673
- 802.3i 673
- 802.3u 673
- 802.3x 673
- 802.3z 673

A

- AAA (Accounting, Authentication, and Authorization)
 - security model 473
- AAA Accounting 473
 - command 474
 - suppress null-username command 474
- AAA Authentication
 - authentication and authorization, local by default 479
 - command 478
 - configuring 477
 - enable 478
 - enable command 478
 - enable method 477
 - line method 477
 - local method 477
 - none method 477
 - radius method 477
 - tacacs+ 477
- AAA Authorization
 - AAA new-model enabled by default 479
- ABR
 - definition 371
- access-class 492
- ACL
 - definition 71
 - IP ACL definition 71
 - RADIUS 485

- ANSI/TIA-1057 316
- Area Border Router. *See* ABR.
- Authentication
 - implementation 476
 - TACACS+ 491
 - VTY 500
- Authorization
 - TACACS+ 491
 - VTY 500
- auto negotiation 250, 252
- auto negotiation, line card 250
- Auto-command 486

B

- Bare Metal Provisioning 101
- base VLAN 392
- BPDU 466
- Bridge MIB
 - STP implementation 578
- Bridge Protocol Data Units. *See* BPDU.

C

- CAM Profiling, When to Use 109
- CLI
 - case sensitivity 32
 - editing commands 32
 - partial keywords 32
- CLI Modes
 - LINE 26
- community port 392
- community VLAN 392
- coredumps 669
- crypto key generate 495

D

- debug ip ssh 495
- Default VLAN
 - changing the VLAN id 624
 - implementation 624
 - Layer 2 mode 624
 - remove interface 630
 - remove untagged interface 625
 - untagged interfaces 624, 627
- directed broadcast 263
- disable-on-sfm failure 246
- display parameter 34
- DNS 264
- Document conventions 24

E

extended IP ACL 72

F

Fast Convergence after MSTP-Triggered Topology Changes 218

fast-convergence

OSPF 374

File Transfer Protocol. *See* FTP.

flowcontrol 248

forward delay 467, 583

FTOS 362

FTP 61

configuring client parameters 63

configuring server parameters 62

enabling server 62

using VLANs 61

G

GARP VLAN Registration Protocol (GVRP) 205

grep option 33

grep pipe option 398

GVRP (GARP VLAN Registration Protocol) 205

H

hello time 467, 583

host port 392

Hot Lock ACL 72

Hybrid ports 629

hybrid ports 630

I

I-D (Internet Draft) Compliance 674

Idle Time 485

IEEE 802.1q 205

IEEE 802.1Q tag 625

IEEE Compliance 673

IEEE Standard 802.3ad 231

implicit deny 71

Interface modes

Layer 2 224

Layer 3 224

Interface Range Macros 241

Interface types

100/1000 Ethernet 219, 224

10-Gigabit Ethernet 219, 224

1-Gigabit Ethernet 219, 224

Loopback 224

management 224

Management Ethernet interface 223

Port Channel 224

VLAN 224

interface types

null interface 224

interfaces

auto negotiation setting 250, 252

clearing counters 258

commands allowed when part of a port channel 233

configuring secondary IP addresses 261

determining configuration 225

member of port channel 236

viewing Layer 3 interfaces 222

viewing only Layer 2 interfaces 254

Inter-VLAN routing 229

considerations 229

IP ACLs

applying IP ACL to interface 82

configuring extended IP ACL 79

configuring filter without sequence number 80

configuring standard IP ACL 76, 77

deleting a filter 77, 78

extended IP ACLs 72, 78

standard IP ACL 72

types 72

viewing configuration 76

IP addresses

assigning IP address to interface 226

assigning to interface 260

assigning to port channel 238

assigning to VLAN 630

composition 259

configuring static routes 261

IP fragmentation 246

ip local-proxy-arp command 398

IP prefix lists

"permit all" statement 87

applying to OSPF routes 89

applying to RIP routes 89

configuring filter 86

configuring filters without seq command 87

definition 85

deleting filter 87, 88

implementation 86

permit default route only 87

rules 86

using the le and ge parameters 85

IP routing

VLANs 625

ip scp topdir 495

ip ssh authentication-retries 495

- ip ssh connection-rate-limit 495
- ip ssh hostbased-authentication enable 496
- ip ssh password-authentication enable 496
- ip ssh pub-key-file 496
- ip ssh rhostsfile 496
- ip ssh rsa-authentication 496
- ip ssh rsa-authentication enable 496
- ip ssh server command 494
- IP version 4 259
- isolated port 392
- isolated VLAN 392

L

- LAG. See Port Channels.
- Layer 2 mode
 - configuring 225
- Layer 2 protocols
 - configuring 224
- Layer 3 mode
 - enable traffic 225
- Layer 3 protocols
 - configuring 225
- line card, auto negotiation 250
- Link Aggregation Group 231
- Link Debounce Timer 247
- Link Layer Discovery Protocol (LLDP) 313
- link MTU
 - configuring 248
- Link State Advertisements. *See LSAs.*
- LLDP 313
- LLDP-MED 316
- logging
 - consolidating messages 60
 - including timestamp 61
 - UNIX system logging facility 59
- Loopback interface
 - configuring 230
 - defaults 224
 - definition 230
 - deleting interface 230
 - viewing configuration 230
- LSAs 355
 - AS Boundary 362
 - AS External 362
 - Network 362
 - Network Summary 362
 - NSSA External 362
 - Opaque Area-local 362
 - Opaque Link-local 363
 - Router 362
 - types supported 362

M

- management interface 224
 - configuring a management interface 227
 - configuring IP address 227
 - definition 227
- management interface, switch 223
- max age 467
- MIB Location 682
- minimum oper up links in a port channel 237
- mirror, port 385, 609
 - remote port mirroring 610
- monitor interfaces 243
- MTU
 - configuring MTU values for Port Channels 249
 - configuring MTU values for VLANs 249
 - definition 246
 - link MTU
 - configuring 248
- MTU Size, Configuring on an Interface 248

N

- NIC teaming 309
- no-more 34
- no-more parameter 34
- NTP
 - configuring authentication 600
 - configuring source address 600
- null 224
- null interface
 - available command 230
 - definition 230
 - entering the interface 230
 - information 224

O

- Open Shortest Path First 355
- OSFP Adjacency with Cisco Routers 365
- OSPF 355
 - backbone area 369
 - changing authentication parameters 377
 - changing interface parameters 376
 - configuring a passive interface 373
 - configuring a stub area 371
 - configuring network command 369
 - debugging OSPF 382
 - default 366
 - disabling OSPF 367, 369
 - enabling routing 367
 - LSA throttling 361
 - redistributing routes 378, 379

- restarting OSPF 367, 369
- router ID 370
- using loopback interfaces 371
- using prefix lists 378
- viewing configuration of neighboring router 381
- viewing interface areas 370

P

- passwords
 - configuring password 481
- port channel
 - definition 231
- port channel (LAG), configure 233
- port channel, minimum oper up links 237
- Port Channels
 - configuring MTU values 249
 - member of VLANs 626
- Port channels
 - benefits 231
 - defaults 224
- port channels
 - adding physical interface 233, 234
 - assigning IP address 238
 - commands allowed on individual interfaces 233
 - configuring 233
 - containing 100/1000 and GE interfaces 232
 - IP routing 238
 - placing in Layer 2 mode 233
 - reassigning interfaces 236
- port cost 467
- port mirror 385, 609
 - remote 610
- Port Monitoring Commands
 - Important Points to Remember 385
- port priority 467, 583
- port types (private VLAN) 392
- port-based VLANs 625
 - assigning IP address 630
 - benefits 625
 - creating VLAN 626
 - definition 625
 - deleting VLAN 627
 - enabling tagging 628
 - interface requirements 625
 - IP routing 625
 - number supported 625
 - remove interface 629
 - remove tagged interface 629
 - tag frames 628
 - tagged interface
 - member of multiple VLANs 629

- Portfast 469, 584
- Prefix list. *See IP Prefix list.*
- primary VLAN 392
- Private VLAN (PVLAN) 391
- private-vlan mapping secondary-vlan command 393
- Privilege Level 486
- privilege levels
 - and CLI commands 480
 - definition 479
 - number of levels available 479
 - privilege level 0 definition 479
 - privilege level 1 definition 479
 - privilege level 15 definition 479
- promiscuous port 392
- Proxy ARP
 - default 268

Q

- QoS
 - dot1p-priority values 418
 - purpose of input policies 423
 - rate limit outgoing traffic 420
- QSFP port splitting 245

R

- RADIUS
 - changing an optional parameter 488
 - configuration requirements 485
 - configuring global communication parameter 488
 - debugging RADIUS transactions 489, 491
 - definition 484
 - deleting a server host 488
 - specifying a server host 488, 492
 - viewing RADIUS configuration 489
- RADIUS authentication 479
- RADIUS Authentication and Authorization 485
- radius-server host command 478
- rate-interval command 255
- remote port mirroring 610
- RFC 1058 435
- RFC 1858 493
- RFC 2138 485
- RFC 2338 634
- RFC 2453 435
- RFC 3128 493
- RFC 791 259
- RFC 959 61
- RFC Compliance 674
- RIP
 - adding routes 440

- auto summarization default 436
- changing RIP version 440
- configuring interfaces to run RIP 438
- debugging RIP 444
- default values 436
- default version 437
- disabling RIP 438
- ECMP paths supported 436
- enabling RIP 437
- route information 439
- setting route metrics 443
- summarizing routes 443
- timer values 436
- version 1 description 435
- version default on interfaces 436
- RIPv1 435
- RIPv2 436
- root bridge 466, 583
- route maps
 - configuring match commands 96
 - configuring set commands 96
 - creating 93
 - creating multiple instances 94
 - default action 93
 - definition 92
 - deleting 94
 - implementation 92
 - implicit deny 92
 - redistributing routes 97
 - tagging routes 97
- RSA 496

S

- SCP 494
- SCP/SSH server 494
- searching show commands 34
 - display 34
 - grep 34
- secondary VLAN 392
- Secure Shell (SSH) 494
- show accounting command 475
- show arp command 399
- show crypto 496
- show hardware commands (S60) 653
- show interfaces command 255
- show interfaces switchport command 254
- show ip protocols command 445, 448
- show ip rip database command 445, 448
- show ip route command 445, 448
- show ip ssh client-pub-keys 496

- show ip ssh command 494
- show ip ssh rsa-authentication 496
- show vlan command 399
- Spanning Tree group. *See* STG.
- SSH 494
 - debug 495
 - display 494
 - host-keys 496
 - ssh command 494
- SSHv2 server 496
- standard IP ACL 72
- static route 261
- STG
 - changing parameters 467, 583
 - default 467
 - port cost 467
 - root bridge 466, 583
- STP
 - benefits 577
 - bridge priority 470, 587, 588, 589
 - default 467
 - definition 577
 - disabling STP 463, 580
 - forward delay 467, 583
 - hello time 467, 583
 - interfaces 464, 580
 - max age 467
 - port cost 468, 584
 - port ID 578
 - port priority 467, 468, 583, 584
 - Portfast 469, 584
 - root bridge 470, 587, 588, 589
- switchport mode private-vlan command 394

T

- TACACS+ 489
 - authentication 491
 - authorization 491
 - deleting a server host 493
 - selecting TACACS+ as a login authentication method 490
 - servers and access classes 491
- tacacs-server host command 478
- Tag Control Information 626
- Tag Header 626
 - definition 625
 - Ethernet frame 626
- tagged interfaces
 - member of multiple VLANs 629

TCP Tiny and Overlapping Fragment Attack, Protection Against 493
 TDR (Time Domain Reflectometer) 244
 Telnet 492
 Telnet Daemon, Enabling and Disabling 500
 Time Domain Reflectometer (TDR) 244
 Time to Live (TTL) 329
 trunk port 392
 TTL 329

U

user level
 definition 479
 user name
 configuring user name 481
 username command 482

V

virtual IP addresses 637
 Virtual LANs. *See* VLAN.
 Virtual Router Identifier. *See* VRID.
 Virtual Router Redundancy Protocol. *See* VRRP.
 VLAN configuration, automatic 205
 VLAN Protocol Identifier 626
 VLAN types 392
 VLAN types (private VLAN) 391
 VLANs 224, 623
 adding a port channel 237
 adding interface 625
 assigning IP address 630
 benefits 623
 configuring MTU values 249
 defaults 624
 definition 623
 enabling tagging 628
 FTP 61, 630
 hybrid ports 629
 IP routing 229, 625
 Layer 2 interfaces 627
 port-based 625
 removing a port channel 237
 removing tagged interface 625, 629
 SNMP 630
 tagged interfaces 626, 627
 TFTP 630
 untagged interfaces 627
 viewing configured 627
 VLSM 259
 VLSM (Variable Length Subnet Masks) 435
 VRRP 633

advertisement interval 642
 benefits 635
 changing advertisement interval 642
 configuring priority 640
 configuring simple authentication 641
 definition 633
 disabling preempt 642
 MAC address 633
 monitoring interface 644
 simple authentication 640
 transmitting VRRP packets 637
 virtual IP addresses 637
 virtual router 636
 VRID 633, 636
 VTY lines
 access class configuration 500
 access classes and TACACS+ servers 491
 assigning access classes by username 501
 deny all, deny incoming subnet access-class application 501
 deny10 ACLs, support for remote authentication and authorization 492
 line authentication, support for 501
 local authentication and authorization 500
 local authentication and authorization, local database source of access class 500
 radius authentication, support for 501
 remote authentication and authorization 491
 remote authentication and authorization, 10.0.0.0 subnets 492
 remote authentication and local authorization 501
 TACACS+ authentication, support for local authorization 501

W

When to Use CAM Profiling 109

