

# FASTPATH 8.7

- ▶ Document Revision 1.2
- ▶ Date: December 2021

## ▶ FASTPATH 8.7 - Configuration Guide

### Disclaimer

Kontron would like to point out that the information contained in this manual may be subject to alteration, particularly as a result of the constant upgrading of Kontron products. This document does not entail any guarantee on the part of Kontron with respect to technical processes described in the manual or any product characteristics set out in the manual. Kontron assumes no responsibility or liability for the use of the described product(s), conveys no license or title under any patent, copyright or mask work rights to these products and makes no representations or warranties that these products are free from patent, copyright or mask work right infringement unless otherwise specified. Applications that are described in this manual are for illustration purposes only. Kontron makes no representation or warranty that such application will be suitable for the specified use without further testing or modification. Kontron expressly informs the user that this manual only contains a general description of processes and instructions which may not be applicable in every individual case. In cases of doubt, please contact Kontron.

This manual is protected by copyright. All rights are reserved by Kontron. No part of this document may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), without the express written permission of Kontron. Kontron points out that the information contained in this manual is constantly being updated in line with the technical alterations and improvements made by Kontron to the products and thus this manual only reflects the technical status of the products by Kontron at the time of publishing.

Brand and product names are trademarks or registered trademarks of their respective owners.

© 2021 by Kontron S&T AG

Kontron S&T AG

Lise-Meitner-Straße 3-5  
86156 Augsburg  
Germany  
[www.kontron.com](http://www.kontron.com)



## Revision History

Rev. Index	Brief Description of Changes	Date of Issue
0.9	Draft version	2021-04-16
1.0	Released version	2021-04-28
1.1	Removed chapter „Booting the switch“	2021-04-28
1.2	Rework for FASTPATH 8.7	2021-12-08

## Intended Use

THIS DEVICE AND ASSOCIATED SOFTWARE ARE NOT DESIGNED, MANUFACTURED OR INTENDED FOR USE OR RESALE FOR THE OPERATION OF NUCLEAR FACILITIES, THE NAVIGATION, CONTROL OR COMMUNICATION SYSTEMS FOR AIRCRAFT OR OTHER TRANSPORTATION, AIR TRAFFIC CONTROL, LIFE SUPPORT OR LIFE SUSTAINING APPLICATIONS, WEAPONS SYSTEMS, OR ANY OTHER APPLICATION IN A HAZARDOUS ENVIRONMENT, OR REQUIRING FAIL-SAFE PERFORMANCE, OR IN WHICH THE FAILURE OF PRODUCTS COULD LEAD DIRECTLY TO DEATH, PERSONAL INJURY, OR SEVERE PHYSICAL OR ENVIRONMENTAL DAMAGE (COLLECTIVELY, "HIGH RISK APPLICATIONS").

You understand and agree that your use of Kontron devices as a component in High Risk Applications is entirely at your risk. To minimize the risks associated with your products and applications, you should provide adequate design and operating safeguards. You are solely responsible for compliance with all legal, regulatory, safety, and security related requirements concerning your products. You are responsible to ensure that your systems (and any Kontron hardware or software components incorporated in your systems) meet all applicable requirements. Unless otherwise stated in the product documentation, the Kontron device is not provided with error-tolerance capabilities and cannot therefore be deemed as being engineered, manufactured or setup to be compliant for implementation or for resale as device in High Risk Applications. All application and safety related information in this document (including application descriptions, suggested safety measures, suggested Kontron products, and other materials) is provided for reference only.

## Customer Support

Find Kontron contacts by visiting: <http://www.kontron.com/support>.

## Customer Service

As a trusted technology innovator and global solutions provider, Kontron extends its embedded market strengths into a services portfolio allowing companies to break the barriers of traditional product lifecycles. Proven product expertise coupled with collaborative and highly-experienced support enables Kontron to provide exceptional peace of mind to build and maintain successful products.

[www.kontron.com](http://www.kontron.com)

For more details on Kontron's service offerings such as: enhanced repair services, extended warranty, Kontron training academy, and more visit <http://www.kontron.com/support-and-services/services>.

## Customer Comments

If you have any difficulties using this CLI Reference Manual, discover an error, or just want to provide some feedback, please send a message to Kontron. Detail any errors you find. We will correct the errors or problems as soon as possible and post the revised Reference Manual on our website.

## Terms and Conditions

Kontron warrants products in accordance with defined regional warranty periods. For more information about warranty compliance and conformity, and the warranty period in your region, visit <http://www.kontron.com/terms-and-conditions>.

Kontron sells products worldwide and declares regional General Terms & Conditions of Sale, and Purchase Order Terms & Conditions. Visit <http://www.kontron.com/terms-and-conditions>.

For contact information, refer to the corporate offices contact information on the last page of this user guide or visit our website CONTACT US.

## Table of Contents

1/ Introduction .....	15
1.1 About This Document .....	15
1.1.1 Document Conventions .....	15
1.2 About FASTPATH Software Modules .....	15
1.3 Configuring Data Center and BGP Features .....	16
2/ Getting Started .....	16
2.1 Connecting the Switch to the Network .....	16
2.2 Understanding the User Interfaces .....	17
2.2.1 Using the Web Interface .....	18
2.2.2 Using the Command-Line Interface .....	23
3/ Getting Started with Stacking .....	24
3.1 Understanding Switch Stacks .....	24
3.1.1 Switch Stack Membership .....	24
3.1.2 Stack Manager Election and Re-Election .....	25
3.1.3 Stack Member Numbers .....	25
3.1.4 Stack Member Priority Values .....	25
3.2 Switch Stack Software Compatibility Recommendations .....	25
3.3 Incompatible Software and Stack Member Image Upgrades .....	26
3.4 Switch Stack Configuration Files .....	26
3.5 Switch Stack Management Connectivity .....	26
3.5.1 Connectivity to the Switch Stack Through Console Ports .....	26
3.5.2 Connectivity to the Switch Stack Through Telnet .....	26
3.6 General Practices .....	26
3.7 Initial Installation and Power-up of a Stack .....	27
3.8 Removing a Unit from the Stack .....	27
3.9 Adding a Unit to an Operating Stack .....	27
3.10 Replacing the Stack Member with a New Unit .....	28
3.11 Renumbering Stack Members .....	28
3.12 Moving a Manager to a Different Unit in the Stack .....	29
3.13 Removing a Manager Unit from an Operating Stack .....	29
3.14 Initiating a Warm Failover of the Manager Unit .....	29
3.15 Merging Two Operational Stacks .....	29
3.16 Preconfiguration .....	29
3.17 Stack Links .....	30
4/ Configuring System Information .....	31
4.1 Viewing the Dashboard .....	31
4.2 Viewing ARP Cache .....	33
4.3 Viewing Inventory Information .....	34
4.4 Viewing the System Firmware Status .....	35
4.4.1 Dual Image Status .....	35
4.4.2 Dual Image Configuration and Upgrade .....	35
4.4.3 AutoInstall .....	37
4.5 Viewing System Resources .....	38
4.6 Selecting the SDM Template .....	40

4.7	Defining General Device Information .....	41
4.7.1	System Description .....	41
4.7.2	Switch Configuration .....	43
4.7.3	IP Address Conflict Detection .....	43
4.7.4	IPv4 Network Connectivity Configuration .....	44
4.7.5	IPv6 Network Connectivity .....	46
4.7.6	Network Port IPv6 Neighbors .....	47
4.7.7	Service Port IPv4 .....	48
4.7.8	Service Port IPv6 .....	49
4.7.9	Service Port IPv6 Neighbors .....	51
4.7.10	DHCP Client Options .....	52
4.7.11	System Connectivity .....	52
4.7.12	Telnet Session .....	53
4.7.13	Outbound Telnet Configuration .....	54
4.7.14	Serial Port .....	55
4.7.15	CLI Banner Configuration .....	56
4.7.16	HTTP Configuration .....	57
4.7.17	HTTPS Configuration .....	57
4.7.18	SSH Configuration .....	59
4.7.19	Management Access Control and Administration List .....	60
4.7.20	User Accounts .....	62
4.7.21	Authentication Server Users .....	66
4.7.22	User Domain Name .....	68
4.7.23	Accounting List Configuration .....	71
4.7.24	Accounting List Configuration .....	73
4.7.25	Authentication List Summary .....	74
4.7.26	Select Authentication List .....	75
4.7.27	Authorization List Configuration .....	77
4.7.28	Line Password .....	79
4.7.29	Enable Password .....	80
4.7.30	Password Rules .....	81
4.7.31	Last Password Result .....	82
4.7.32	Denial of Service .....	83
4.8	Configuring and Searching the Forwarding Database .....	85
4.8.1	Switch Configuration .....	85
4.9	Managing Logs .....	86
4.9.1	Log Configuration .....	86
4.9.2	Buffered Log .....	88
4.9.3	Event Log .....	89
4.9.4	Hosts Log Configuration .....	90
4.9.5	Syslog Source Interface Configuration .....	92
4.9.6	Persistent Log .....	92
4.10	Configuring Email Alerts .....	93
4.10.1	Email Alert Global Configuration .....	93
4.10.2	Email Alerts Server Configuration .....	95
4.10.3	Email Alert Statistics .....	96
4.10.4	Email Alert Subject Configuration .....	96
4.10.5	Email Alerts To Address Configuration .....	97
4.11	Configuring and Viewing Device Slot Information .....	98
4.11.1	Slot Card Configuration .....	98
4.11.2	Slot Supported Cards .....	100
4.12	Configuring Power Over Ethernet (PoE) and PoE Statistics .....	101
4.12.1	PoE Configuration .....	101
4.12.2	PoE Port Configuration .....	102
4.12.3	PoE Port Statistics .....	104
4.13	Viewing Device Port Information .....	105
4.13.1	Port Summary .....	105

4.13.2	Port Description .....	107
4.13.3	Port Cable Test .....	108
4.13.4	Mirroring .....	109
4.13.5	Mirroring Summary .....	113
4.13.6	Expandable Ports .....	115
4.13.7	Green Mode Statistics .....	116
4.13.8	Green Ethernet EEE Interface History Table .....	116
4.14	Configuring sFlow .....	117
4.14.1	sFlow Agent Summary .....	117
4.14.2	sFlow Receiver Configuration .....	118
4.14.3	sFlow Poller Configuration .....	119
4.14.4	sFlow Sampler Configuration .....	121
4.14.5	sFlow Source Interface Configuration .....	122
4.15	Defining SNMP Parameters .....	123
4.15.1	SNMP v1 and v2 .....	123
4.15.2	SNMP v3 .....	124
4.15.3	SNMP Community Configuration .....	124
4.15.4	Trap Receiver v1/v2 Configuration .....	126
4.15.5	Trap Receiver v3 Configuration .....	127
4.15.6	SNMP Supported MIBs .....	128
4.15.7	SNMP Access Control Group .....	129
4.15.8	SNMP User Security Model .....	131
4.15.9	SNMP Source Interface Configuration .....	132
4.16	Viewing System Statistics .....	133
4.16.1	Switch Statistics .....	133
4.16.2	Port Summary .....	134
4.16.3	Port Detailed Statistics .....	135
4.16.4	Port DHCPv6 Client Statistics .....	139
4.16.5	Time Based Group Statistics .....	140
4.16.6	Time Based Flow Statistics .....	141
4.16.7	Time Based Statistics .....	143
4.17	Using System Utilities .....	144
4.17.1	System Reset .....	144
4.17.2	Ping .....	144
4.17.3	Ping IPv6 .....	145
4.17.4	TraceRoute .....	147
4.17.5	IP Address Conflict Detection .....	148
4.17.6	File Transfer .....	149
4.17.7	Digital Signature Verification .....	153
4.17.8	Core Dump .....	153
4.17.9	Core Dump Test .....	155
4.18	Managing SNMP Traps .....	155
4.18.1	System Trap Log .....	155
4.18.2	System Trap Flags .....	156
4.19	Managing the DHCP Server .....	157
4.19.1	DHCP Server Global Configuration .....	157
4.19.2	DHCP Server Excluded Addresses .....	158
4.19.3	DHCP Server Pool Summary .....	159
4.19.4	DHCP Server Pool Configuration .....	162
4.19.5	DHCP Server Pool Options .....	164
4.19.6	DHCP Server Bindings .....	167
4.19.7	DHCP Server Statistics .....	167
4.19.8	DHCP Server Conflicts Information .....	169
4.20	Configuring Time Ranges .....	170
4.20.1	Time Range Summary .....	170
4.20.2	Time Range Entry Summary .....	170
4.21	Configuring DNS .....	172

4.21.1	Global Configuration .....	172
4.21.2	DNS IP Mapping Configuration .....	173
4.21.3	DNS Source Interface Configuration .....	175
4.22	Configuring SNTP Settings .....	175
4.22.1	SNTP Global Configuration .....	176
4.22.2	SNTP Global Status .....	177
4.22.3	SNTP Server Configuration .....	178
4.22.4	SNTP Server Status .....	179
4.22.5	SNTP Source Interface Configuration .....	180
4.23	Configuring the Time Zone .....	181
4.23.1	Time Zone Configuration .....	182
4.23.2	Summer Time Configuration .....	183
4.24	Configuring and Viewing ISDP Information .....	184
4.24.1	ISDP Global Configuration .....	185
4.24.2	ISDP Cache Table .....	186
4.24.3	ISDP Interface Configuration .....	186
4.24.4	Statistics .....	187
4.25	Link Dependency .....	188
4.25.1	Link Dependency Group Status .....	188
4.26	Link Local Protocol Filtering .....	189
4.26.1	LLPF Interface Configuration .....	190
5/	Configuring Switching Information .....	191
5.1	Managing VLANs .....	191
5.1.1	VLAN Status .....	191
5.1.2	VLAN Port Configuration .....	193
5.1.3	VLAN Port Summary .....	194
5.1.4	Switchport Summary .....	195
5.1.5	VLAN Internal Usage .....	197
5.1.6	Configure VLAN Statistics .....	198
5.1.7	Reset VLAN Configuration .....	199
5.1.8	RSPAN Configuration .....	199
5.2	Configuring UDLD .....	200
5.2.1	UDLD Interface Configuration .....	201
5.3	MAC Based VLAN Status .....	202
5.4	Double VLAN (DVLAN) Tunneling .....	203
5.4.1	DVLAN Configuration .....	203
5.4.2	DVLAN Summary .....	204
5.4.3	DVLAN Interface Summary .....	205
5.5	IP Subnet Based VLAN .....	206
5.6	Protocol Based VLAN Configuration .....	207
5.6.1	Status .....	207
5.6.2	Configuration .....	208
5.7	Private VLAN .....	209
5.7.1	Private VLAN Configuration .....	209
5.7.2	Private VLAN Association .....	210
5.7.3	Private VLAN Interface .....	211
5.8	Voice VLAN Configuration .....	213
5.9	Voice VLAN Interface .....	214
5.10	Virtual Port Channel Configuration .....	215
5.10.1	Interface Configuration .....	218
5.10.2	Statistics .....	220
5.11	Port Auto Recovery .....	222
5.11.1	Port Auto Recovery Configuration .....	222
5.12	Creating MAC Filters .....	225

5.12.1	MAC Filter Configuration .....	225
5.13	Configuring Dynamic ARP Inspection .....	226
5.13.1	DAI Configuration .....	227
5.13.2	DAI VLAN Configuration .....	227
5.13.3	DAI Interface Configuration .....	228
5.13.4	DAI ARP ACL Configuration .....	229
5.13.5	Add Access Control List .....	230
5.13.6	Add ACL Rule Configuration .....	231
5.13.7	DAI ACL Summary .....	232
5.13.8	DAI Statistics .....	232
5.14	GARP Configuration .....	234
5.14.1	Switch Configuration .....	234
5.14.2	Port Configuration .....	234
5.15	Configuring DHCP Snooping .....	236
5.15.1	Global DHCP Snooping Configuration .....	236
5.15.2	DHCP Snooping VLAN Configuration .....	236
5.15.3	DHCP Snooping Interface Configuration .....	237
5.15.4	DHCP Snooping Static Bindings .....	238
5.15.5	DHCP Snooping Dynamic Bindings .....	239
5.15.6	DHCP Snooping Persistent Configuration .....	240
5.15.7	DHCP Snooping Statistics .....	240
5.15.8	DHCP L2 Relay Global Configuration .....	241
5.15.9	DHCP L2 Relay Interface Configuration .....	242
5.15.10	DHCP L2 Relay VLAN Configuration .....	243
5.15.11	DHCP L2 Relay Interface Statistics .....	243
5.15.12	DHCP Snooping IP Source Guard Interface Configuration .....	244
5.15.13	DHCP Snooping IP Source Guard Bindings .....	245
5.16	Configuring IGMP Snooping .....	246
5.16.1	Global Configuration and Status .....	246
5.16.2	Interface Configuration .....	247
5.16.3	Source Specific Multicast .....	248
5.16.4	VLAN Status .....	249
5.16.5	Multicast Router Configuration .....	250
5.16.6	Multicast Router VLAN Status .....	251
5.16.7	Multicast Router VLAN Configuration .....	251
5.17	Configuring IGMP Snooping Querier .....	252
5.17.1	Configuration .....	252
5.17.2	VLAN Configuration .....	253
5.17.3	VLAN Status .....	254
5.18	Configuring MLD Snooping .....	254
5.18.1	Global Configuration and Status .....	255
5.18.2	Interface Configuration .....	255
5.18.3	Source Specific Multicast .....	256
5.18.4	VLAN Status .....	257
5.18.5	Multicast Router Configuration .....	258
5.18.6	Multicast Router VLAN Status .....	259
5.19	Configuring MLD Snooping Querier .....	260
5.19.1	Configuration .....	260
5.19.2	VLAN Configuration .....	260
5.19.3	VLAN Status .....	261
5.20	Creating Port Channels .....	262
5.20.1	Port Channel Summary .....	263
5.20.2	Port Channel Configuration .....	264
5.20.3	Port Channel Statistics .....	265
5.21	Viewing Multicast Forwarding Database Information .....	266
5.21.1	MFDB Table .....	266
5.21.2	GMRP Table .....	267

5.21.3	IGMP Snooping Table .....	268
5.21.4	MLD Snooping Table .....	269
5.21.5	Source Specific Multicast .....	269
5.21.6	Source Specific Multicast Status .....	270
5.21.7	MFDB Statistics .....	271
5.22	Multicast VLAN Registration .....	271
5.22.1	MVR Global Configuration .....	271
5.22.2	MVR Group Status .....	272
5.22.3	MVR Interface Status .....	273
5.22.4	MVR Statistics .....	274
5.23	Configuring Protected Ports .....	275
5.24	Priority Flow Control .....	276
5.24.1	Priority Flow Control Configuration .....	276
5.24.2	Priority Flow Control Statistics .....	277
5.25	Configuring Spanning Tree Protocol .....	278
5.25.1	Switch Configuration/Status .....	279
5.25.2	CST Configuration .....	280
5.25.3	CST Port Configuration .....	281
5.25.4	MST Configuration .....	283
5.25.5	MST Port Configuration .....	284
5.25.6	Spanning Tree Statistics .....	286
5.25.7	PVST Global .....	287
5.25.8	PVST VLAN .....	287
5.25.9	PVST Interface .....	289
5.25.10	PVST Statistics .....	290
5.26	Mapping 802.1p Priority .....	291
5.27	Configuring Port Security .....	292
5.27.1	Port Security Administration .....	293
5.27.2	Port Security Interface Configuration .....	293
5.27.3	VLAN MAC Locking .....	294
5.27.4	Port Security Statically Configured MAC Addresses .....	295
5.27.5	Port Security Dynamically Learned MAC Addresses .....	296
5.28	Managing LLDP .....	296
5.28.1	Global Configuration .....	296
5.28.2	LLDP Interface Configuration .....	297
5.28.3	Local Devices .....	299
5.28.4	Remote Devices .....	300
5.28.5	Statistics .....	301
5.28.6	LLDP-MED .....	303
5.29	Loop Protection .....	307
5.29.1	Loop Protection Configuration .....	307
5.30	IEEE 802.1ag Connectivity Fault Management (CFM) .....	309
5.30.1	Dot1ag Global Configuration .....	309
5.30.2	Dot1ag Maintenance Domain (MD) Configuration .....	310
5.30.3	Dot1ag Maintenance Association (MA) Configuration .....	311
5.30.4	Dot1ag Maintenance Association End-Point (MEP) Configuration .....	312
5.30.5	Dot1ag Remote Maintenance Association End-Point (RMEP) Configuration .....	313
5.30.6	Dot1ag Maintenance Intermediate Point (MIP) Configuration .....	315
5.30.7	Dot1ag Remote Maintenance Association End-Point Summary .....	316
5.30.8	Dot1ag L2 Ping .....	316
5.30.9	Dot1ag L2 Traceroute .....	317
5.30.10	Dot1ag L2 Traceroute Cache .....	318
5.30.11	Dot1ag Statistics .....	319
5.31	IEEE 802.3ah Ethernet in the First Mile .....	320
5.31.1	Dot3ah Configuration .....	320
5.31.2	Dot3ah Link Monitor Configuration .....	321
5.31.3	Dot3ah Remote Loopback Configuration .....	323



5.32	Data Center Features .....	323
5.32.1	Data Center Bridging Exchange Protocol .....	324
5.32.2	FIP Snooping .....	324
5.32.3	Congestion Notification (Qau) .....	324
5.32.4	Enhanced Transmission Selection .....	325
5.32.5	OpenFlow .....	325
5.33	802.1AS .....	325
5.33.1	802.1AS Configuration .....	325
5.33.2	802.1AS Port Summary .....	326
5.33.3	802.1AS Statistics .....	328
5.34	Multiple Registration Protocol Configuration .....	330
5.34.1	MRP Configuration .....	330
5.34.2	MRP Interface Configuration .....	332
5.34.3	Qav Mapping .....	333
5.34.4	Qav Parameters .....	333
5.34.5	MSRP Reservation Parameters .....	334
5.34.6	MSRP Streams Information .....	336
5.34.7	MSRP Statistics .....	337
5.34.8	MMRP Statistics .....	338
5.34.9	MVRP Statistics .....	339
5.35	IP Device Tracking .....	340
5.35.1	Device Tracking Global Configuration .....	340
5.35.2	Device Tracking Summary .....	341
5.35.3	Device Tracking Interface Configuration .....	342
5.35.4	Device Tracking Statistics .....	343
6/	Configuring Routing .....	344
6.1	Configuring ARP .....	344
6.1.1	ARP Create .....	344
6.1.2	ARP Table Configuration .....	345
6.2	Configuring Global IP Settings .....	346
6.2.1	Configuration .....	346
6.2.2	Interface Summary .....	349
6.2.3	Interface Configuration .....	351
6.2.4	IP Loopback Configuration .....	353
6.2.5	IP Statistics .....	354
6.3	Router .....	356
6.3.1	Route Table .....	356
6.3.2	Configured Routes .....	357
6.3.3	Summary .....	359
6.3.4	ECMP Group .....	362
6.4	Configuring IPv6 Settings .....	362
6.4.1	IPv6 Global Configuration .....	362
6.4.2	IPv6 Interface Summary .....	364
6.4.3	IPv6 Interface Configuration .....	366
6.4.4	IPv6 Loopback Configuration .....	367
6.4.5	IPv6 Global Address Table .....	368
6.4.6	IPv6 Global Address Configuration .....	369
6.4.7	IPv6 Statistics .....	371
6.4.8	IPv6 Detailed Statistics .....	371
6.4.9	IPv6 Neighbor Table .....	375
6.5	Configuring IPv6 Routes .....	376
6.5.1	IPv6 Route Table .....	376
6.5.2	IPv6 Configured Routes .....	377
6.5.3	IPv6 ECMP Groups Summary .....	378
6.5.4	IPv6 Route Summary .....	379

6.6	Configuring DHCPv6 .....	381
6.6.1	DHCPv6 Global Configuration .....	381
6.6.2	DHCPv6 Pool Summary .....	381
6.6.3	DHCPv6 Pool Configuration .....	382
6.6.4	DHCPv6 Interface Summary .....	383
6.6.5	DHCPv6 Interface Configuration .....	384
6.6.6	DHCPv6 Binding Summary .....	385
6.6.7	DHCPv6 Statistics .....	386
6.6.8	DHCPv6 Server Conflicts Information .....	387
6.7	Configuring Policy Based Routing .....	388
7/	Managing Device Security .....	389
7.1	Captive Portal .....	389
7.1.1	Captive Portal Global Configuration .....	389
7.1.2	Captive Portal Instance Configuration .....	390
7.1.3	Captive Portal Page Design .....	391
7.1.4	Captive Portal Group Configuration .....	394
7.1.5	Captive Portal Local User Configuration .....	395
7.1.6	Captive Portal Interface Association .....	396
7.1.7	Captive Portal Traps .....	396
7.1.8	Captive Portal Global Status .....	397
7.1.9	Captive Portal Activation and Activity Status .....	398
7.1.10	Captive Portal Interface Activation Status .....	399
7.1.11	Captive Portal Interface Capability Status .....	399
7.1.12	Captive Portal Client Summary .....	400
7.1.13	Captive Portal Client Statistics .....	401
7.1.14	Captive Portal Client Interface Summary .....	402
7.2	Port Access Control .....	403
7.2.1	Global Port Access Control Configuration .....	403
7.2.2	Port Access Control Port Summary .....	404
7.2.3	Port Configuration .....	406
7.2.4	Port Details .....	408
7.2.5	Statistics .....	411
7.2.6	Client Summary .....	413
7.2.7	Privileges Summary .....	414
7.2.8	History Log Summary .....	415
7.3	RADIUS Settings .....	416
7.3.1	RADIUS Configuration .....	416
7.3.2	Server Statistics .....	420
7.3.3	Named Accounting Server Status .....	422
7.3.4	Accounting Statistics .....	422
7.3.5	Clear Statistics .....	423
7.3.6	Source Interface Configuration .....	424
7.4	TACACS+ Settings .....	425
7.4.1	TACACS+ Server Summary .....	425
7.4.2	TACACS+ Server Configuration .....	426
7.4.3	TACACS+ Source Interface Configuration .....	427
7.5	Authentication Manager .....	427
7.5.1	Authentication Manager Configuration .....	427
7.5.2	Authentication Manager Interface Configuration .....	429
7.5.3	Authentication Tiering .....	432
7.5.4	Authenticated Clients .....	434
7.5.5	Authenticated Statistics .....	435
7.5.6	Authenticated History .....	437
7.6	Media Access Control Security .....	438
7.6.1	MACsec Key Agreement Policy Summary .....	438

7.6.2	MKA Key Chain Summary .....	441
7.6.3	MKA Key Configuration .....	442
7.6.4	MKA Session Summary .....	444
7.6.5	MACsec Port Summary .....	447
7.6.6	MACsec Port Configuration .....	447
7.6.7	MKA Global Statistics .....	449
7.6.8	MKA Interface Statistics .....	451
<b>8/</b>	<b>Configuring Quality of Service .....</b>	<b>453</b>
8.1	Configuring Access Control Lists .....	453
8.1.1	IP Access Control Lists .....	453
8.1.2	IPv6 ACL Rules .....	465
8.2	Configuring Auto VoIP .....	467
8.2.1	Auto VoIP Global Configuration .....	467
8.2.2	OUI Table Summary .....	468
8.2.3	OUI Based Auto VoIP .....	469
8.2.4	Protocol Based Auto VoIP .....	470
8.3	Configuring Class of Service .....	471
8.3.1	IP DSCP Mapping Configuration .....	471
8.3.2	Interface Configuration .....	472
8.3.3	Interface Queue Configuration .....	473
8.3.4	CoS Interface Queue Drop Precedence Configuration .....	474
8.3.5	CoS Statistics .....	476
8.4	Configuring DiffServ .....	478
8.4.1	DiffServ Global Configuration and Status .....	478
8.4.2	DiffServ Class Summary .....	479
8.4.3	DiffServ Class Configuration .....	480
8.4.4	DiffServ Policy Summary .....	485
8.4.5	DiffServ Policy Configuration .....	485
8.4.6	DiffServ Service Summary .....	488
8.4.7	DiffServ Service Statistics .....	489
8.4.8	DiffServ Service Policy Statistics .....	490
<b>9/</b>	<b>Appendix: .....</b>	<b>Configuration Examples 492</b>
9.1	Configuring VLANs .....	492
9.1.1	Using the Web Interface to Configure VLANs .....	492
9.1.2	Using the CLI to Configure VLANs .....	494
9.1.3	Using the SNMP to Configure VLANs .....	494
9.2	Configuring Multiple Spanning Tree Protocol .....	496
9.2.1	Using the Web UI to Configure MSTP .....	496
9.2.2	Using the CLI to Configure MSTP .....	497
9.2.3	Using SNMP to Configure MSTP .....	498
9.3	Configuring VLAN Routing .....	499
9.3.1	Using the CLI to Configure VLAN Routing .....	499
9.3.2	Using SNMP to Configure VLAN Routing .....	500
9.4	Configuring Policy Based Routing .....	501
9.4.1	Configuring Policy Based Routing Using the CLI .....	502
9.5	Configuring OSPF .....	504
9.5.1	Using the CLI to Configure OSPF .....	504
9.5.2	Using the CLI to Configure OSPF Areas .....	506
9.5.3	Using the CLI to Configure OSPFv3 Enhancements .....	508
9.6	Configuring 802.1X Network Access Control .....	508
9.6.1	Using the CLI to configure 802.1X Port-Based Access Control .....	509
9.6.2	Using SNMP to configure 802.1X Port-Based Access Control .....	509
9.7	Configuring Authentication Tiering .....	510

9.7.1	Configuring Authentication Tiering Using the Web Interface .....	510
9.7.2	Configuring Authentication Tiering Using the CLI .....	512
9.8	Configuring Differentiated Services for VoIP .....	512
9.8.1	Using the CLI to Configure DiffServ VoIP Support .....	513
9.8.2	Using SNMP to Configure DiffServ VoIP Support .....	514
9.9	Configuring PIM .....	515
9.9.1	Using the CLI to Configure PIM-SMv4 .....	515
9.9.2	Using SNMP to Configure PIM-SMv4 .....	516
9.9.3	Configuring IP Multicast Routing with PIM Sparse Mode .....	516
9.10	Configuring FIP Snooping .....	519
9.10.1	Using the CLI to Configure FIP snooping .....	519
9.11	Configuring OpenFlow .....	521
9.11.1	Enabling and Disabling OpenFlow .....	521
9.11.2	Interacting with the OpenFlow Manager .....	522
9.11.3	Deploying OpenFlow .....	522
9.11.4	OpenFlow Scenarios .....	522
9.11.5	OpenFlow Interaction with Other Functions .....	522
9.11.6	OpenFlow Variants .....	523
9.11.7	Configuring OpenFlow .....	523
9.12	IGMP and MLD Snooping Switches .....	526
9.12.1	Snooping Functionality on a FASTPATH Switch .....	526
9.12.2	Snooping Switch Restrictions .....	528
9.12.3	Configuring IGMP and MLD Snooping .....	528
9.13	Multicast Snooping Example (with IP Multicast Routing) .....	530
9.13.1	Snooping Within a Subnet .....	532
9.13.2	Snooping on a Multicast Router .....	532
9.14	Configuring Port Mirroring .....	534
9.15	Configuring RSPAN .....	534
9.15.1	Configuring RSPAN Using the Web Interface .....	535
9.15.2	Configuring RSPAN Using the CLI .....	540
9.16	Configuring VPC .....	541
9.16.1	Configuring VPC Using the Web Interface .....	542
9.16.2	Configuring VPC Using the CLI .....	545
9.17	Virtual Routing and Forwarding Lite Operation and Configuration .....	547
9.17.1	Overview .....	547
9.17.2	VRF Functionality .....	547
9.17.3	VRF and FASTPATH Feature Support .....	549
9.17.4	VRF Lite Deployment Scenarios .....	550
9.17.5	VRF Configuration Example .....	552
9.18	Bidirectional Forwarding Detection .....	553
9.18.1	Overview .....	553
9.18.2	Configuring BFD .....	554

# 1/ Introduction

## 1.1 About This Document

This guide describes how to configure the FASTPATH® software features by using the Web-based GUI. The FASTPATH architecture accommodates a variety of software modules so that a platform running FASTPATH software can function as a Layer 2 switch in a basic network or a Layer 3 router in a large, complex network.

The information in this guide is intended for any of the following individuals:

- System administrators who are responsible for configuring and operating a network using FASTPATH software
- Software engineers who are integrating FASTPATH software into a router or switch product
- Level 1, Level 2, or both Support providers

To obtain the greatest benefit from this guide, you should have an understanding of the base software and should have read the specification for your networking device platform. You should also have basic knowledge of Ethernet and networking concepts.

### 1.1.1 Document Conventions

The following conventions may be used in this document.

Convention	Description
<b>Bold</b>	User input and actions: for example, type exit, click OK, press Alt+C
<a href="#">Blue Text</a>	Hyperlinked text.
Monospace	Code: #include <iostream> Command line commands and parameters: show network
<i>Monospace italic</i>	Placeholders for required elements: username name
[ ]	Indicates optional command-line parameters: show port [all]
{ }	Indicates a choice in command-line parameters: network protocol {dhcp bootp none}

## 1.2 About FASTPATH Software Modules

The FASTPATH software suite includes the following modules:

- Switching (Layer 2)
- Routing (Layer 3)
- IPv4-IPv6 routing
- Multicast
- Quality of Service
- Management (CLI, Web UI, SNMP, Open Ethernet Networking (OpEN) API, RESTful API, RESTCONF, NETCONF)
- Metro
- Stacking
- Data Center
- BGP
- IPV6 Management
- QoS
- Secure Management
- Service Provider

Not all modules are available for all platforms or software releases. The FASTPATH modules can be applied in various combinations to develop advanced Layer 2/3/4+ products. The user-configurable features available on your switch depend on the installed modules.

## 1.3 Configuring Data Center and BGP Features

For most features, you can use either the web-based user interface or the command-line interface (CLI) to view and configure settings. However, FASTPATH includes several data center features that can be configured only by using the CLI. The following data center features cannot be configured by using the web-based user interface:

- Data Center Bridging Exchange protocol (DCBX)
- FIP Snooping Bridge
- Congestion Notification (IEEE 802.1Qau)
- Enhanced Transmission Selection (ETS), defined in IEEE 802.1Qaz
- OpenFlow

Additionally, the Border Gateway Control (BGP) features can be configured only by using the CLI. There are no web pages within FASTPATH for configuring BGP.

For information about the CLI commands you use to configure Data Center features and BGP, refer to the FASTPATH CLI Command Reference Guide.

## 2/ Getting Started

This section describes how to start the switch and access the user interface.

### 2.1 Connecting the Switch to the Network

To enable remote management of the switch through telnet, a Web browser, or SNMP, you must connect the switch to the network. The switch has no IP address by default, and DHCP is disabled, so you must provide network information by connecting to the switch command-line interface (CLI) by using a local serial connection.

To access the switch over a network you must first configure it with network information (an IP address, subnet mask, and default gateway). You can configure the IP information using any of the following:

- BOOTP
- DHCP
- Terminal interface via the EIA-232 port

After you configure network information, such as the IP address and subnet mask, and the switch is physically and logically connected to the network, you can manage and monitor the switch remotely through SSH, telnet, a Web browser, or an SNMP-based network management system. You can also continue to manage the switch through the terminal interface via the EIA-232 port.

---

#### **NOTICE**

Some switches provide a Service port, an Ethernet port usually located on the back on the switch, as a dedicated management port. On switches without a Service port, you use one of the network ports.

---

After you perform the physical hardware installation, you need to make a serial connection to the switch so that you can do one of the following:

- Manually configure network information for the management interface, or
- Enable the management interface as a DHCP or BootP client on your network (if not already enabled) and then view the network information after it is assigned by the DHCP server.

To connect to the switch and configure or view network information, use the following steps:

1. **Using a straight-through modem cable, connect a VT100/ANSI terminal or a workstation to the console (serial) port.**  
If you attached a PC, Apple, or UNIX workstation, start a terminal-emulation program, such as HyperTerminal or Tera Term.
2. **Configure the terminal-emulation program to use the following settings:**
  - Baud rate: 115200 bps
  - Data bits: 8
  - Parity: none

- Stop bit: 1
  - Flow control: none
3. Power on the switch.
  4. Press the return key, and the `User:` prompt appears.  
 Enter `admin` as the user name. There is no default password. Press ENTER at the password prompt if you did not change the default password.  
 After a successful login, the screen shows the system prompt, for example `(switch)>`.
  5. At the `(switch)>` prompt, enter `enable` to enter the Privileged EXEC command mode. There is no default password to enter Privileged EXEC mode. Press ENTER at the password prompt if you did not change the default password.  
 The command prompt changes to `(switch)#`.
  6. Configure network information.  
 If the unit has a service port:
    - To have the address assigned through DHCP:  
 By default, the port is configured as a DHCP client. If your network has a DHCP server, then you need only to connect the switch to your network.
    - To use BootP, change the protocol by entering:  

```
serviceport protocol bootp
```
    - To disable DHCP/BootP and manually assign an IPv4 address, enter:  

```
serviceport protocol none
serviceport ip <ipaddress> <netmask> [<gateway>]
```
 For example:  

```
serviceport ip 192.168.2.23 255.255.255.0 192.168.2.1
```

    - To disable DHCP/BootP and manually assign an IPv6 address and (optionally) default gateway, enter:  

```
serviceport protocol none
serviceport ipv6 address <address>/<prefix-length> [eui64]
serviceport ipv6 gateway <gateway>
```
    - To view the assigned or configured network address, enter:  

```
show serviceport
```
 If the unit does not have a service port:
    - To use a DHCP server to obtain the IP address, subnet mask, and default gateway information, enter:  

```
network protocol dhcp.
```
    - To use a BootP server to obtain the IP address, subnet mask, and default gateway information, enter:  

```
network protocol bootp.
```
    - To manually configure the IPv4 address, subnet mask, and (optionally) default gateway, enter:  

```
network parms <ipaddress> <netmask> [<gateway>],
```
 For example:  

```
network parms 192.168.2.23 255.255.255.0 192.168.2.1
```

    - To manually configure the IPv6 address, subnet mask, and (optionally) default gateway, enter:  

```
network ipv6 address <address>/<prefix-length> [eui64]
network ipv6 gateway <gateway>
```
    - To view the network information, enter `show network`.
    - To save these changes so they are retained during a switch reset, enter the following command:  

```
copy system:running-config nvram:startup-config
```

After the switch is connected to the network, you can use the IP address for remote access to the switch by using a Web browser or through telnet or SSH.

## 2.2 Understanding the User Interfaces

FASTPATH software includes a set of comprehensive management functions for configuring and monitoring the system by using one of the following methods:

- Web User Interface
- Command-Line Interface (CLI)
- Simple Network Management Protocol (SNMP)
- OpEN API
- RESTful API
- RESTCONF
- NETCONF

Each of the standards-based management methods allows you to configure and monitor the components of the FASTPATH software. The method you use to manage the system depends on your network size and requirements, and on your preference.

---

Not all components can be managed by each interface.

**NOTICE**

---

This guide describes how to use the Web-based interface to manage and monitor the system. For information about how to manage and monitor the system by using the CLI, refer to the FASTPATH CLI Command Reference and the FASTPATH Configuration Guide.

---

The Web configuration and monitoring pages and the CLI commands available for each platform depend on the FASTPATH software version and modules installed. For more information about the modules, see [Chapter 2, Getting Started](#).

**NOTICE**

## 2.2.1 Using the Web Interface

To access the switch by using a Web browser, the browser must meet the following software requirements:

- HTML version 4.0, or later
- HTTP version 1.1, or later
- JavaScript version 1.5, or later

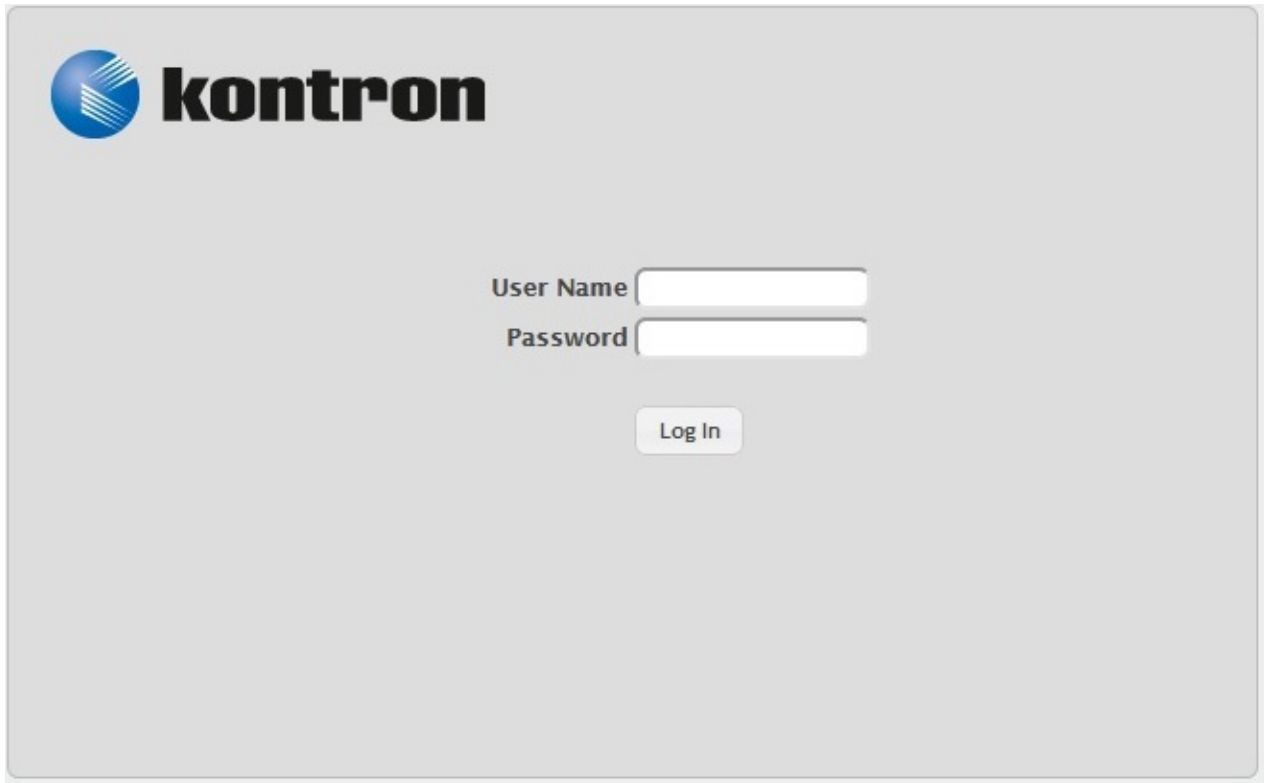
Use the following procedures to log in to the Web Interface shown in [Figure 1](#):

1. **Open a Web browser and enter the IP address of the switch in the Web browser address field.**
2. **Type the user name and password into the fields on the login screen, and then click Login.**

The user name and password are the same as those you use to log on to the command-line interface. By default, the user name is admin, and there is no password. Passwords are case sensitive.



Figure 1: Login Screen



**kontron**

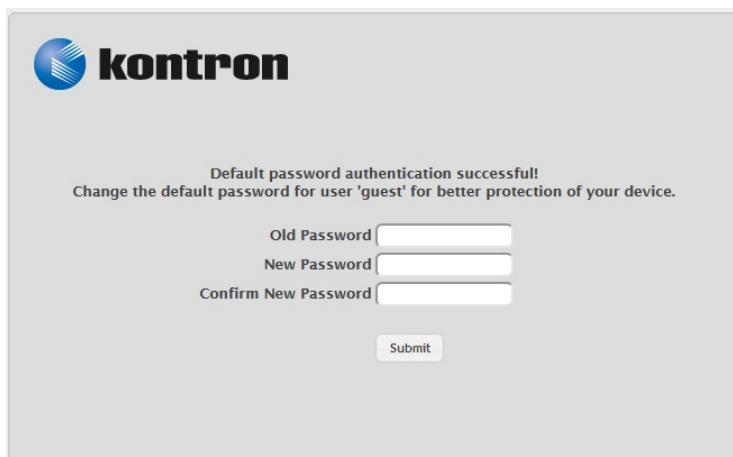
User Name

Password

Log In

3. After the initial default password authentication is successful, the default or pre-configured users (admin and guest) are redirected to a new Web page to change the default password, as shown in the following figure.

Figure 2: Default Password Change Page



**kontron**

Default password authentication successful!  
Change the default password for user 'guest' for better protection of your device.

Old Password

New Password

Confirm New Password

Submit

4. After the default password is successfully changed, the default or pre-configured users (admin and guest) are required to log in again (see [Figure 1](#)) with the new credentials to get access to the device.
5. When the default password is changed, it can be reset to the factory default setting only if the user resets the system configuration on the Reset Configuration page. In which case, the default or pre-configured users will be required to reinitiate the default password change procedure to get access to the device.
6. After the system authenticates you, the System Description page displays.

[Figure 3: "Web Interface Layout," on page 20](#) shows the layout of the FASTPATH software Web interface. Each Web page contains two main areas: the navigation menu, and the configuration status and options.

Figure 3: Web Interface Layout

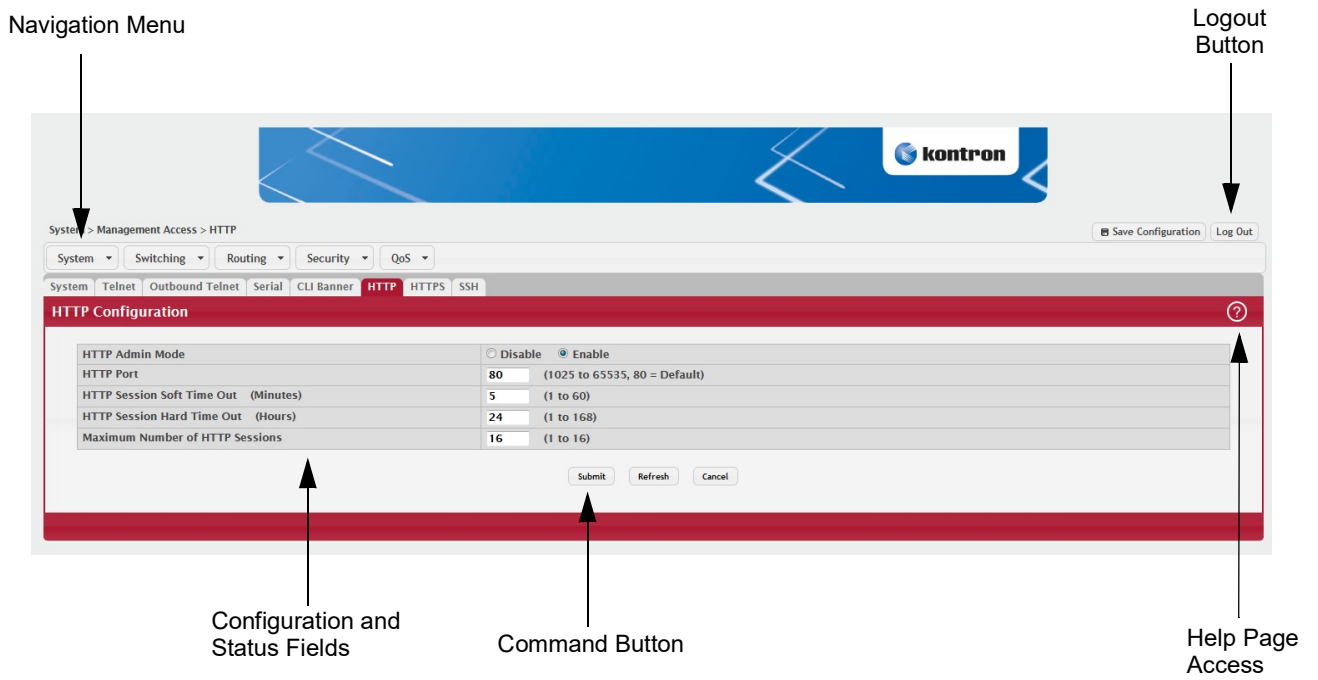
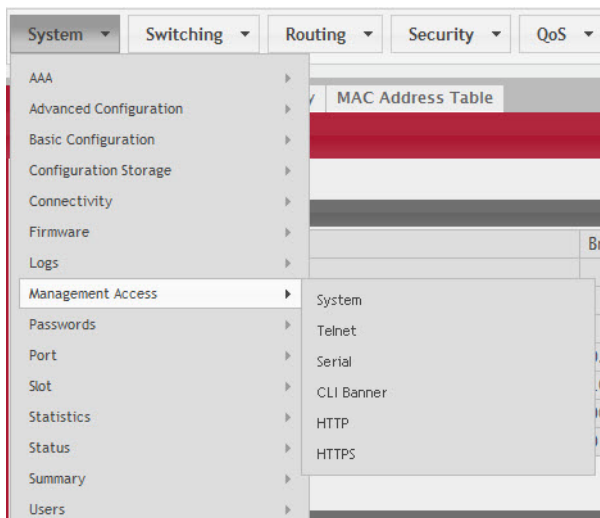


Figure 4: Management Access Menu

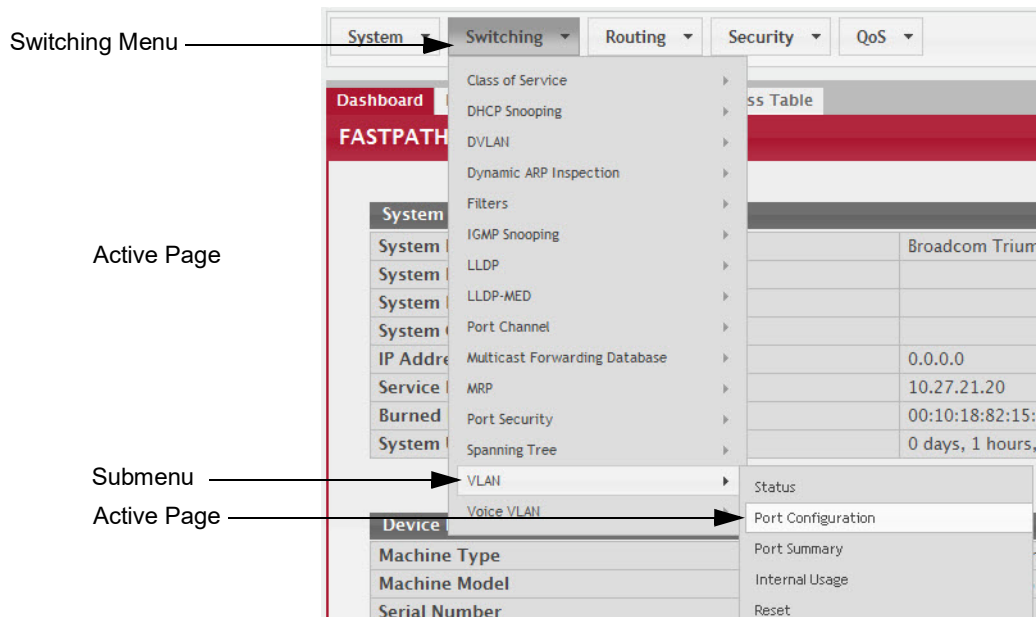


### 2.2.1.1 Navigation Menu

The navigation menu is on the top of the Web interface. The navigation menu contains a list of various device features. The main items in the navigation menu can be expanded to view all the components under a specific feature, or retracted to hide the feature's components.

The navigation menu consists of a combination of main feature menus, submenus, and configuration and status pages. Click the feature menu, such as System or Switching, to view the options in that menu. Each menu contains submenus, HTML pages, or a combination of both. [Figure 5: "Navigation Menu View," on page 21](#) shows an example of a feature menu (Switching), submenu (VLAN), and the active page in the navigation menu (Port Configuration).

**Figure 5: Navigation Menu View**



When you click a menu or submenu, the color turns from gray to red, the menu expands to show its contents, and the arrow on the right side of the menu rotates. If you click a page under a menu or submenu, a new page displays in the main frame.

### 2.2.1.2 Configuration and Status Fields

The main area of the screen displays the fields you use to configure or monitor the switch. On pages that contain configuration options, you can input information into fields or select options from drop-down menus.

Each page contains access to the HTML-based help that explains the fields to configure or view on the page. Many pages also contain command buttons.

Table 1 shows the command buttons that are used throughout the pages in the Web interface.

Table 1: Common Command Buttons

Button	Function
<b>Submit</b>	Sends the updated configuration to the switch. Configuration changes take effect immediately, but changes are not retained across a power cycle unless you save them to the system configuration file. To save the configuration to non-volatile memory, navigate to the System > System Utilities > Save All Applied Changes page and click Save.
<b>Refresh</b>	Refreshes the page with the most current information.
<b>Delete</b>	Removes the selected entry from the running configuration.
<b>Clear</b>	Removes all entries from a table or resets statistical counters to the default value.
<b>Edit</b>	Changes an existing entry.
<b>Remove</b>	Deletes the selected entries.
<b>Clear Counter</b>	Clear all the statistics counters, resetting all switch summary and detailed statistics to default values.
<b>Logout</b>	Ends the session.

**ATTENTION:** Submitting changes makes them effective during the current boot session only. You must save any changes if you want them to be retained across a power cycle (reboot).

### 2.2.1.3 Table Sorting

Tables shown in the web pages now have the ability to be sorted in each column. To sort a column, click at the top of the column to sort by that field. For example, in the Event Log page, clicking on the Event Time will sort the entries by that field.

### 2.2.1.4 Help Page Access

The Help button is always available in the upper right corner of the active page. Click Help to open a new page that contains information about the configuration fields, status fields, and command buttons available on the active page. The online help pages are context sensitive. For example, if the IP Addressing page is open, the help topic for that page displays if you click Help. Figure 6 shows the Help icon.

Figure 6: Help Icon



Figure 3: "Web Interface Layout," on page 20 shows the location of the Help link on the Web interface.

### 2.2.1.5 User-Defined Fields

User-defined fields can contain 1-159 characters, unless otherwise noted on the configuration Web page.

All characters may be used except for the following (unless specifically noted in for that feature):

\                    <  
/                    >|  
\*                    |  
?

## 2.2.2 Using the Command-Line Interface

The command-line interface (CLI) is a text-based way to manage and monitor the system. You can access the CLI by using a direct serial connection or by using a remote logical connection with telnet or SSH.

The CLI groups commands into modes according to the command function. Each of the command modes supports specific software commands. The commands in one mode are not available until you switch to that particular mode, with the exception of the User EXEC mode commands. You can execute the User EXEC mode commands in the Privileged EXEC mode.

To display the commands available in the current mode, enter a question mark (?) at the command prompt. To display the available command keywords or parameters, enter a question mark (?) after each word you type at the command prompt. If there are no additional command keywords or parameters, or if additional parameters are optional, the following message appears in the output:

```
<cr>                               Press Enter to execute the command
```

For more information about the CLI, refer to the FASTPATH CLI Command Reference Guide.

The FASTPATH CLI Command Reference lists each command available from the CLI by the command name and provides a brief description of the command. Each command reference also contains the following information:

- The command keywords and the required and optional parameters.
- The command mode you must be in to access the command.
- The default value, if any, of a configurable setting on the device.

The `show` commands in the document also include a description of the information that the command shows.

### 2.2.2.1 Using SNMP

For FASTPATH software that includes the SNMP module, you can configure SNMP groups and users that can manage traps that the SNMP agent generates.

FASTPATH uses both standard, public MIBs for standard functionality and private MIBs that support additional switch functionality. All private MIBs begin with a "-" prefix. The main object for interface configuration is in -SWITCHING-MIB, which is a private MIB. Some interface configurations also involve objects in the public MIB, IF-MIB.

SNMP is enabled by default. The System Description Web page, which is the page the displays after a successful login and the `show sysinfo` command display the information you need to configure an SNMP manager to access the switch.

Any user can connect to the switch using the SNMPv3 protocol, but for authentication and encryption, you need to configure a new user profile. To configure a profile by using the CLI, refer to the SNMP section in the FASTPATH CLI Command Reference Guide.

To configure an SNMPv3 profile by using the Web interface, use the following steps:

1. **Select System > Configuration > User Accounts from the navigation menu on the left side of the Web interface.**
2. **From the User menu, select Create to create a new user.**
3. **Enter a new user name in the User Name field.**
4. **Enter a new user password in the Password field and then retype it in the Confirm Password field.**  
To use SNMPv3 Authentication for this user, set a password of eight or more alphanumeric characters.
5. **To enable authentication, use the Authentication Protocol menu to select either MD5 or SHA for the authentication protocol.**
6. **To enable encryption, use the Encryption Protocol menu to select DES for the encryption scheme. Then, enter an encryption code of eight or more alphanumeric characters in the Encryption Key field.**
7. **Click Submit.**

To access configuration information for SNMPv1 or SNMPv2, click and click the page that contains the information to configure.

## 3/ Getting Started with Stacking

This section describes the concepts and operating procedures to manage stacked Ethernet switches running FASTPATH.

---

### NOTICE

For complete syntax and usage information for the commands used in this chapter, refer to the FASTPATH CLI Command Reference Guide for this release.

---

### 3.1 Understanding Switch Stacks

A switch stack is a set of up to *n* Ethernet switches connected through their stacking ports. One of the switches controls the operation of the stack and is called the stack manager. All other switches in the stack are stack members. The stack members use stacking technology to behave and work together as a unified system. Layer 2 and Layer 3 protocols present the entire switch stack as a single entity to the network.

---

### NOTICE

The value of *n* when referring to the number of connected Ethernet switches is dependent on the platform being used. Refer to the FASTPATH Scaling Parameters and Values for the correct value for a given platform.

---

The stack manager is the single point of stack-wide management. From the stack manager, you configure:

- System-level (global) features that apply to all stack members
- Interface-level features for all interfaces on any stack member

A switch stack is identified in the network by its network IP address. The network IP address is assigned according to the MAC address of the stack manager. The MAC address used by the switch is the MAC address of the manager. You can see this address by issuing the `show network` command. Every stack member is uniquely identified by its own stack member number.

All stack members are eligible stack managers. Exception: Setting a stack member's priority to 0 (zero) makes it ineligible for manager selection. When the stack is formed, one of the units is automatically selected as the Standby for the stack. The standby of the stack takes over as Manager if the current Manager fails. The standby of the stack can also be configured using the `standby <unit-number>` command.

The stack manager contains the saved and running configuration files for the switch stack. The configuration files include the system-level settings for the switch stack and the interface-level settings for all stack members. Each stack member retains a copy of the saved file for backup purposes.

If the manager is removed from the stack, the standby of the stack will take over and will then run from that saved configuration.

You can use these methods to manage switch stacks:

- Web interface
- Command line interface (CLI) over a serial connection to the console port of the manager
- A network management application through the Simple Network Management Protocol (SNMP)

#### 3.1.1 Switch Stack Membership

A switch stack has up to *n* stack members, including the manager, connected through their stacking ports. A switch stack always has one stack manager.

A standalone switch is a switch stack with one stack member that also operates as the stack manager. You can connect one standalone switch to another to create a switch stack containing two stack members, with one of them being the stack manager. You can connect standalone switches to an existing switch stack to increase the stack membership.

If you replace a stack member with an identical model, the new switch functions with exactly the same configuration as the replaced switch, assuming that the new switch is using the same member number as the replaced switch. By default, FASTPATH configures the new member.

The operation of the switch stack continues uninterrupted during membership changes unless you remove the stack manager.

### 3.1.2 Stack Manager Election and Re-Election

The stack manager is elected or re-elected based on one of these factors and in the order listed:

- The switch that is currently the stack manager
- The switch with the highest stack member priority value

---

#### NOTICE

Assign the highest priority value to the switch that you prefer to be the stack manager. This ensures that the switch is re-elected as stack manager if a re-election occurs.

---

- The switch with the higher MAC address

A stack manager retains its role unless one of these events occurs:

- The stack manager is removed from the switch stack
- The stack manager is reset or powered off
- The stack manager has failed
- The switch stack membership is increased by adding powered-on standalone switches or switch stacks

In the case of a manager re-election, the new stack manager becomes available after a few seconds.

If a new stack manager is elected and the previous stack manager becomes available, the previous stack manager does not resume its role as stack manager.

### 3.1.3 Stack Member Numbers

A stack member number (1 to n) identifies each member in the switch stack. The member number also determines the interface-level configuration that a stack member uses. You can display the stack member number by using the `show switch` Privileged EXEC command.

A new, out-of-the-box switch (one that has not joined a switch stack or has not been manually assigned a stack member number) ships with a default stack member number of 1. When it joins a switch stack, its default stack member number changes to the lowest available member number in the stack.

Stack members in the same switch stack cannot have the same stack member number. Every stack member, including a standalone switch, retains its member number until you manually change the number or unless the number is already being used by another member in the stack. See [Section 3.11: "Renumbering Stack Members"](#) and [Section 3.15: "Merging Two Operational Stacks"](#).

### 3.1.4 Stack Member Priority Values

Starting with FASTPATH 7.2, you can set the stack member's priority in the range 0 to 15.

---

#### NOTICE

Setting the switch priority to 0 (zero) makes it ineligible for manager selection.

---

## 3.2 Switch Stack Software Compatibility Recommendations

All stack members must run the same FASTPATH software version to ensure compatibility between stack members. The software versions on all stack members, including the stack manager, must be the same. This helps ensure full compatibility in the stack protocol version among the stack members.

If a stack member is running a software version that is not the same as the stack manager, then the stack member joins the stack but stays in code incompatible status (the stack unit is not allowed to join the stack as a fully functional member). Use the `show switch` command to list the stack members and the software versions. The new unit will be visible. The administrator can load the code to that new unit and reset the unit. The ports on the unit in software mismatch state do not come up.

### 3.3 Incompatible Software and Stack Member Image Upgrades

You can upgrade a switch that has an incompatible software image by using the `copy {active | backup} unit://<unit-number>/{active | backup}` command from config stack mode. It copies the software image from an existing stack member to the one with incompatible software. Because that switch does not automatically reload, issue a `reload` command to that switch and it joins the stack as a fully functioning member.

### 3.4 Switch Stack Configuration Files

The configuration files record settings for all global and interface specific settings that define the operation of the stack and individual members. Once a save to the configuration is issued, all stack members store a copy of the configuration settings. If a stack manager becomes unavailable, any stack member assuming the role of stack manager will operate from the saved configuration files.

When a new, out-of-box switch joins a switch stack, it uses the system-level settings of that switch stack. If the switch to store this system-level configuration, you must issue the following command:

```
copy system:running-config nvram:startup-config
(in Privileged EXEC)
```

This will save passwords and all other changes to the device.

If you do not save the configuration by doing this command, all configurations will be lost when a power cycle is performed on the networking device or when the networking device is reset.

#### NOTICE

After downloading a configuration file to a stack, you must perform a configuration save operation from the FASTPATH user interface (that is, the `copy` command shown above) to distribute this configuration to non-management units in the stack. This is also true of SSH key files and SSL certificate files. From the command line interface, the following command can be used: `copy system:running-config nvram:startup-config (in Privileged EXEC)`

You back up and restore the stack configuration in the same way as you would for standalone switch configuration.

### 3.5 Switch Stack Management Connectivity

You manage the switch stack and the stack member interfaces through the stack manager. You can use the web interface, CLI, and SNMP. You cannot manage stack members on an individual switch basis.

#### 3.5.1 Connectivity to the Switch Stack Through Console Ports

You can connect to the stack manager through the console port of the stack manager only.

#### 3.5.2 Connectivity to the Switch Stack Through Telnet

You can also telnet to the stack manager using the command `telnet <ipaddress>` then `login`.

### 3.6 General Practices

The following practices are recommended:

- When issuing a command (such as `move management`, or `renumber`), allow the command to fully complete before issuing the next command. For example, if you issue a `reset` to a stack member, use the `show port` command to verify that the unit has re-merged with the stack, and all ports are joined before issuing the next command.
- When physically removing or relocating a unit, always power down the unit before disconnecting stack cables.
- When reconnecting stack cables, connect them before powering up the unit, if possible. Tighten all connector screws, where applicable, to ensure a good connection.

The following sections provide switch stack configuration scenarios. Most of the scenarios assume at least two switches are connected through their stacking ports.



## 3.7 Initial Installation and Power-up of a Stack

Use the following steps to install and power-up a stack of switches:

1. Install units in rack whenever possible to prevent the units and cables from being disturbed
2. Install all stacking cables. Fully connect all cables, including the redundant stack link. Install a redundant link because this provides stack resiliency.
3. Identify the unit to be the manager. Power this unit up first.
4. To set up a stack, complete the following steps:
  - a. Make sure there is a FASTPATH image on each box.
  - b. If the image does not exist or needs to be updated, use TFTP or xmodem to perform the update operation.
5. Monitor the console port. Allow this unit to come up to the login prompt. If the unit has the default configuration, it should come up as unit #1, and will automatically become a manager unit. If not, renumber the unit as desired.
6. If desired, preconfigure other units to be added to the stack. See [Section 3.16: "Preconfiguration"](#).
7. Power on a second unit, making sure it is adjacent (next physical unit in the stack) to the unit already powered up. This will ensure the second unit comes up as a member of the stack, and not a Manager of a separate stack.
8. Monitor the manager unit to see that the second unit joins the stack. Use the `show switch` command to determine when the unit joins the stack. It will be assigned a unit number (unit #2, if it has the default configuration.)
9. If desired, renumber this stack unit. See [Section 3.11: "Renumbering Stack Members"](#) for recommendations for renumbering stack members.
10. Repeat steps 6 through 8 to add additional members to the stack. Always power on a unit adjacent to the units already in the stack.

## 3.8 Removing a Unit from the Stack

Use the following steps to remove a switch from the stack:

1. Make sure the redundant stack connection is in place and functional. All stack members should be connected in a logical ring.
2. Power down the unit to be removed.
3. Disconnect the stacking cables
4. If the unit is not to be replaced, reconnect the stack cable from the stack member above to the stack member below the unit being removed.
5. Remove the unit from the rack.
6. If desired, remove the unit from the configuration by issuing the command: `no member <unit-id>` in Stack mode.

Using the Web Interface, you delete a member of the stack through the Stacking - Unit Configuration page. To delete a member, select the unit number on the Switch ID menu and click the Delete button.

## 3.9 Adding a Unit to an Operating Stack

Use the following steps to add a switch to a stack of switches while the stack is running:

1. Make sure that the redundant stack link is in place and functional. All stack members should be connected in a logical ring.
2. Preconfigure the new unit, if desired.
3. Install the new unit in the rack. (Assumes installation below the bottom-most unit, or above the top-most unit).
4. Disconnect the redundant stack cable that connects the last unit in the stack back up to the first unit in the stack at the new position in the ring where the new unit is to be inserted.
5. Connect this cable to the new unit, following the established order of connections. In other words, use the redundant stack cable to connect from the first box in the stack to the last.

6. Power up the new unit. Verify, by monitoring the manager unit console port, that the new unit successfully joins the stack by using the `show switch` command in EXEC mode. The new unit should always join as a member (never as manager; the existing manager of the stack should not change).
7. If the FASTPATH software version of the newly added member is not the same as the existing stack, update the software image.

Adding a powered-up standalone unit to an operational stack is similar to merging two operational stacks where the standalone unit is a stack of one unit. See [Section 3.15: "Merging Two Operational Stacks"](#) for more details.

Using the Web Interface, you create a new member for the stack through the Stacking - Unit Configuration page. To create a new member, select the create option from the Switch ID pull-down menu.

## 3.10 Replacing the Stack Member with a New Unit

There are two options here. If a stack member of a certain model number is replaced with another unit of the same model, follow these steps:

1. Follow the process in [Section 3.8: "Removing a Unit from the Stack"](#) to remove the desired stack member.
2. Follow the process in [Section 3.9: "Adding a Unit to an Operating Stack"](#) to add a new member to the stack with the following exceptions:
  - Insert the new member in the same position in the stack as the one removed.
  - The preconfiguration described in step 2 of [Section 3.9: "Adding a Unit to an Operating Stack"](#) is not required.

If a stack member is replaced with a unit of a different model number, follow these steps:

1. Follow the process in [Section 3.8: "Removing a Unit from the Stack"](#) to remove the desired stack member.
2. Remove the now-absent stack member from the configuration by issuing the `no member` command in Config Stack mode.
3. Add the new stack unit to the stack using the process described in [Section 3.9: "Adding a Unit to an Operating Stack"](#). The unit can be inserted into the same position as the unit just removed, or the unit can be inserted at the bottom of the stack. In either case, make sure all stack cables are connected with the exception of the cable at the position where the new unit is to be inserted to ensure that the stack does not get divided into two separate stacks, causing the election of a new manager.

## 3.11 Renumbering Stack Members

1. If particular numbering is required, assign specific numbers to stack members when they are first installed and configured in the stack, if possible.
2. If the desired stack unit number for a particular unit is unused, a unit can be renumbered simply by using the `switch <oldunit-id> renumber <newunit-id>` CLI command in Global Config mode.
3. Renumbering a non-manager unit requires a unit reset for the renumbering to take effect. Renumbering a manager unit requires a reset of all the switches in the stack for the renumbering to take effect.
4. If the newunit-id has been preconfigured, you may need to remove the newunit-id from the configuration before renumbering the unit.
5. If reassignment of multiple existing stack unit numbers is necessary, there are a number of implications in terms of mismatching of configuration. In this case, power down all units except the manager and add back one at a time using the procedure in [Section 3.9: "Adding a Unit to an Operating Stack"](#).

Using the Web Interface, you renumber a switch through the Stacking - Unit Configuration page. To renumber a switch:

1. Select the switch you want to renumber from the Switch ID menu
2. Type the new number into the Switch ID input box and click a button to submit

## 3.12 Moving a Manager to a Different Unit in the Stack

Use the following steps to change the stack manager from the current switch to a new switch in the stack:

1. Using the `movemanagement` command, move the manager to the desired unit number. The operation may take three minutes or longer depending on the stack size and configuration. The command is `movemanagement <fromunit-id><tounit-id>` in Config Stack mode.
2. Make sure that you can log in on the console attached to the new manager. Use the `show switch` command to verify that all units rejoined the stack.
3. Reset the stack with the `reload` command in Privileged EXEC mode after moving the manager.

## 3.13 Removing a Manager Unit from an Operating Stack

Use the following steps to remove the manager unit from the stack during operation:

1. Move the designated manager to a different unit in the stack using the [Section 3.12: "Moving a Manager to a Different Unit in the Stack"](#) procedure on this page.
2. Using the procedure [Section 3.8: "Removing a Unit from the Stack"](#), remove the unit from the stack.

## 3.14 Initiating a Warm Failover of the Manager Unit

You can use the `initiate failover` command to initiate a warm restart. This command reloads the management unit, triggering the standby unit to take over. As the standby management unit takes over, the system continues to forward end-user traffic. The end-user data streams may lose a few packets during the failure, but they do not lose their IP sessions, such as VoIP calls.

If there no standby unit is available when the `initiate failover` command is issued, the command fails with an error message stating that no standby unit exists. If the standby unit is not ready for a warm restart, the command fails with a similar error message. The `move management` command triggers a cold restart, even if the target unit is the backup unit.

## 3.15 Merging Two Operational Stacks

The recommended procedure for merging two operational stacks is as follows:

1. Always power off all units in one stack before connecting to another stack.
2. Add the units as a group by unplugging one stacking cable in the operational stack and physically connecting all unpowered units.
3. Completely cable the stacking connections, making sure the redundant link is also in place.

Two operational stacks can also be merged by reconnecting stack cables without powering down all units in one stack. Connecting a powered-up standalone unit to an existing stack leads to same behavior as when merging two operational stacks. In such cases, the Manager re-election is done based on the rules listed in [Section 3.1.2: "Stack Manager Election and Re-Election"](#). One of the two managers wins the election and the losing stack manager resets itself and all its member units. After the reset, all the losing stack members join the winning stack to form a single stack. The winning stack remains functional through the merge process. If the stack merge is performed in this way, then it is strongly recommended that the user set the priority of the desired winner stack manager to a higher value than the stack manager that should lose the election.

## 3.16 Preconfiguration

This section is intended to explain how to configure units. Units do not necessarily have to be preconfigured to be added to the stack.

1. **General information:** All configuration on the stack, except unit numbers, is stored on the management unit. This means that a stack unit may be replaced with another device of the same type without having to reconfigure the switch. Unit numbers are stored independently on each switch, so that after power cycling the stack, the units always come back with the same unit numbers. The unit type associated with each unit number may be learned by the management unit automatically as the units are connected or preconfigured by the administrator.

2. Issue the `member <unit-id> <switchindex>` command to preconfigure a unit from the config stack mode. Supported unit types are shown by the `show supported switchtype` command.
  - To display supported switches:  
Use Privileged EXEC mode  
Enter the command `show supported switchtype <x>` where x is the SID.
  - To add a new member (see [Section 3.9: "Adding a Unit to an Operating Stack"](#)):  
Use Config stack mode  
Enter the `member <unit-id>` command
3. Next, configure the unit you just defined with configuration commands, just as if the unit were physically present.
4. Ports for the preconfigured unit come up in detached state and can be seen with the `show port all` command in Privileged EXEC mode. The detached ports may now be configured for VLAN membership and any other port-specific configuration.
5. After a unit type is preconfigured for a specific unit number, attaching a unit with a different unit type for this unit number causes the switch to report an error. The Privileged Exec mode `show switch` command indicates config mismatch for the new unit and the ports on that unit do not come up. To resolve this situation, you may change the unit number of the mismatched unit, using the procedure in [Section 3.11: "Renumbering Stack Members"](#), or delete the preconfigured unit type using the `no member <unit-id>` command from the config stack mode.

## 3.17 Stack Links

Use the Stack Summary page to configure the Stack Trunk Hash mode on all HiGig™ trunks across the units in the stack. To navigate to the Stack Summary page, click Stacking > Base > Summary in the navigation menu.

Figure 7: Stack Summary

The screenshot shows the 'Stacking > Base > Summary' page. At the top, there are navigation tabs for System, Switching, Routing, Security, QoS, and Stacking. Below these are sub-tabs for Summary, Unit Configuration, Supported Switches, and Firmware Update. The 'Stack Summary' section features a dropdown for 'Stack Trunk Hash Mode' set to 'Source/Destination MAC'. Below this is a table with columns: Switch ID, Status, Management Status, Standby Switch, Preconfigured Model Identifier, and Plugged-in Model Identifier. The table contains one row with the following data: Switch ID 1, Status OK, Management Status Management Switch, Standby Switch, Preconfigured Model Identifier BCM-56970, and Plugged-in Model Identifier BCM-56970. At the bottom of the table, there are navigation buttons: First, Previous, 1, Next, Last. Below the table are buttons for Submit, Refresh, Add, Edit, and Remove.

Switch ID	Status	Management Status	Standby Switch	Preconfigured Model Identifier	Plugged-in Model Identifier
1	OK	Management Switch		BCM-56970	BCM-56970

Select one of the available modes from the list and click the Submit button. When one of the Dynamic load balance modes is configured and dynamic load balance is enabled on LAGs, the configuration is applied after reboot.

## 4/ Configuring System Information

Use the features in the System feature menu to define the switch's relationship to its environment.

### 4.1 Viewing the Dashboard

After a successful login, the Dashboard page displays. This page provides a brief overview of the system.

To navigate to the Dashboard, click System > Summary > Dashboard in the navigation menu.

Figure 8: System Dashboard

The screenshot displays the FASTPATH System Dashboard. At the top, there are navigation tabs for System, Switching, Routing, Security, and QoS. Below these are sub-tabs for Dashboard, Description, Inventory, and MAC Address Table. The main content area is divided into several sections:

- System Information:** A table with fields like System Description, System Name, System Location, System Contact, IP Address, Burned In MAC Address, Service Port IP Address, Service Port MAC Address, and System Up Time.
- Device Information:** A table with fields like Machine Type, Machine Model, Serial Number, FRU Number, Maintenance Level, Software Version, and Operating System.
- System Resource Usage:** Shows CPU Utilization (60 Second Average) at 42% and Memory Usage at 56% with progress bars.
- Disk Space Utilization:** Shows Total Disk Space (Kbytes) at 126,976, Free Disk Space (Kbytes) at 85,036, Used Disk Space (Kbytes) at 41,940, and Disk Usage at 33% with a progress bar.
- Logged In Users:** A table with columns for User Name, Connection From, and Idle Time, showing two active users.
- Recent Log Entries:** A table with columns for Log Time, Severity, and Description, showing several session-related log entries.

A Refresh button is located at the bottom of the dashboard area. The footer of the page indicates the copyright year 2016.

**Table 2: Dashboard Fields**

Field	Description
System Description	The product name of this device.
System Name	The configured name used to identify this device.
System Location	The configured location of this device.
System Contact	The configured contact person for this device.
IP Address	The IP address assigned to the network interface. The network interface is the logical interface that allows remote management of the device via any of the front-panel switch ports.
Burned In MAC Address	The device burned-in universally-administered media access control (MAC) address of the base system.
Service Port IP Address	The IP address assigned to the service port. The service port provides remote management access to the device. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.
Service Port MAC Address	The device burned-in universally-administered media access control (MAC) address of the service port.
System Up Time	The time in days, hours, minutes and seconds since the system was last reset.
Machine Type	The device hardware type or product family.
Machine Model	The model identifier, which is usually related to the Machine Type.
Serial Number	The unique device serial number.
FRU Number	The field replaceable unit number.
Maintenance Level	The device hardware change level identifier.
Software Version	The release.version.maintenance number of the software currently running on the device. For example, if the release is 1, the version is 2 and the maintenance number is 4, this version number is displayed as 1.2.4.
Operating System	The device operating system type and version identification information.
CPU Utilization (60 Second Average)	The percentage of CPU utilization for the entire system averaged over the past 60 seconds.
Memory Usage	The percentage of total available system memory (RAM) that is currently in use.
Disk Usage	The percentage of total available disk space that is currently in use.
Logged In Users	A brief summary indicating all other users currently logged into the device. The Idle Time field gives an indication of user activity, with a smaller time value denoting more recent access to the system.
Recent Log Entries	A brief list of the newest entries recorded in the system log.

Click Refresh to reload the page and refresh the Dashboard.

## 4.2 Viewing ARP Cache

The ARP cache is a table maintained locally in each station on a network. ARP cache entries are learned by examining the source information in the ARP packet payload fields, regardless of whether it is an ARP request or response. Thus, when an ARP request is broadcast to all stations on a LAN segment or virtual LAN (VLAN), every recipient has the opportunity to store the sender's IP and MAC address in their respective ARP cache. The ARP response, being unicast, is normally seen only by the requester, who stores the sender information in its ARP cache. Newer information always replaces existing content in the ARP cache.

The ARP cache can support 1024 entries, although this size is user-configurable to any value less than 1024. When multiple network interfaces are supported by a device, as is typical of a router, either a single ARP cache is used for all interfaces, or a separate cache is maintained per interface. While the latter approach is useful when network addressing is not unique per interface, this is not the case for Ethernet MAC address assignment so a single ARP cache is employed.

To display the system ARP cache, click System > Status > ARP Cache page in the navigation menu.

**Figure 9: ARP Cache**

The screenshot shows the 'ARP Cache' page in a web interface. At the top, there are tabs for 'ARP Cache', 'Resource Status', and 'Resource Configuration'. The main content area has a red header with 'ARP Cache' and a help icon. Below the header, there is a filter section with 'Display All rows', 'Showing 1 to 1 of 1 entries', and a 'Filter:' input field. The table below has three columns: 'MAC Address', 'IP Address', and 'Interface'. The first row contains the values '00:16:9C:E1:D8:00', '10.27.8.1', and 'Management'. Below the table are navigation links: 'First', 'Previous', '1', 'Next', 'Last'. At the bottom of the table area are two buttons: 'Refresh' and 'Clear Entries'. The footer of the page reads '© Broadcom Corporation 2000-2012'.

**Table 3: ARP Cache Fields**

Field	Description
MAC Address	Displays the physical (MAC) address of the system in the ARP cache.
IP Address	Displays the IP address associated with the system's MAC address.
Interface	Displays the unit, slot, and port number being used for the connection. For non-stacking systems, only the slot and port number is displayed. For units that have a service port, the service port will be listed as Management in this field.

Click Refresh to reload the page and refresh the ARP cache view. Click Clear Entries to clear all entries from the table. The table will be repopulated as new addresses are learned.

## 4.3 Viewing Inventory Information

Use the Inventory Information page to display the switch's Vital Product Data, which is stored in non-volatile memory at the factory.

To display the inventory information, click System > Summary > Inventory page in the menu.

Figure 10: Inventory Information

System Inventory Information	
Management Unit Number	1
System Description	Broadcom Triumph2 56634 Development System – 48 GE, 4 TENGIG, R.8.4.1, Linux 2.6.34.6
Machine Type	Broadcom Triumph2 56634 Development System – 48 GE, 4 TENGIG
Machine Model	BCM-56634-48
Serial Number	none
FRU Number	
Part Number	BCMS6634
Maintenance Level	A
Manufacturer	0xbc00
Burned In MAC Address	00:10:18:82:15:2F
Software Version	R.8.4.1
Operating System	Linux 2.6.34.6
Network Processing Device	BCMS6634_B0
Additional Packages	FASTPATH BGP-4 FASTPATH QoS FASTPATH IPv6 FASTPATH IPv6 Management FASTPATH Stacking

Refresh

© Broadcom Corporation 2000-2012

Table 4: Inventory Information Fields

Field	Description
Management Unit Number	Unit number that corresponds to the stack manager. This field is available only on switches that support stacking.
System Description	The product name of this switch.
Machine Type	The machine type of this switch.
Machine Model	The model within the machine type.
Serial Number	The unique serial number for this switch.
FRU Number	The field replaceable unit number.
Part Number	The manufacturing part number.
Maintenance Level	The identification of the hardware change level.
Manufacturer	The two-octet code that identifies the manufacturer.
Burned In MAC Address	The burned-in universally administered MAC address of this switch.
Software Version	The release version.maintenance number of the code currently running on the switch. For example, if the release is 1, the version is 2 and the maintenance number is 4, the format is 1.2.4.
Operating System	The operating system currently running on the switch.
Network Processing Device	Identifies the network processor hardware.
Additional Packages	A list of the optional software packages installed on the switch, if any. For example, FASTPATH IPv6, or FASTPATH Multicast.



## 4.4 Viewing the System Firmware Status

The pages in the Firmware folder allow you to view and monitor the system firmware status. The Firmware folder has links to the following pages.

### 4.4.1 Dual Image Status

The Dual Image feature allows the switch to have two FASTPATH software images in permanent storage. One image is the active image, and the second image is the backup. This feature reduces the system down-time during upgrades and downgrades. You can use the Dual Image Status page to view information about the system images on the device.

To display the Dual Image Status page, click System > Firmware > Status in the navigation menu.

Figure 11: Dual Image Status



Unit	Active	Backup	Current Active	Next Active
1	<none>	<none>	<none>	<none>

Image Description	
Active	
Backup	

© Broadcom Corporation 2000–2012

Table 5: Dual Image Status Fields

Field	Description
Unit	Displays the unit ID of the switch.
Active	Displays the version of the active code file.
Backup	Displays the version of the backup code file.
Current Active	Displays the currently active image on this unit.
Next Active	Displays the image to be used on the next restart of this unit.
Active Description	Displays the description associated with the active code file.
Backup Description	Displays the description associated with the backup code file.

Click Refresh to display the latest information from the router.

For information about how to update or change the system images, see [Section 4.17: "Using System Utilities"](#).

### 4.4.2 Dual Image Configuration and Upgrade

Use the Dual Image Configuration and Upgrade feature to transfer a new firmware (code) image to the device, select which image to load during the next boot cycle, and add a description to each image on the device. The device uses the HTTP protocol to transfer the image, and the image is saved as the backup image.

To display the Dual Image Configuration and Upgrade page, click System > Firmware > Configuration and Upgrade in the navigation menu.

Figure 12: Dual Image Configuration and Upgrade

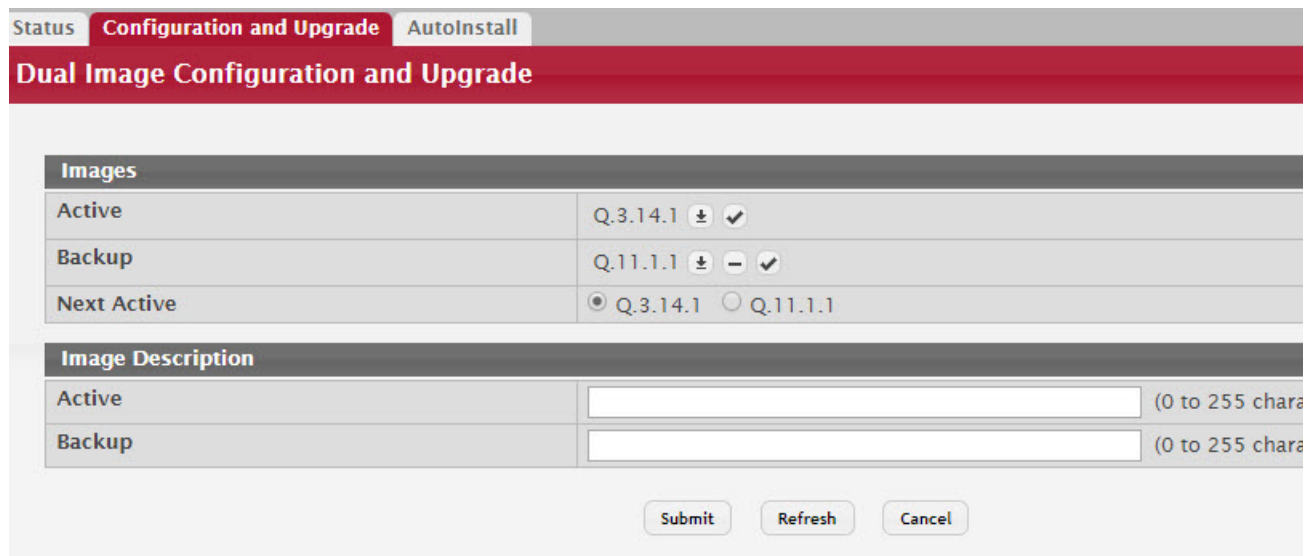


Table 6: Dual Image Configuration and Upgrade Fields

Field	Description
Unit	Use this field to select the unit with the code image to activate, upgrade, delete, or describe.
Active	<p>The active code file version. Use the icons to the right of the field to perform the file transfer.</p> <ul style="list-style-type: none"> <li>To transfer a new code image to the device, click the File Transfer icon. The Firmware Upgrade window opens. Click Choose File to browse to the file to transfer. After you select the appropriate file, click Begin Transfer to launch the HTTP transfer process. The active image is overwritten by the file that you transfer.</li> <li>To verify that the active image has a proper digital signature, click the Check icon. You must confirm the action before the image is verified. <input checked="" type="checkbox"/></li> </ul>
Backup	<p>The backup code file version. Use the icons to the right of the field to perform the following tasks:</p> <ul style="list-style-type: none"> <li>To transfer a new code image to the device, click the File Transfer icon. The Firmware Upgrade window opens. Click Choose File to browse to the file to transfer. After you select the appropriate file, click Begin Transfer to launch the HTTP transfer process. If a backup image already exists on the device, it is overwritten by the file that you transfer.</li> <li>To delete the backup image from permanent storage, click the – (minus) icon. You must confirm the action before the image is deleted.</li> <li>To verify that the backup image has a proper digital signature, click the Check icon. You must confirm the action before the image is verified. <input checked="" type="checkbox"/></li> </ul>
Next Active	Use this field to select the image version to load the next time this unit reboots.
Active Description	Use this field to specify a description to associate with the image that is currently the active code file.
Backup Description	Use this field to specify a description to associate with the image that is currently the backup code file.
Select File	<p>These three are all described in Help but I don't see them in the UI</p> <p>Use this field to provide option to browse to the directory where the file is located and select the file to transfer to the device.</p>
Digital Signature Verification	When this option is checked, the file download will be verified with the digital signature.
Status	Provides information about the status of the file transfer.

### 4.4.3 AutoInstall

The AutoInstall feature enables the configuration of a switch automatically when the device is turned on and, during the boot process, no configuration file is found in device storage. By communicating with a DHCP server, AutoInstall obtains an IP address for the switch and an IP address for a TFTP server. AutoInstall attempts to download a configuration file from the TFTP server and install in on the switch.

The DHCP server that the switch communicates with must provide the following information:

- The IP address and subnet mask (option 1) to be assigned to the switch.
- The IP address of a default gateway (option 3), if needed for IP communication.
- The identification of the TFTP server from which to obtain the boot file. This is given by any of the following fields, in the priority shown (highest to lowest):
  - The **sname** field of the DHCP reply.
  - The **hostname** of the TFTP server (option 66). Either the TFTP address or name is specified (not both) in most network configurations. If a TFTP hostname is given, a DNS server is required to translate the name to an IP address.
  - The **IP address** of the TFTP server (option 150).
  - The **address** of the TFTP server supplied in the **siaddr** field.
  - The **name** of the configuration file (boot file or option 67) to be downloaded from the TFTP server. The boot file name must have a file type of \*.cfg.
- The IP addresses of DNS name servers (option 6). The IP addresses of DNS name servers should be returned from the DHCP server only if the DNS server is in the same LAN as the switch performing AutoInstall. A DNS server is needed to resolve the IP address of the TFTP server if only the sname or option 66 values are returned to the switch.

After obtaining IP addresses for both the switch and the TFTP server, the AutoInstall feature attempts to download a host-specific configuration file using the boot file name specified by the DHCP server. If the switch fails to obtain the file, it will retry indefinitely.

To display the AutoInstall page, click System > Firmware> AutoInstall.

Figure 13: AutoInstall

AutoInstall Configuration	
Admin Mode	<input type="radio"/> Start <input checked="" type="radio"/> Stop
Persistent Mode	<input type="checkbox"/>
AutoSave Mode	<input type="checkbox"/>
AutoReboot Mode	<input checked="" type="checkbox"/>
Retry Count	<input type="text" value="3"/> (1 to 3)
Status	AutoInstall is completed.

© Broadcom Corporation 2000-2012

Table 7: AutoInstall Fields

Field	Definition
Admin Mode	The current administrative mode of the AutoInstall feature: <ul style="list-style-type: none"> <li>• Start — AutoInstall is enabled, and the feature will attempt to automatically configure the device during the next boot cycle.</li> <li>• Stop — AutoInstall is disabled. The automatic process will begin only if no configuration file is located during the next boot cycle.</li> </ul>
Persistent Mode	If this option is selected, the settings you configure on this page are automatically saved to persistent memory in the startup-config file when you apply the changes. If this option is not selected, the device treats these settings like any other applied changes (i.e. the changes are not retained across a reboot unless you save the configuration).
AutoSave Mode	If this option is selected, the downloaded configuration is automatically saved to persistent storage. If this option is not selected, you must explicitly save the downloaded configuration in non-volatile memory for the configuration to be available for the next reboot.
AutoReboot Mode	If this option is selected, the switch automatically reboots after a new image is successfully downloaded and makes the downloaded image the active image. If this option is not selected, the device continues to boot with the current image. The downloaded image will not become the active image until the device reboots.
Retry Count	When attempting to retrieve the DHCP-specified configuration file, this value represents the number of times the TFTP client on the device tries to use unicast requests before reverting to broadcast requests.
Status	The current status of the AutoInstall process.

Click Refresh to display the most recently configured AutoInstall state from the switch.

## 4.5 Viewing System Resources

Use the System Resources page to display the following memory information for the switch:

- Free memory
- Allocated memory
- CPU utilization by task
- Total CPU utilization at the following intervals:
  - Five seconds
  - One minute
  - Five minutes

To display the Resource Status page, click System > Status > Resource Status in the navigation menu.

Figure 14: System Resource Status

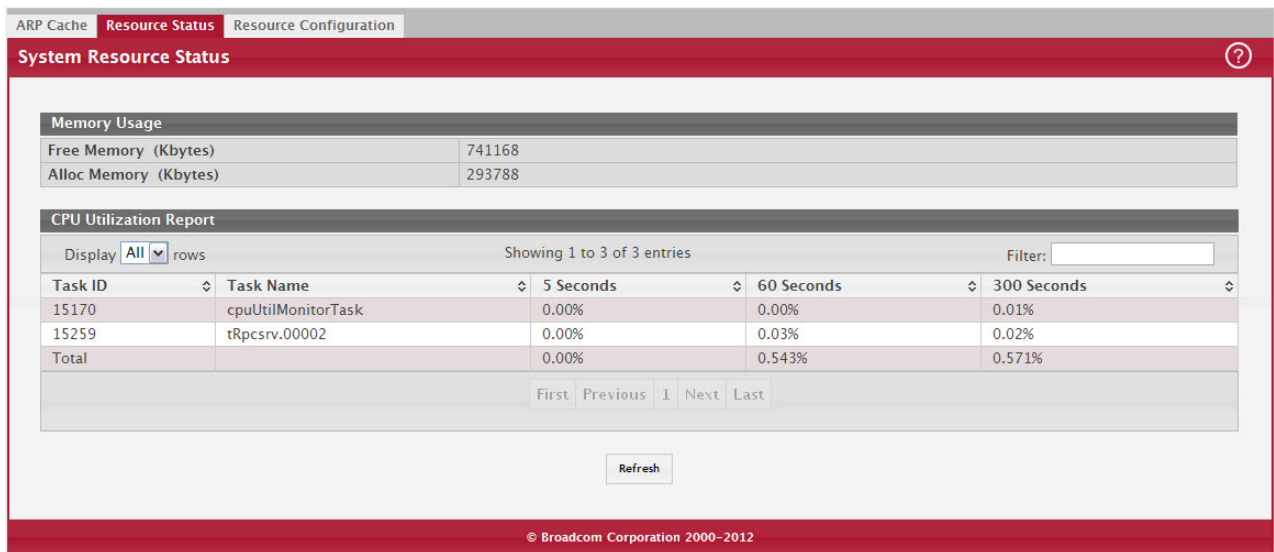


Table 8: System Resource Status Fields

Field	Description
Free Memory	Displays the available Free Memory on the switch.
Alloc Memory	Displays the allocated Memory for the switch.
Task Id	Displays the Id of running tasks.
Task Name	Displays the name of the running tasks.
CPU Utilization Report	Displays the Total CPU Utilization in terms of percentage. Total CPU Utilization is shown in the following intervals: <ul style="list-style-type: none"> <li>• 5 seconds</li> <li>• 60 seconds</li> <li>• 300 seconds</li> </ul>

To display the Resource Configuration page, click System > Status > Resource Configuration in the navigation menu.

Figure 15: System Resources Configuration

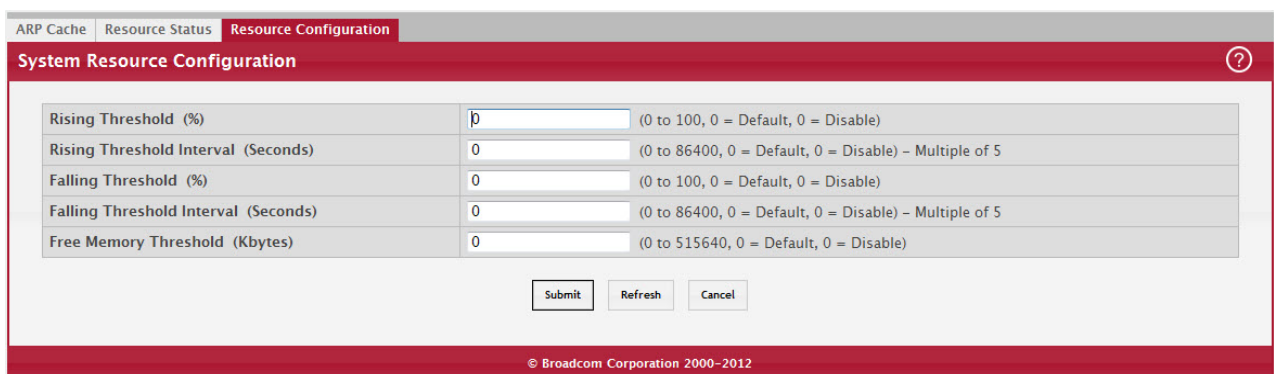


Table 9: System Resource Configuration Fields

Field	Description
Rising Threshold	The CPU Rising utilization threshold in percentage. A zero percent threshold indicates CPU Utilization Notification feature is disabled.
Rising Threshold Interval	The CPU Rising threshold interval in seconds. The time interval is configured in multiples of 5. A time interval of zero seconds indicates CPU Utilization Notification feature is disabled.
Falling Threshold	The CPU Falling utilization threshold in percentage. Configuration of this field is optional. If configured, the Falling threshold value must be equal to or less than the Rising threshold value. If not configured, it takes the same value as the Rising threshold.
Falling Threshold Interval	The CPU Falling threshold interval in seconds. Configuration of this field is optional. If configured, the Falling interval value must be equal to or less than the Rising interval value. If not configured, it takes the same value as the Rising interval. The time interval is configured in multiples of 5.
Free Memory Threshold	The CPU Free Memory threshold in kilobytes. A zero threshold value indicates CPU Free Memory Notification feature is disabled.

Click Submit to send the updated configuration to the switch. Click Refresh to update the page with the most current information. Click Cancel exit the page.

## 4.6 Selecting the SDM Template

A Switch Database Management (SDM) template is a description of the maximum resources a switch or router can use for various features. Different SDM templates allow different combinations of scaling factors, enabling different allocations of resources depending on how the device is used. In other words, SDM templates enable you to reallocate system resources to support a different mix of features based on your network requirements.

### NOTICE

If you attach a unit to a stack and its template does not match the stack's template, then the new unit will automatically reboot using the template used by other stack members. To avoid the automatic reboot, you may first set the template to the template used by existing members of the stack. Then power off the new unit, attach it to the stack, and power it on.

To display the SDM Template Preference page, click System > Advanced Configuration > SDM > SDM in the navigation menu.

Figure 16: SDM Template Preference

The screenshot shows the 'SDM Template Preference' page. At the top, there is a red header with 'SDM' and a question mark icon. Below the header, there are two configuration fields: 'Active SDM Template' set to 'IPv4-routing Default' and 'SDM Template on the Next Reload' set to 'IPv4-routing Default' with a refresh icon. A 'Summary' table is displayed below, showing resource usage for two templates: 'IPv4-routing Default' and 'Data Center Plus - IPv4'. The table has columns for ARP Entries, IPv4 Unicast Routes, IPv6 NDP Entries, IPv6 Unicast Routes, ECMP Next Hops, IPv4 Multicast Routes, and IPv6 Multicast Routes. A 'Refresh' button is located below the table. At the bottom, there is a copyright notice: '© Broadcom Corporation 2000-2013'.

SDM Template	ARP Entries	IPv4 Unicast Routes	IPv6 NDP Entries	IPv6 Unicast Routes	ECMP Next Hops	IPv4 Multicast Routes	IPv6 Multicast Routes
IPv4-routing Default	256	128	0	0	4	0	0
Data Center Plus - IPv4	256	128	0	0	4	0	0

Table 10: SDM Template Preference

Field	Description
Active SDM Template	Displays the SDM Template that is currently active.
SDM Template on the Next Reload	Select the template that will become active after the next reboot: <ul style="list-style-type: none"> <li>Dual IPv4 and IPv6 — filters subsequent template choices to those that support both IPv4 and IPv6. There is only one such template, and it is selected using the keyword default.</li> <li>IPv4 Routing (default) — filters subsequent template choices to those that support IPv4, and not IPv6. The default IPv4-only template maximizes the number of IPv4 unicast routes, while limiting the number of ECMP next hops in each route to 4. The data-center template supports increases the number of ECMP next hops to 16 and reduces the number of routes.</li> <li>IPv4 Data Center — this template sets the IPv6 scaling factors to 0 and sets the IPv4 unicast scaling factors to those used in builds with IPv6. The IPv4 Data Center template increases the maximum number of ECMP next hops from 4 to 16.</li> </ul>
SDM Template (Summary)	Identifies the available templates.
ARP Entries	The maximum number of entries in the IPv4 Address Resolution Protocol (ARP) cache for routing interfaces.
IPv4 Unicast Routes	The maximum number of IPv4 unicast forwarding table entries.
IPv6 NDP Entries	The maximum number of IPv6 Neighbor Discovery Protocol (NDP) cache entries.
IPv6 Unicast Routes	The maximum number of IPv6 unicast forwarding table entries.
ECMP Next Hops	The maximum number of next hops that can be installed in the IPv4 and IPv6 unicast forwarding tables.
IPv4 Multicast Routes	The maximum number of IPv4 multicast forwarding table entries.
IPv6 Multicast Routes	The maximum number of IPv6 multicast forwarding table entries.

Click Refresh to display the latest information from the router.

## 4.7 Defining General Device Information

The Configuration submenu in the System menu contains links to pages that allow you to configure device parameters.

### 4.7.1 System Description

After a successful login, the System Description page displays. Use this page to configure and view general device information.

To display the System Description page, click System > Summary > Description in the navigation menu.

Figure 17: System Description

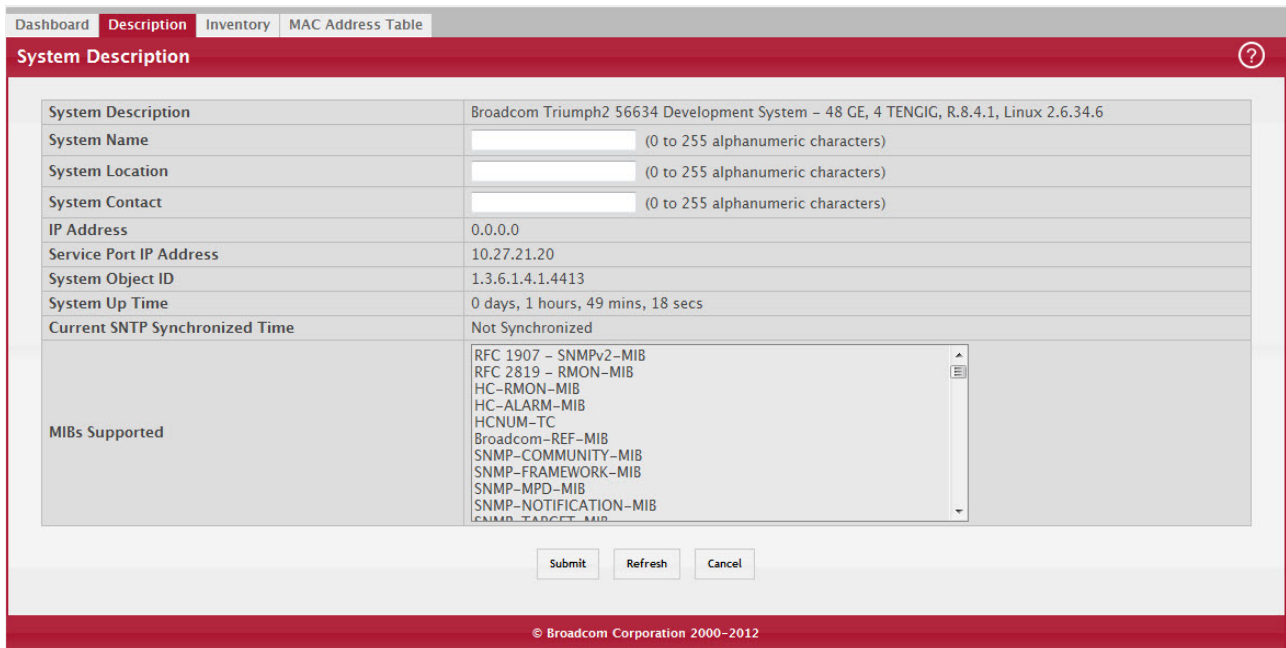


Table 11: System Description Fields

Field	Description
System Description	The product name of this switch.
System Name	Enter the name you want to use to identify this switch. You may use up to 31 alphanumeric characters. The factory default is blank.
System Location	Enter the location of this switch. You may use up to 31 alpha-numeric characters. The factory default is blank.
System Contact	Enter the contact person for this switch. You may use up to 31 alpha-numeric characters. The factory default is blank.
IP Address	The IP Address assigned to the network interface. To change the IP address, see <a href="#">Section 4.7.4: "IPv4 Network Connectivity Configuration"</a> .
Service Port IP Address	The IP address assigned to the service port. The service port provides remote management access to the device. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.
System Object ID	The base object ID for the switch's enterprise MIB.
System Up Time	Displays the number of days, hours, and minutes since the last system restart.
Current SNMP Synchronized Time	Displays currently synchronized SNMP time in UTC. If no SNMP server has been configured and the time is not synchronized, this field displays <i>Not Synchronized</i> . To specify an SNMP server, see <a href="#">Section 4.22: "Configuring SNMP Settings"</a> .
MIBs Supported	Displays the list of MIBs supported by the management agent running on this switch.

### 4.7.1.1 Defining System Information

1. Open the System Description page.
2. Define the following fields: System Name, System Contact, and System Location.
3. Scroll to the bottom of the page and click Submit.

The system parameters are applied, and the device is updated.



If you want the switch to retain the new values across a power cycle, you must perform a save.

## NOTICE

### 4.7.2 Switch Configuration

IEEE 802.3x flow control works by pausing a port when the port becomes oversubscribed and dropping all traffic for small bursts of time during the congestion condition. This can lead to high-priority and/or network control traffic loss. When 802.3x flow control is enabled, lower speed switches can communicate with higher speed switches by requesting that the higher speed switch refrains from sending packets. Transmissions are temporarily halted to prevent buffer overflows.

To display the Switch Configuration page, click System > Basic Configuration > Switch in the navigation menu.

Figure 18: Switch 802.3x Flow Control

The screenshot shows the 'Switch Configuration' page. At the top, there is a red header with the title 'Switch Configuration' and a help icon. Below the header, there are two main configuration fields:

- 802.3x Flow Control Mode:** This field has two radio buttons: 'Disable' (which is selected) and 'Enable'.
- MAC Address Aging Interval (Seconds):** This field is a text input box containing the value '300', with a range '(10 to 1000000)' indicated to the right.

At the bottom of the configuration area, there are three buttons: 'Submit', 'Refresh', and 'Cancel'. A footer at the very bottom of the page reads '© Broadcom Corporation 2000-2012'.

Table 12: Switch Configuration Fields

Field	Description
IEEE 802.3x Flow Control Mode	The 802.3x flow control mode on the switch. IEEE 802.3x flow control works by pausing a port when the port becomes oversubscribed. This allows lower-speed switches to communicate with higher-speed switches. A lower-speed or congested switch can send a PAUSE frame requesting that the peer device refrain from sending packets. Transmissions are temporarily halted to prevent buffer overflows. The options are as follows: <ul style="list-style-type: none"> <li>• Disabled – The switch does not send PAUSE frames if the port buffers become full.</li> <li>• Enabled – The switch can send PAUSE frames to a peer device if the port buffers become full.</li> </ul>
MAC Address Aging Interval	The MAC address table (forwarding database) contains static entries, which never age out, and dynamically-learned entries, which are removed if they are not updated within a given time. Specify the number of seconds a dynamic address should remain in the MAC address table after it has been learned.

If you change the mode, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

### 4.7.3 IP Address Conflict Detection

Use the IP Address Conflict Detection page to run the IP Address Conflict Detection tool, which detects IP address conflicts for IPv4 addresses. When a conflict is detected, the switch updates the status on the page, generates an SNMP trap, and a logs a message noting the conflict.

To display the IP Address Conflict Detection page, click System > Utilities > IP Address Conflict in the navigation menu.

Figure 19: IP Address Conflict Detection



Table 13: IP Address Conflict Detection Fields

Field	Description
IP Address Conflict Currently Exists	Shows whether an address conflict has been detected since status was last reset.
Last Conflicting IP Address	The IP address of the interface that was last found to be conflicting. This field displays only if a conflict has been detected since the switch was last reset.
Last Conflicting MAC Address	The MAC address of the remote host associated with the IP address that was last found to be conflicting. This field displays only if a conflict has been detected since the switch was last reset.
Time Since Conflict Detected	The time elapsed (displayed in days, hours, minutes, seconds) since the last conflict was detected (provided a reset did not occur in the meantime). This field displays only if a conflict has been detected since the switch was last reset.

To run the tool and check for possible address conflicts, click Run Conflict Detection. If the conflict detection status is true, click Reset Conflict Detection Status to clear the information and run the tool again.

#### 4.7.4 IPv4 Network Connectivity Configuration

The network interface is the logical interface used for in-band connectivity with the switch via any of the switch's front panel ports. The configuration parameters associated with the switch's network interface do not affect the configuration of the front panel ports through which traffic is switched or routed.

The IPv4 Network Connectivity page allows you to change the IPv4 information using the Web interface. To access the page, click System > Connectivity > IPv4 in the navigation menu.

Figure 20: Network Connectivity Configuration for IPv4

Table 14: Network Connectivity Configuration for IPv4 Fields

Field	Description
Network Configuration Protocol	Specify what the switch should do following power-up. The factory default is None. The options are as follows: <ul style="list-style-type: none"> <li>None: Do not send any requests following power-up.</li> <li>Bootp: Transmit a Bootp request.</li> <li>DHCP: Transmit a DHCP request.</li> </ul>
DHCP Client Identifier	The DHCP Client Identifier (Option 61) is used by DHCP clients to specify their unique identifier. DHCP servers use this value to index their database of address bindings. This value is expected to be unique for all clients in an administrative domain. The Client Identifier string will be displayed beside the check box once DHCP is enabled on the port on which the Client Identifier option is selected. This web page will need to be refreshed once this change is made.
IP Address	The IP address of the network interface. The factory default value is 0.0.0.0 <b>Note:</b> Each part of the IP address must start with a number other than zero. For example, IP addresses 001.100.192.6 and 192.001.10.3 are not valid.
Subnet Mask	The IP subnet mask for the interface. The factory default value is 0.0.0.0.
Default Gateway	The default gateway for the IP interface. The factory default value is 0.0.0.0.
MAC Address Type	Specify whether the burned-in or the locally administered MAC address should be used for in-band connectivity. The factory default is to use the burned-in MAC address
Burned-in MAC Address	This read-only field displays the MAC address that is burned-in to the network card at the factory. This MAC address is used for in-band connectivity if you choose not to configure a locally administered address.
Locally Administered MAC Address	Specifies a locally administered MAC address for in-band connectivity instead of using the burned-in universally administered MAC address. In addition to entering an address in this field, you must also set the MAC address type to locally administered. Enter the address as twelve hexadecimal digits (6 bytes) with a colon between each byte. Bit 1 of byte 0 must be set to a 1 and bit 0 to a 0, i.e. byte 0 must have a value between x'40' and x'7F'.
Management VLAN ID	Specify the management VLAN ID of the switch. It may be configured to any value in the range of (1 to 4093). The management VLAN is used for management of the switch. This field is configurable for administrative users and read-only for other users.

If you change any of the network connectivity parameters, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

Click Renew DHCP IPv4 Address to force the interface to release the current DHCP-assigned information and submit a request for new information.

### 4.7.5 IPv6 Network Connectivity

IPv6 is the next generation of the Internet Protocol. With 128-bit addresses, versus 32-bit addresses for IPv4, IPv6 solves the address depletion issues seen with IPv4 and removes the requirement for Network Address Translation (NAT), which is used in IPv4 networks to reduce the number of globally unique IP addresses required for a given network. Its aggregate addresses can dramatically reduce the size of the global routing table through well known address combinations. Security is more integrated and network configuration is simplified yet more flexible.

IPv6 can coexist with IPv4. As with IPv4, IPv6 routing can be enabled on physical and VLAN interfaces. Each L3 routing interface can be used for IPv4, IPv6, or both. IP protocols running over L3 (for example, UDP and TCP) do not change with IPv6. For this reason, a single CPU stack is used for transport of both IPv4 and IPv6, and a single sockets interface provides access to both. Routing protocols are capable of computing routes for one or both IP versions.

CLI commands are not available for all the IPv6 pages.

**NOTICE**

Use the IPV6 Network Connectivity page to configure and view IPv6 information on the network interface. The network interface is the logical interface that allows remote management of the device via any of the front-panel switch ports. To enable management of the device over an IPv6 network by using a Web browser, SNMP, Telnet, or SSH, you must first configure the device with the appropriate IPv6 information. The configuration parameters associated with the network interface do not affect the configuration of the front-panel ports through which traffic is switched or routed.

To display the page, click System > Connectivity > IPv6 in the navigation menu.

Figure 21: IPv6 Network Connectivity Configuration

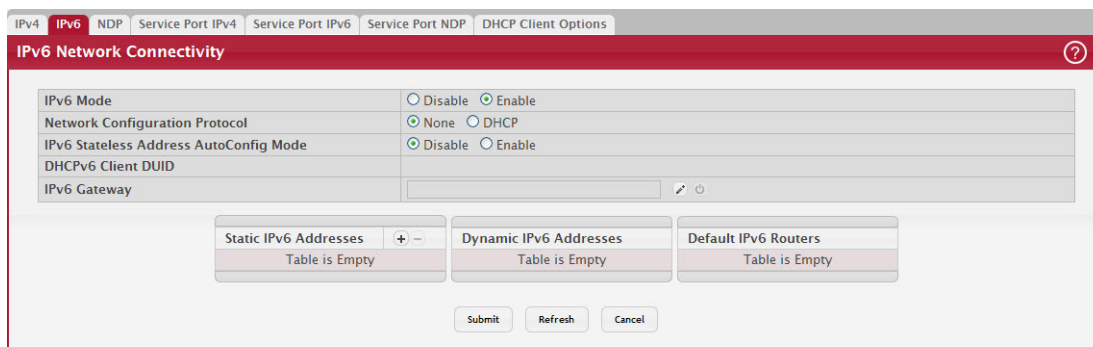


Table 15: IPv6 Network Connectivity Configuration Fields

Field	Description
IPv6 Mode	Enables or disables the IPv6 administrative mode on the network interface.
Network Configuration Protocol	Specify whether the device should attempt to acquire network information from a DHCPv6 server. Selecting None disables the DHCPv6 client on the network interface.
IPv6 Stateless Address AutoConfig Mode	Sets the IPv6 stateless address autoconfiguration mode on the network interface. <ul style="list-style-type: none"> <li>Enabled – The network interface can acquire an IPv6 address through IPv6 Neighbor Discovery Protocol (NDP) and through the use of Router Advertisement messages.</li> <li>Disabled – The network interface will not use the native IPv6 address autoconfiguration features to acquire an IPv6 address.</li> </ul>
DHCPv6 Client DUID	The client identifier used by the DHCPv6 client (if enabled) when sending messages to the DHCPv6 server.

Table 15: IPv6 Network Connectivity Configuration Fields (Continued)

Field	Description
IPv6 Gateway	The default gateway for the IPv6 network interface. To configure this field, click the Edit icon in the row. To reset the field to the default value, click the Reset icon in the row.
Static IPv6 Addresses	Lists the manually configured static IPv6 addresses on the network interface. Use the buttons available in this table to perform the following tasks: To add an entry to the list, click the + (plus) button to open the Add IPv6 Address dialog and provide the following: <ul style="list-style-type: none"> <li>New IPv6 Address – Specify the IPv6 address to add to the interface.</li> <li>EUI Flag – Select this option to enable the Extended Universal Identifier (EUI) flag for IPv6 address, or clear the option to omit the flag.</li> </ul> To delete an entry from the list, click the – (minus) button associated with the entry to remove. To delete all entries from the list, click the – (minus) button in the heading row.
Dynamic IPv6 Addresses	Lists the IPv6 addresses on the network interface that have been dynamically configured through IPv6 autoconfiguration or DHCPv6.
Default IPv6 Routers	Lists the IPv6 address of each default router that has been automatically configured through IPv6 router discovery.

Click Refresh to update the information on the screen.

#### 4.7.6 Network Port IPv6 Neighbors

This page provides information about IPv6 neighbors the device has discovered through the network interface by using the Neighbor Discovery Protocol (NDP) and the manually configured static network port IPv6 neighbors.

To access this page, click System > Connectivity > IPv6 Neighbors.

Figure 22: Network Port IPv6 Neighbors

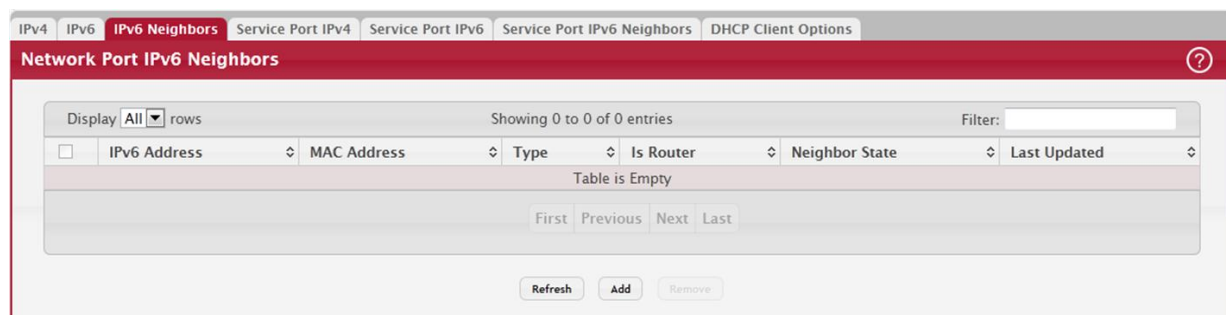


Table 16: Network Port IPv6 Neighbors Fields

Field	Description
IPv6 Address	Displays the IP address of the neighbor.
MAC Address	Displays the MAC address of the neighbor.
IS Router	Indicates whether the neighbor is a router. If the neighbor is a router, the value is TRUE. If the neighbor is not a router, the value is FALSE.
Type	The type of the neighbor entry, which is one of the following: <ul style="list-style-type: none"> <li>Static–The neighbor entry is manually configured.</li> <li>Dynamic–The neighbor entry is dynamically resolved.</li> <li>Local–The neighbor entry is a local entry.</li> <li>Other–The neighbor entry is an unknown entry.</li> </ul>

Table 16: Network Port IPv6 Neighbors Fields (Continued)

Field	Description
Neighbor State	<p>Specifies the state of the neighbor cache entry. Following are the states for dynamic entries in the IPv6 neighbor discovery cache:</p> <ul style="list-style-type: none"> <li>• Reachable—Positive confirmation was received within the last Reachable Time milliseconds that the forward path to the neighbor was functioning properly. While in REACH state, the device takes no special action as packets are sent.</li> <li>• Stale—More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. While in STALE state, the device takes no action until a packet is sent.</li> <li>• Delay—More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. A packet was sent within the last DELAY_FIRST_PROBE_TIME seconds. If no reachability confirmation is received within DELAY_FIRST_PROBE_TIME seconds of entering the DELAY state, send a neighbor solicitation message and change the state to PROBE.</li> <li>• Probe—A reachability confirmation is actively sought by resending neighbor solicitation messages every RetransTimer milliseconds until a reachability confirmation is received.</li> <li>• Unknown—The reachability status cannot be determined.</li> </ul>
Last Updated	Displays the time since the address was confirmed to be reachable.

Use the buttons to perform the following tasks:

- To add network port static IPv6 neighbor entry, click Add and configure the desired settings.
- To remove network port static IPv6 neighbor entries, select each static neighbor entry to remove and click Remove.

After you click Add or Edit, a window opens and allows you to configure Network Port IPv6 Neighbor settings.

You can configure the IPv6 Address and the MAC Address.

Figure 23: Add Network Port IPv6 Neighbor

Table 17: Add Network Port IPv6 Neighbor Fields

Field	Description
IPv6 Address	Use this field to enter the IP address of the neighbor.
MAC Address	Use this field to enter the MAC address of the neighbor.

#### 4.7.7 Service Port IPv4

Some platforms have a built-in service port that can serve as a dedicated network management interface. For systems that have the service port, the Service Port IPv4 Configuration page allows you to configure network information for the switch.

To access the Service Port Configuration page, click System > Connectivity > Service Port IPv4 in the navigation menu.

Figure 24: Service Port IPv4 Configuration

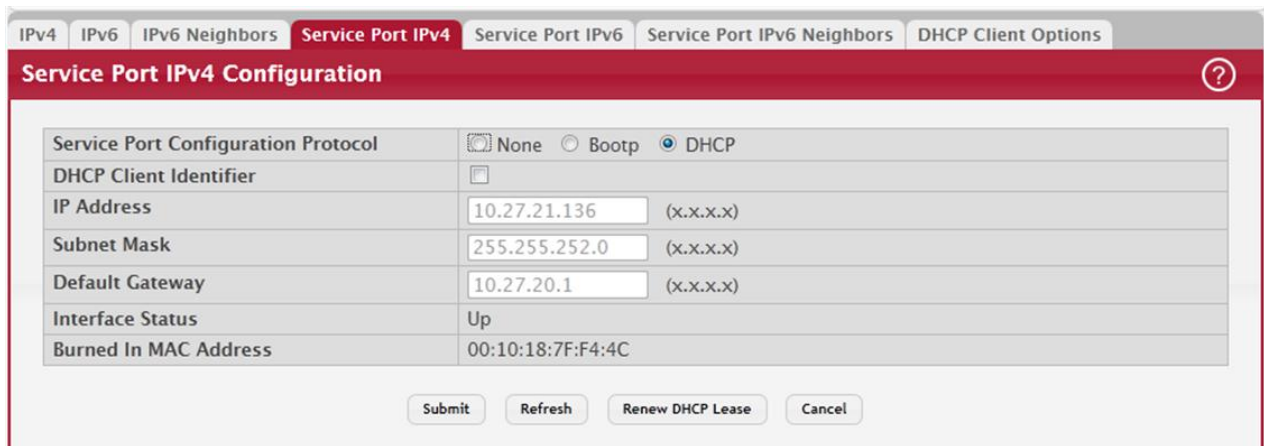


Table 18: Service Port IPv4 Configuration Fields

Field	Description
IPv4 Fields	These display IPv4 configuration information.
Service Port Configuration Protocol	Specify what the switch should do following power-up. The factory default is None. The options are as follows: <ul style="list-style-type: none"> <li>• None: Do not send any requests following power-up.</li> <li>• BootP: Transmit a BootP request.</li> <li>• DHCP: Transmit a DHCP request.</li> </ul>
IP Address	The IP address of the network interface. The factory default value is 0.0.0.0 <b>Note:</b> Each part of the IP address must start with a number other than zero. For example, IP addresses 001.100.192.6 and 192.001.10.3 are not valid.
Subnet Mask	The IP subnet mask for the interface. The factory default value is 0.0.0.0.
Default Gateway	The default gateway for the IP interface. The factory default value is 0.0.0.0.
Burned-in MAC Address	This read-only field displays the MAC address that is burned-in to the network card at the factory. This MAC address is used for in-band connectivity if you choose not to configure a locally administered address.

To renew the IPv4 address learned from a DHCP server on the service port, click Renew DHCP IPv4 Address.

If you change any of the parameters on this page, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

### 4.7.8 Service Port IPv6

Use this page to configure IPv6 network information on the service port. The service port is a dedicated Ethernet port for out-of-band management of the device. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.

To access the Service Port Configuration page, click System > Connectivity > Service Port IPv6 in the navigation menu.



Figure 25: Service Port IPv6 Configuration

Table 19: Service Port IPv6 Configuration Fields

Field	Description
IPv6 Mode	Enables or disables IPv6 mode on the interface.
Service Port Configuration Protocol	Specify whether the device should attempt to acquire network information from a DHCPv6 server. Selecting None disables the DHCPv6 client on the service port.
IPv6 Stateless Address AutoConfig Mode	Sets the IPv6 stateless address autoconfiguration mode on the service port. <ul style="list-style-type: none"> <li>Enabled – The service port can acquire an IPv6 address through IPv6 Neighbor Discovery Protocol (NDP) and through the use of Router Advertisement messages.</li> <li>Disabled – The service port will not use the native IPv6 address autoconfiguration features to acquire an IPv6 address.</li> </ul>
DHCPv6 Client DUID	The client identifier used by the DHCPv6 client (if enabled) when sending messages to the DHCPv6 server.
Static IPv6 Addresses	Lists the manually configured static IPv6 addresses on the service port interface. Use the buttons available in this table to perform the following tasks: <ul style="list-style-type: none"> <li>To add an entry to the list, click the + (plus) button to open the Add IPv6 Address dialog and provide the following: <ul style="list-style-type: none"> <li>New IPv6 Address – Specify the IPv6 address to add to the service port interface.</li> <li>EUI Flag – Select this option to enable the Extended Universal Identifier (EUI) flag for IPv6 address, or clear the option to omit the flag.</li> </ul> </li> <li>To delete an entry from the list, click the – (minus) button associated with the entry to remove.</li> <li>To delete all entries from the list, click the – (minus) button in the heading row.</li> </ul>
Dynamic IPv6 Addresses	Lists the IPv6 addresses on the service port interface that have been dynamically configured through IPv6 autoconfiguration or DHCPv6.
Default IPv6 Routers	Lists the IPv6 address of each default router that has been automatically configured through IPv6 router discovery.

To renew the IPv6 address learned from a DHCP server on the service port, click Renew DHCP IPv6 Address.

If you change any of the parameters on this page, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.



### 4.7.9 Service Port IPv6 Neighbors

This page provides information about IPv6 neighbors the device has discovered through the service port by using the Neighbor Discovery Protocol (NDP). The manually configured static service port IPv6 neighbors are also displayed.

To display the page, click System > Connectivity > Service Port IPv6 Neighbors

Figure 26: Service Port IPv6 Neighbors



Table 20: Service Port IPv6 Neighbors Fields

Field	Description
IPv6 Addresses	Displays the IP address of the neighbor.
MAC Address	Displays the MAC address of the neighbor.
Type	The type of the neighbor entry, which is one of the following: <ul style="list-style-type: none"> <li>• Static–The neighbor entry is manually configured.</li> <li>• Dynamic–The neighbor entry is dynamically resolved.</li> <li>• Local–The neighbor entry is a local entry.</li> <li>• Other–The neighbor entry is an unknown entry.</li> </ul>
Is Router	Indicates whether the neighbor is a router. If the neighbor is a router, the value is TRUE. If the neighbor is not a router, the value is FALSE.
Neighbor State	Specifies the state of the neighbor cache entry. Following are the states for dynamic entries in the IPv6 neighbor discovery cache: <ul style="list-style-type: none"> <li>• Reachable–Positive confirmation was received within the last Reachable Time milliseconds that the forward path to the neighbor was functioning properly. While in REACH state, the device takes no special action as packets are sent.</li> <li>• Stale–More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. While in STALE state, the device takes no action until a packet is sent.</li> <li>• Delay–More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. A packet was sent within the last DELAY_FIRST_PROBE_TIME seconds. If no reachability confirmation is received within DELAY_FIRST_PROBE_TIME seconds of entering the DELAY state, send a neighbor solicitation message and change the state to PROBE.</li> <li>• Probe–A reachability confirmation is actively sought by resending neighbor solicitation messages every RetransTimer milliseconds until a reachability confirmation is received.</li> <li>• Unknown–The reachability status cannot be determined.</li> </ul>
Last Updated	Displays the time since the address was confirmed to be reachable.

Use the buttons to perform the following tasks:

- To add service port static IPv6 neighbor entry, click Add and configure the desired settings.
- To remove service port static IPv6 neighbor entries, select each static neighbor entry to remove and click Remove.

After you click Add or Edit, a window opens and allows you to configure Service Port IPv6 Neighbor settings.

You can configure the IPv6 Address and the MAC Address.

Figure 27: Add Service Port IPv6 Neighbor

Table 21: Add Service Port IPv6 Neighbor Fields

Field	Description
IPv6 Address	Use this field to enter the IP address of the neighbor.
MAC Address	Use this field to enter the MAC address of the neighbor.

#### 4.7.10 DHCP Client Options

Use the DHCP Client Options page to configure DHCP client settings on the system.

To access the DHCP Client Options page, click System > Connectivity > DHCP Client Options in the navigation menu.

Figure 28: DHCP Client Options

Table 22: DHCP Client Options Fields

Field	Description
DHCP Vendor Class ID Mode	Enables/Disables the vendor class identifier mode.
DHCP Vendor Class ID String	The string added to DHCP requests as Option-60, that is, Vendor Class Identifier option.

#### 4.7.11 System Connectivity

Use the System Connectivity page to control access to the management interface by administratively enabling or disabling various access methods.

To display the System Connectivity page, click System > Management Access > System in the navigation menu.

Figure 29: System Connectivity Configuration

The screenshot shows the 'System Connectivity' configuration page. At the top, there are navigation tabs for System, Telnet, Outbound Telnet, Serial, CLI Banner, HTTP, HTTPS, and SSH. The 'System' tab is selected. Below the tabs, the page is titled 'System Connectivity'. The configuration is organized into sections:

- HTTP:** HTTP Admin Mode is set to  Enable.
- Telnet:** Telnet Server Admin Mode is set to  Enable. Allow New Sessions is checked with a .
- Outbound Telnet:** Allow New Sessions is checked with a .
- Secure HTTP:** HTTPS Admin Mode is set to  Disable.
- Secure Shell:** SSH Admin Mode is set to  Disable.

At the bottom of the page, there are three buttons: Submit, Refresh, and Cancel.

Table 23: System Connectivity Configuration Fields

Field	Description
HTTP Admin Mode	Enables or disables the HTTP administrative mode. When this mode is enabled, the device management interface can be accessed through a web browser using the HTTP protocol.
Telnet Server Admin Mode	Enables or disables the telnet administrative mode. When this mode is enabled, the device command-line interface (CLI) can be accessed through the telnet port. Disabling this mode disconnects all existing telnet connections and shuts down the telnet port in the device.
Telnet—Allow New Sessions	Enables or disables new telnet sessions. When this option is disabled, the system does not accept any new telnet sessions, but existing telnet sessions are unaffected.
Outbound Telnet—Allow New Sessions	Enables or disables new outbound telnet sessions. When this option is disabled, initiating telnet sessions from the system is not allowed.
HTTPS Admin Mode	Enables or disables the administrative mode of secure HTTP. When this mode is enabled, the device management interface can be accessed through a web browser using the HTTPS protocol.
SSH Admin Mode	Enables or disables the administrative mode of SSH. When this mode is disabled, all existing SSH connections remain connected until timed-out or logged out, but new SSH connections cannot be established.

If you change any of the parameters, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

#### 4.7.12 Telnet Session

Telnet is a terminal emulation TCP/IP protocol. ASCII terminals can be virtually connected to the local device through a TCP/IP protocol network. Telnet is an alternative to a local login terminal where a remote login is required.

The switch supports up to five simultaneous telnet sessions. All CLI commands can be used over a telnet session.

The Telnet Session Configuration page allows you to control inbound telnet settings on the switch. Inbound telnet sessions originate on a remote system and allow a user on that system to connect to the switch CLI.

To display the Telnet Session Configuration page, click System > Management Access > Telnet in the navigation menu.

Figure 30: Telnet Session Configuration

Field	Value	Range
Admin Mode	<input type="radio"/> Disable <input checked="" type="radio"/> Enable	
Telnet Port	23	(1 to 65535, 23 = Default)
Session Timeout (Minutes)	5	(1 to 160)
Maximum Number of Sessions	5	(0 to 5)
Allow New Sessions	<input checked="" type="checkbox"/>	

Table 24: Telnet Session Configuration Fields

Field	Description
Admin Mode	Enables or disables the telnet administrative mode. When enabled, the device may be accessed through the telnet port (23). Disabling this mode value disconnects all existing telnet connections and shuts down the telnet port in the device.
Telnet Port	The TCP port number on which the telnet server listens for requests. Existing telnet login sessions are not affected by a change in this value, although establishment of any new telnet sessions must use the new port number. <b>Note:</b> Before changing this value, check your system, e.g., using netstat, to make sure the desired port number is not currently being used by any other service.
Session Timeout (Minutes)	Specify how many minutes of inactivity should occur on a telnet session before the session is logged off. You may enter any number from 1 to 160. The factory default is 5. <b>Note:</b> When you change the timeout value, the new value is applied to all active and inactive sessions immediately. Any sessions that have been idle longer than the new timeout value are disconnected immediately.
Maximum Number of Sessions	From the drop-down menu, select how many simultaneous telnet sessions to allow. The maximum is 4, which is also the factory default. A value of 0 indicates that no outbound Telnet session can be established.
Allow New Sessions	Controls whether to allow new telnet sessions: <ul style="list-style-type: none"> <li>• Yes: Permits new telnet sessions until the maximum number allowed is reached.</li> <li>• No: New telnet sessions will not be allowed, but existing sessions are not disconnected.</li> </ul>

If you change any of the telnet parameters, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

### 4.7.13 Outbound Telnet Configuration

This page displays the current value of the outbound Telnet settings on the device. An outbound Telnet session is a Telnet session initiated from the CLI of the device to the Telnet client on a remote device.

To display the Outbound Telnet Configuration page, click System > Management Access > Outbound Telnet in the navigation menu.

Figure 31: Outbound Telnet Configuration

Table 25: Outbound Telnet Configuration Fields

Field	Description
Allow New Sessions	Controls whether new outbound Telnet sessions are allowed. Setting this value to Disable disallows any new outbound Telnet sessions from starting (although existing Telnet sessions are unaffected).
Maximum Number of Sessions	The maximum number of allowed outbound Telnet sessions from the device simultaneously.
Session Timeout	Outbound telnet session inactivity timeout value, in minutes. An outbound Telnet session is closed automatically if there is no activity within the configured amount of time.

If you change any of the outbound telnet parameters, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

#### 4.7.14 Serial Port

The Serial Port Configuration page allows you to change the switch's serial port settings. For a terminal or terminal emulator to communicate with the switch, the serial port settings on both devices must be the same. Some settings on the switch cannot be changed.

To view or configure the serial port settings on the switch, click System > Management Access > Serial in the navigation menu.

Figure 32: Serial Port

Table 26: Serial Port Fields

Field	Description
Serial Port Time Out (Minutes)	Indicates how many minutes of inactivity should occur on a serial port connection before the switch closes the connection. Enter a number between 0 and 160. The factory default is 5. Entering 0 disables the timeout.
Baud Rate (bps)	Select the default baud rate for the serial port connection from the menu. The factory default is 115200 baud
Character Size (Bits)	The number of bits in a character. This is always 8.
Parity	The parity method used on the serial port. It is always None.
Stop Bits	The number of stop bits per character. Its is always 1.
Flow Control	Whether hardware flow control is enabled or disabled. It is always disabled.

If you change any data, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

#### 4.7.15 CLI Banner Configuration

Use the CLI Banner Configuration page to configure a message that appears before the user prompt as a Pre-login banner. The message configured shows up on Telnet, SSH and Console connections.

To access the CLI Banner Configuration page, click System > Management Access > CLI Banner in the navigation menu.

Figure 33: CLI Banner Configuration

Table 27: CLI Banner Configuration Fields

Field	Description
CLI Banner Message	Text area for creating, viewing, or updating the CLI banner message. To to create the CLI banner message, type the desired message in the text area. If you reach the end of the line, the text wraps to the next line. The line might not wrap at the same location in the CLI. To create a line break (carriage return) in the message, press the Enter key on the keyboard. The line break in the text area will be at the same location in the banner message when viewed through the CLI.
Clear (Button)	Clears the CLI banner message from the device. After you click Clear, you must confirm the action. You can also clear the CLI banner by deleting the text in the CLI Banner Message field and clicking Submit.

Click Submit to apply the changes to the system. You must perform a save to make the changes persist across a reboot.

## 4.7.16 HTTP Configuration

Use the HTTP Configuration page to configure the HTTP server settings on the system.

To access the HTTP Configuration page, click System > Management Access > HTTP in the navigation menu.

Figure 34: HTTP Configuration

Table 28: HTTP Configuration Fields

Field	Description
HTTP Admin Mode	This select field is used to Enable or Disable the Administrative Mode of HTTP. The currently configured value is shown when the web page is displayed. The default value is Enable. If you disable the HTTP admin mode, access to the web interface is limited to secure HTTP, which is disabled by default.
HTTP Port	The TCP port number on which the HTTP server listens for requests. Existing HTTP login sessions are closed whenever this value is changed. All new HTTP sessions must use the new port number. <b>Note:</b> Before changing this value, check your system, e.g., using netstat, to make sure the desired port number is not currently being used by any other service.
HTTP Session Soft Timeout	This field is used to set the inactivity timeout for HTTP sessions. The value must be in the range of (1 to 60) minutes. A value of zero corresponds to an infinite timeout. The default value is 5 minutes. The currently configured value is shown when the web page is displayed.
HTTP Session Hard Time-out	This field is used to set the hard timeout for HTTP sessions. This timeout is unaffected by the activity level of the session. The value must be in the range of (1 to 168) hours. A value of zero corresponds to an infinite timeout. The default value is 24 hours. The currently configured value is shown when the web page is displayed.
Maximum Number of HTTP Sessions	This field is used to set the maximum allowable number of HTTP sessions. The value must be in the range of (0 to 16). The default value is 16. The currently configured value is shown when the web page is displayed.

If you make changes to the page, click Submit to apply the changes to the system.

## 4.7.17 HTTPS Configuration

Use this page to view and modify the Secure HTTP (HTTPS) settings on the device. HTTPS increases the security of web-based management by encrypting communication between the administrative system and the device.

To access the HTTPS Configuration page, click System > Management Access > HTTPS in the navigation menu.

Figure 35: HTTPS Configuration

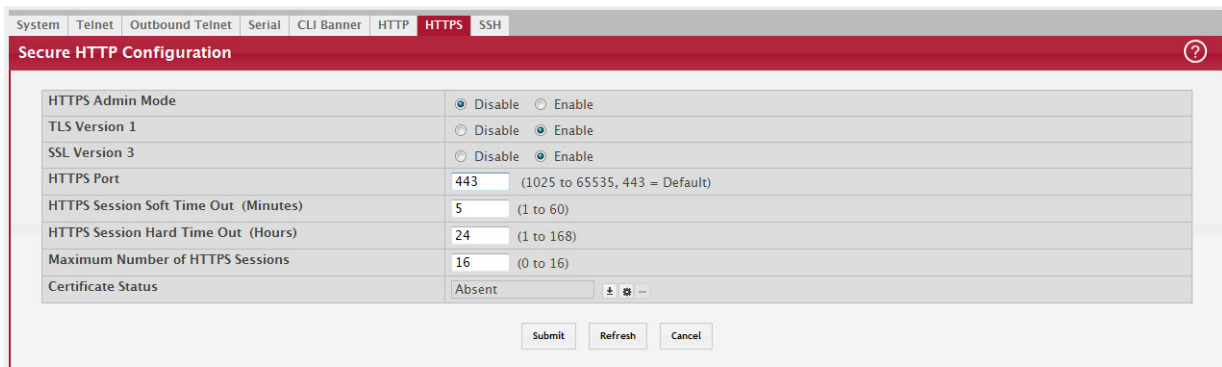


Table 29: HTTPS Configuration Fields

Field	Description
HTTPS Admin Mode	Enables or disables the HTTPS administrative mode. When this mode is enabled, the device can be accessed through a web browser using the HTTPS protocol.
TLS Version 1	Enables or disables Transport Layer Security Version 1.0. When this option is enabled, communication between the web browser on the administrative system and the web server on the device is sent through TLS 1.0.
SSL Version 3	Enables or disables Secure Sockets Layer Version 3.0. When this option is enabled, communication between the web browser on the administrative system and the web server on the device is sent through SSL 3.0. SSL must be administratively disabled while downloading an SSL certificate file from a remote server to the device.
HTTPS Port	The TCP port number that HTTPS uses. <b>Note:</b> Before changing this value, check your system (for example, using <code>netstat</code> ) to make sure that the desired port number is not currently being used by any other service.
HTTPS Session Soft Time Out (Minutes)	HTTPS session inactivity timeout value. A logged-in user that does not exhibit any HTTPS activity for this amount of time is automatically logged out of the HTTPS session.
HTTPS Session Hard Time Out (Hours)	HTTPS session hard timeout value. A user connected to the device via an HTTPS session is automatically logged out after this amount of time regardless of the amount of HTTPS activity that occurs.
Maximum Number of HTTPS Sessions	The maximum number of HTTPS sessions that can be connected to the device simultaneously.
Certificate Status	The status of the SSL certificate generation process. <ul style="list-style-type: none"> <li>• Present – The certificate has been generated and is present on the device</li> <li>• Absent – Certificate is not available on the device</li> <li>• Generation In Progress – An SSL certificate is currently being generated.</li> </ul>
Download Certificates (Button)	Allows you to download an SSL certificate file from a remote system to the device. <b>Note:</b> To download SSL certificate files, SSL must be administratively disabled.
Generate Certificate (Button)	Generates an SSL certificate to use for secure communication between the web browser and the embedded web server on the device.
Delete Certificates (Button)	Deletes the SSL certificate. This button is available only if an SSL certificate is present on the device.



If you make changes to the page, click Submit to apply the changes to the system.

When you click the Download Certificates button, a Download Certificates window appears with the following fields.

**Table 30: Download Certificates Fields**

Field	Description
File Type	Specify the type of file to transfer from the device to a remote system.
Certificate Index	Specify the Certificate Index number from 1 to 8. Enter 0 for None.
Select File	Provides the option to browse to the directory where the file is located, and select the file to transfer to the device.
Status	Provides information about the status of the file transfer.

## 4.7.18 SSH Configuration

Use this page to view and modify the Secure Shell (SSH) server settings on the device. SSH is a network protocol that enables access to the CLI management interface by using an SSH client on a remote administrative system. SSH is a more secure access method than Telnet because it encrypts communication between the administrative system and the device. This page also allows you to download or generate SSH host keys for secure CLI-based management.




To access the SSH Configuration page, click System > Management Access > SSH in the navigation menu.

**Figure 36: SSH Configuration**

**Table 31: SSH Configuration Fields**

Field	Description
SSH Admin Mode	Enables or disables the SSH server administrative mode. When this mode is enabled, the device can be accessed by using an SSH client on a remote system.
SSH Port	The TCP port number on which the SSH server listens for requests. Existing SSH login sessions are not affected by a change in this value, although establishment of any new SSH sessions must use the new port number. <b>Note:</b> Before changing this value, check your system, e.g., using netstat, to make sure the desired port number is not currently being used by any other service.

Table 31: SSH Configuration Fields (Continued)

Field	Description
SSH Version 2	When this option is selected, the SSH server on the device can accept connections from an SSH client using SSH-2 protocol. If the option is clear, the device does not allow connections from clients using the SSH-2 protocol. <b>Note:</b> This is the only supported SSH version and is enabled by default. Clearing this option is not permitted.
SSH Connections Currently in Use	The number of active SSH sessions between remote SSH clients and the SSH server on the device.
Maximum number of SSH Sessions Allowed	The maximum number of SSH sessions that may be connected to the device simultaneously.
SSH Session Timeout (minutes)	The SSH session inactivity timeout value. A connected user that does not exhibit any SSH activity for this amount of time is automatically disconnected from the device.
RSA Key Status	The status of the SSH-1 Rivest-Shamir-Adleman (RSA) key file or SSH-2 RSA key file (PEM Encoded) on the device, which might be Present, Absent, or Generation in Progress. <b>Note:</b> FASTPATH supports only SSHv2. SSH1-RSA keys can be downloaded, but they cannot be used.
DSA Key Status	The status of the SSH-2 Digital Signature Algorithm (DSA) key file (PEM Encoded) on the device, which might be Present, Absent, or Generation in Progress.
ECDSA Key Status	The status of the SSH-2 Elliptic Curve Digital Signature Algorithm (ECDSA) key file (PEM Encoded) on the device, which might be Present, Absent, or Generation in Progress.
ECDSA Key Length	The length of the SSH-2 ECDSA key file used to generate the PEM Encoded key file on the device, if present.
Download Certificates (Button) 	Use this button to download an SSH-1 RSA, SSH-2 RSA, or SSH-2 DSA key file from a remote system to the device. After you click the button, a Download Certificate window opens. Select the file type to download, browse to the location on the remote system, and select the file to upload. Then, click Begin Transfer. The Status field provides information about the file transfer. <b>Note:</b> FASTPATH supports only SSHv2. SSH1-RSA keys can be downloaded, but they cannot be used.
Generate Certificate (Button) 	Use this button to manually generate an RSA key or DSA key on the device.
Delete Certificates (Button) 	Use this button to delete an RSA key or DSA key that has been downloaded to the device or manually generated on the device.

If you make changes to the page, click Submit to apply the changes to the system.

#### 4.7.19 Management Access Control and Administration List

Use this page to create and configure a management access list to help secure access to the switch management features. The Management Access Control and Administration List (MACAL) feature is used to ensure that only known and trusted devices are allowed to remotely manage the switch via TCP/IP.

MACALs can be applied only to in-band ports and cannot be applied to the service port.

To access the Management Access List Configuration page, click System > Management Security > Access Profile in the navigation menu.

Figure 37: Management Access List Configuration

This Management Access List Configuration page provides the capability to add, edit, and remove MACALs.

### NOTICE

Profile rules cannot be added or modified when a profile is active. To add or edit a profile, the Active Profile field must be set to None.

- To add a new MACAL, click Add. The Add Profile Rule dialog box opens. Specify the rule criteria in the available fields.
- To edit an existing rule, select the appropriate check box or click the row to select the account and click Edit. The Edit Profile Rule box opens. Modify the rule criteria as needed.
- To remove a Profile Rule, select one or more table entries and click Remove to delete the selected entries.

Table 32: User Accounts Fields

Field	Description
Access Profile	Profile name for the Management Access Control list. One user defined Access Profile can be created.
Active Profile	Currently enabled profile name.
Packets Filtered	The number of packets filtered due to matching a rule in the MACAL.
Interface	The port/interface or trunk ID.
Management Method	The types of action will be taken on access control list. <ul style="list-style-type: none"> <li>• Permit: To allow conditions for the management access list.</li> <li>• Deny: To deny conditions for the management access list.</li> </ul> In the Add or Edit Profile Rule dialog, this is specified by using the Action field.
Source IP Address	IP Address of device which needs to permit or deny in the management access list.
Subnet Mask	Specifies the network mask of the source IP address.
VLAN	The VLAN ID.
Port Channel	Port channels, also known as Link Aggregation Groups (LAGs), allow one or more full-duplex Ethernet links of the same speed to be aggregated together.

Table 32: User Accounts Fields (Continued)

Field	Description
Service	The type of service to permit or deny: <ul style="list-style-type: none"> <li>• ANY</li> <li>• TELNET</li> <li>• HTTP</li> <li>• HTTPS</li> <li>• SNMP</li> <li>• SSH</li> <li>• TFTP</li> <li>• SNTF</li> </ul>
Priority	Priority for the rule. Duplicates are not allowed.

### 4.7.20 User Accounts

By default, the switch contains two user accounts:

- admin, with 'Read/Write' privileges
- guest, with 'Read Only' privileges

Both of these accounts have blank passwords by default. The names are not case sensitive.

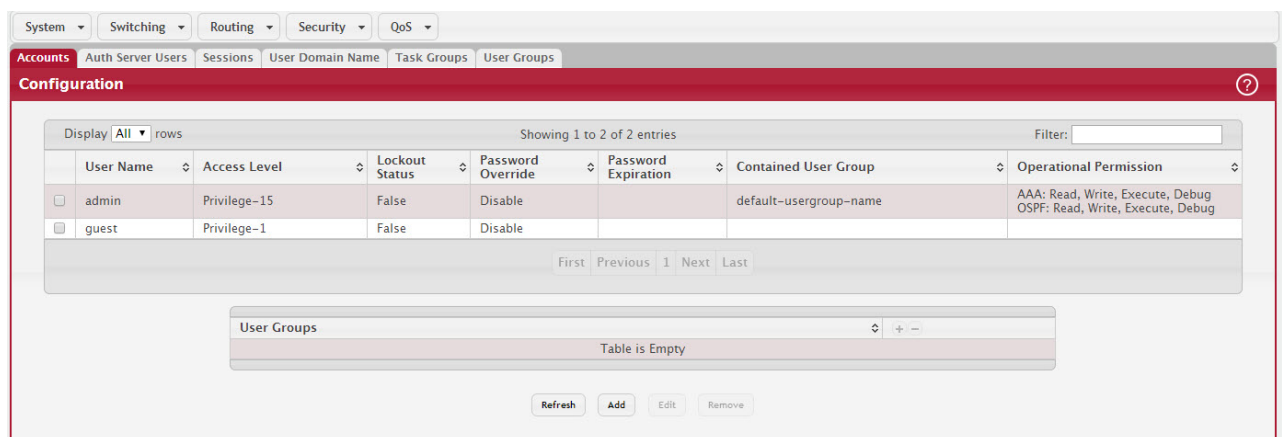
If you log on to the switch with the user account that Read/Write privileges (i.e., as admin), you can use the User Accounts page to assign passwords and set security parameters for the default accounts. You can also add up to five read-only accounts. You can delete all accounts except for the Read/Write account.

**NOTICE**

Only a user with Read/Write privileges may alter data on this screen, and only one account can exist with Read/Write privileges.

To access the User Accounts page, click System > Users > Accounts in the navigation menu.

Figure 38: User Accounts



This User Accounts page provides the capability to add, edit, and remove user accounts.

- To add a user, click Add. The Add new user dialog box opens. Specify the new account information in the available fields.
- To edit an existing user, select the appropriate check box or click the row to select the account and click Edit. The Edit existing user dialog box opens. Modify the account information as needed.
- To remove a user, select one or more table entries and click Remove to delete the selected entries.

Table 33: User Accounts Fields

Field	Description
User Name	Enter the name you want to give to the new account. (You can only enter data in this field when you are creating a new account.) User names are up to 32 alphanumeric characters in length and are not case sensitive. Valid characters include all the alphanumeric characters and the dash ('-') and underscore ('_') characters. User name default is not valid. <b>Note:</b> You can change the Read/Write user name from <i>admin</i> to something else, but when you click <b>Submit</b> , you must re-authenticate with the new username.
Password	Enter the optional new or changed password for the account. It will not display as it is typed, only asterisks (*) or dots(.) will show based on the browser used. Passwords must be greater than eight characters and can be up to 64 characters in length, and are case sensitive.
Confirm	Enter the password again, to confirm that you entered it correctly. This field will not display, but will show asterisks (*)
Access Level	Indicates the access or privilege level for this user. The options are: <ul style="list-style-type: none"> <li>• Read Write - The user can view and modify the configuration.</li> <li>• Read Only - The user can view the configuration but cannot modify any fields.</li> <li>• Suspended - The user exists but is not permitted to log on to the device.</li> </ul>
Password Override	Identifies the password override complexity status for this user. <ul style="list-style-type: none"> <li>• Enable - The system does not check the strength of the password.</li> <li>• Disable - When configuring a password, it is checked against the Strength Check rules configured for passwords.</li> </ul>
Password Expiration	Indicates the date when this user's current password will expire. This is determined by the date the password was created and the number of days specified in the aging Password Aging setting on the Password Management page.
Contained User Group	The associated user groups for the user.
Operational Permissions	The operational task permissions for the user. In addition to the fields described above, the User Groups table will be populated when you click on each row. To configure this user group, click the Add icon in the header row. To remove the user group, click the Reset icon in the row.
Password Strength	Shows the status of password strength check.
Encrypt password	Select this option to encrypt the password before it is stored on the device.

#### 4.7.20.1 Adding a User Account

Use the following procedures to add a user account. The system supports one Read/Write user and five Read Only users.

1. From the User menu, click Add. The Add new user dialog is displayed.

Figure 39: Add new user

**Add new user**

Unencrypted passwords must be from 8 to 64 characters in length.  
Encrypted passwords must be 128 hexadecimal characters for AES and 34 characters for MD5 salt hash.

User Name	<input type="text"/> (1 to 64)
Password	<input type="text"/>
Confirm	<input type="text"/>
Access Level	<input type="radio"/> Privilege-0 <input checked="" type="radio"/> Privilege-1 <input type="radio"/> Privilege-15
Password Override	<input type="checkbox"/>
Password Strength	Disabled
Encryption Type	<input checked="" type="radio"/> AES <input type="radio"/> MD5-Salt
Encrypted Password	<input type="checkbox"/>

**Submit** **Cancel**

2. Enter a user name and password for the new user, then reenter the password in the Confirm Password field.
3. Select the Access Level option.
4. Select the Password Override option to enable or disable the password override complexity status for this user.
5. Select the option to specify the encryption type. Options are AES or MD5-Salt. The default encryption type is AES.
6. Select the encrypted password option to encrypt the password before it is stored on the device.
7. Click Submit to update the switch with the values on this screen.

If you want the switch to retain the new values across a power cycle, you must perform a save.

#### 4.7.20.2 Changing User Account Information

You cannot add or delete the Read/Write user, but you can change the user name and password. To change the password for an existing account or to overwrite the user name on an existing account, use the following procedures.

1. From the User menu, select the user to change. The screen is refreshed.
2. Click Edit to change the user settings.
3. This button is only visible when you have selected a user account with 'Read Only' access. You cannot delete the 'Read/Write' user.
4. The Edit existing user dialog is displayed.

Figure 40: Edit existing user

**Unencrypted passwords must be from 8 to 64 characters in length.  
Encrypted passwords must be 128 hexadecimal characters for AES and 34 characters for MD5 salt hash.**

User Name	u2
Password	<input type="text"/>
Confirm	<input type="text"/>
Access Level	<input type="radio"/> Privilege-0 <input type="radio"/> Privilege-1 <input checked="" type="radio"/> Privilege-15
Lockout Status	False
Unlock User Account	<input type="checkbox"/>
Password Override	<input type="checkbox"/>
Password Strength	Disabled
Encryption Type	<input checked="" type="radio"/> AES <input type="radio"/> MD5-Salt
Encrypted Password	<input type="checkbox"/>

1. To alter the user name or, delete the existing name in the User Name field and enter the new user name.
2. To change the password, delete any asterisks (\*) in the Password and Confirm fields, and then enter and confirm the new password. Unencrypted passwords must be from 8 to 64 characters in length and are case sensitive. Encrypted passwords must be exactly 128 hexadecimal characters for AES and 34 characters for MD5-Salt hash, if specified in encrypted format. Be sure the password conforms to the allowed number of characters.
3. Select the Access Level option.
4. Select the Password Override option to enable or disable the password override complexity status for this user.
5. Select the option to specify the encryption type. Options are AES or MD5-Salt. The default encryption type is AES.
6. Select the encrypted password option to encrypt the password before it is stored on the device.
7. Click Submit to update the switch with the values on this screen.

If you want the switch to retain the new values across a power cycle, you must perform a save.

### 4.7.20.3 Removing a User Account

Use the following procedures to remove any of the Read Only user accounts.

1. From the User menu, select the user to remove.

The screen refreshes.

2. Click Remove to delete the user.

This button is only visible when you have selected a user account with 'Read Only' access. You cannot remove the 'Read/Write' user.

If you want the switch to retain the new values across a power cycle, you must perform a save.

## 4.7.21 Authentication Server Users

Use the Auth Server Users page to add and remove users from the local authentication server user database. For some security features, such as IEEE 802.1X port-based authentication, you can configure the device to use the locally stored list of usernames and passwords to provide authentication to users instead of using an external authentication server.

---

### NOTICE

The preconfigured users, admin and guest, are assigned to a pre-configured list named defaultList, which you cannot delete. All newly created users are also assigned to the defaultList until you specifically assign them to a different list.

---

You can create a text file that contains a list of IAS users to add to the database and then download the file to the switch. The following script is an example of an IAS user text file that contains three users:

```
configure
aaa ias-user username client-1
password my-password1
exit
aaa ias-user username client-2
password aa5c6c251fe374d5e306c62496c3bcf6 encrypted
exit
aaa ias-user username client-3
password 1f3ccb1157
exit
```

After the download completes, client-1, client-2, and client-3 are added to the IAS database. The password for client-2 is encrypted.

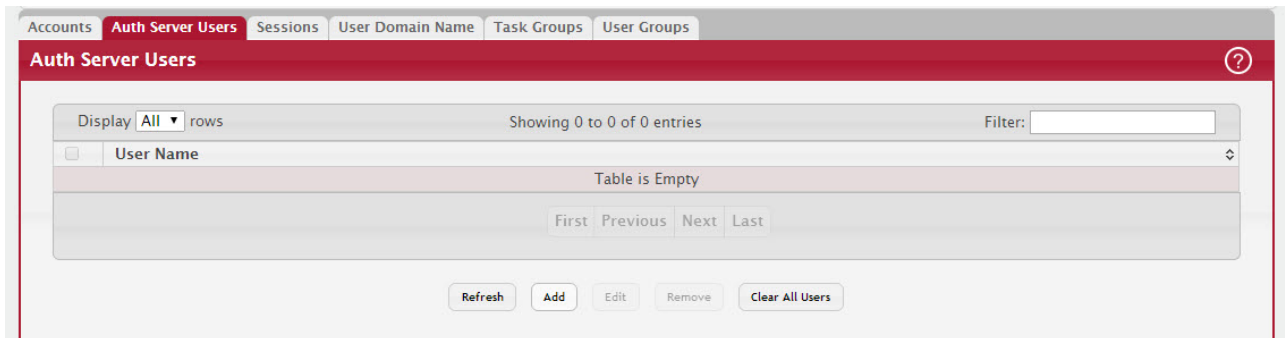
When Dot1x authentication is enabled on the ports and the authentication method is LOCAL, port access is allowed only to users in this database that provide the correct name and password.

Use the buttons to perform the following tasks:

- To add a new authentication server user, click Add.
- To add a user to the local authentication server database, click Add and complete the required information.
- To change the password information for an existing user, select the user to update and click Edit.
- To delete a user from the database, select each user to delete and click Remove.
- To remove all users from the database, click Clear All Users.



Figure 41: Auth Server Users



When Add is selected from Auth Server Users list, the Add New User page displays.

Figure 42: Add New User

Table 34: Add New Authentication User Fields

Field	Description
User Name	A unique name used to identify this user account. You configure the User Name when you add a new user.
Password Required	Select this option to indicate that the user must enter a password to be authenticated. If this option is clear, the user is required only to enter a valid user name.
Password	Specify the password to associate with the user name (if required).
Confirm	Re-enter the password to confirm the entry.
Encrypted	Select this option to encrypt the password before it is stored on the device.

Click Submit to apply the changes to the system. You must perform a save to make the changes persist across a reboot.

#### 4.7.21.1 Logged-in Sessions

The Sessions page identifies the users that are logged in to the management interface of the device. The page also provides information about their connections.

To access the Logged In Session page, click System > Users > Sessions in the navigation menu.

Figure 43: Logged In Sessions

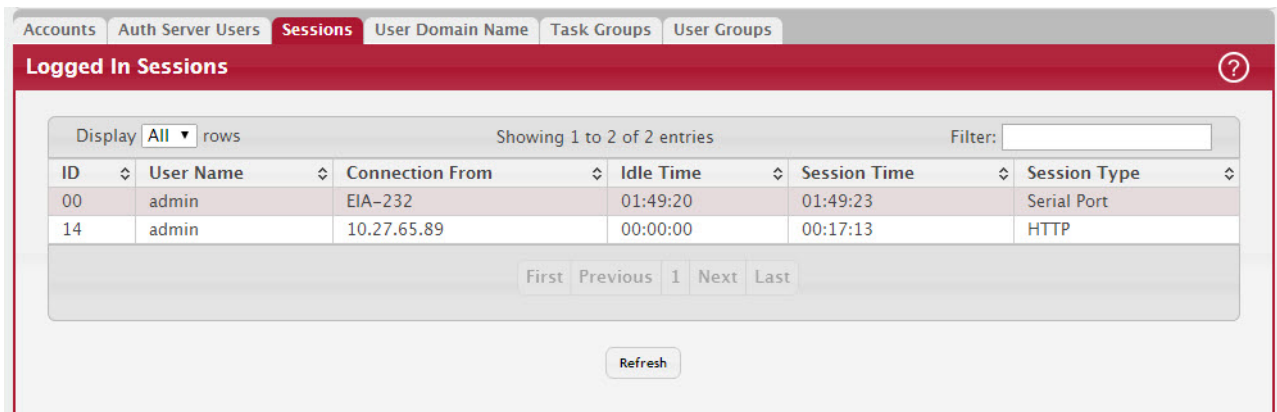


Table 35: Logged In Sessions Fields

Field	Description
ID	The unique ID of the session.
User Name	The name that identifies the user account.
Connection From	Identifies the administrative system that is the source of the connection. For remote connections, this field shows the IP address of the administrative system. For local connections through the console port, this field shows the communication standard for the serial connection.
Idle Time	Shows the amount of time in hours, minutes, and seconds that the logged-on user has been inactive.
Session Time	Shows the amount of time in hours, minutes, and seconds since the user logged onto the system.
Session Type	Shows the type of session, which can be Telnet, Serial, SSH, HTTP, or HTTPS.


Click Refresh to update the information on the screen.

### 4.7.22 User Domain Name

Use this page to configure the domain name to send to the authentication server, along with the user name and password, to authenticate a user attempting to access the device management interface. Domain name authentication is supported when user authentication is performed by a Remote Authentication Dial-In User Server (RADIUS) server or a TACACS+ server.

To access the User Domain Name page, click System > Users > User Domain Name in the navigation menu.

Figure 44: User Domain Name



**Table 36: User Domain Name Fields**

Field	Description
User Domain Name Mode	The administrative mode of domain name authentication on the device. When enabled, the domain name is included when the user name and password are sent to the authentication server. The domain name can be input by the user in the User Name field on the login screen in a domain-name\username format, or the domain name can be specified in the Domain Name field.
Domain Name	The domain name to send to the authentication server when the user does not provide one in the User Name field during logon. When only the username is provided, the device sends the username as domain-name\username, where domain-name is the string configured in this field. To configure the domain name, click the Edit icon and specify the desired string. To reset the field to its default value, click the Reset icon and confirm the action.

### 4.7.22.1 Task Group

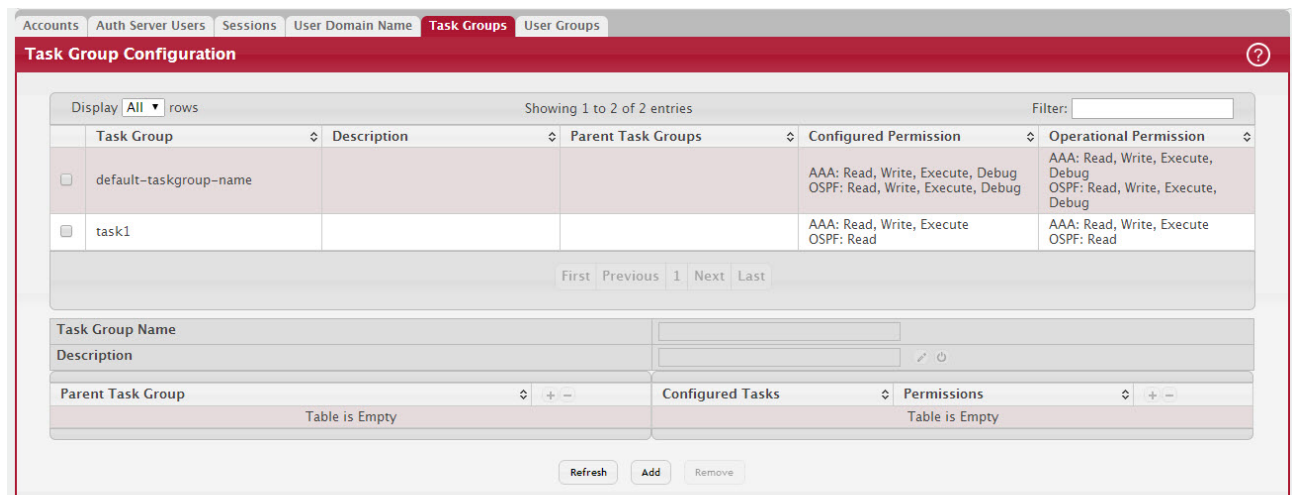
The Task Group Configuration page allows you to add, edit, and remove task groups. Task groups allow users to have different permission levels (read, write, execute, debug) at a per-component level. Task-based authorization uses the concept of components/tasks to define permission for commands for a given user. Users are assigned to User Groups that are, in turn, associated with Task Groups. Each Task Group is then associated with one or more tasks/components.

The user admin, and any other user with privilege level 15, are part of the default user-group or task-group and have read, write, execute, and debug access to commands from all components.

This feature is supported only for users who are authenticated locally via the Web interface.

To access the Task Group page, click System > Users > Task Group in the navigation menu.

**Figure 45: Task Group Configuration**



**Table 37: Task Group Configuration Fields**

Field	Description
Task Group	The task group name.
Description	The associated description for task group name.
Parent Task Groups	The associated parent task groups for task group name. To configure this parent task group, click the Add icon in the header row. To remove the parent task group, click the Reset icon in the row.
Configured Permission	The configured task permissions for task group.

**Table 37: Task Group Configuration Fields (Continued)**

Field	Description
Configured Tasks	The list of task names. To configure this task, click the Add icon in the header row. To remove the task, click the Reset icon in the row. The tasks available are platform and package dependent.
Permissions	The task permissions. <ul style="list-style-type: none"> <li>• Read</li> <li>• Write</li> <li>• Debug</li> <li>• Execute</li> </ul>

Use the buttons to perform the following:

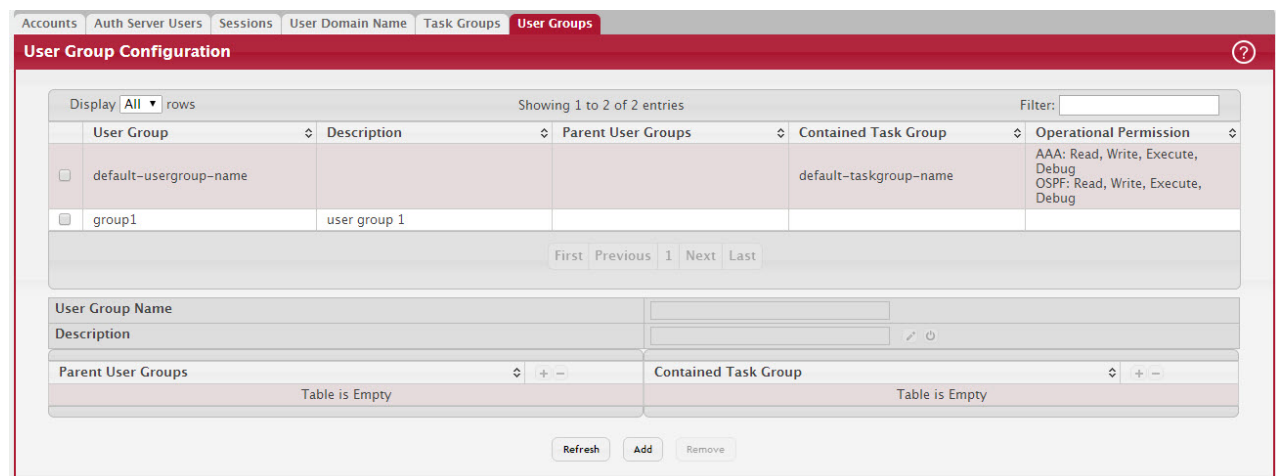
- To add a task group, click Add and specify a name for the group.
- To remove a task group, select the check box associate to the group to remove and click Remove.
- Click Refresh to update the information on the screen.

### 4.7.22.2 User Group

The User Group Configuration page allows you to add, edit, and remove user groups.

To access the User Group page, click System > Users > User Group in the navigation menu.

**Figure 46: User Group Configuration**



**Table 38: User Group Configuration Fields**

Field	Description
User Group	The user group name.
Description	The associated description for User group name.
Parent User Groups	The associated parent user groups for user group. To configure this parent user group, click the Add icon in the header row. To remove the parent user group, click the Reset icon in the row.
Configured Permission	The configured task permissions for task group.

**Table 38: User Group Configuration Fields (Continued)**

Field	Description
Contained Task Group	The associated task groups for user group. To configure this task group, click the Add icon in the header row. To remove the task group, click the Reset icon in the row.
Operational Permission	The operational task permissions for the user group. <ul style="list-style-type: none"> <li>• Read</li> <li>• Write</li> <li>• Debug</li> <li>• Execute</li> </ul>

Use the buttons to perform the following:

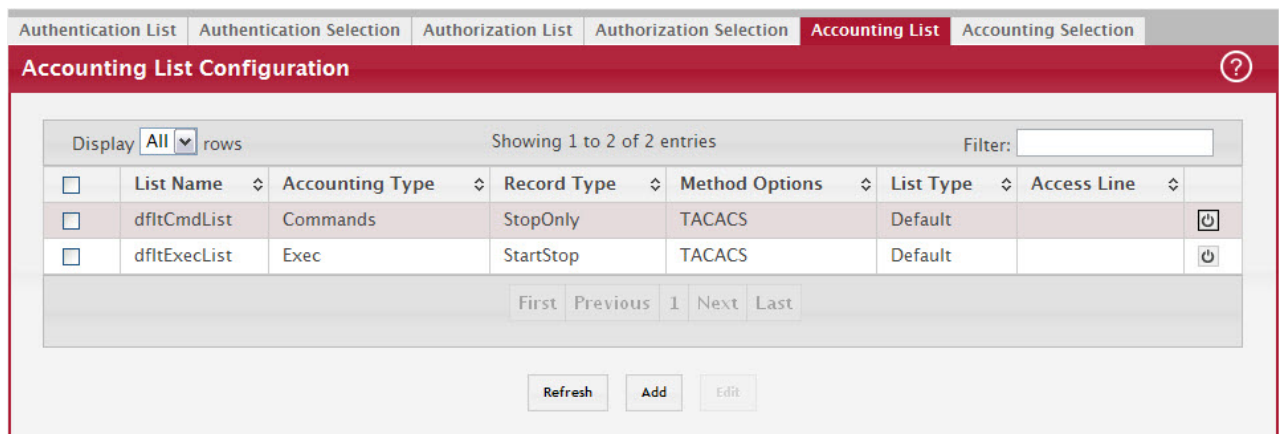
- To add a user group, click Add and specify a name for the group.
- To remove a user group, select the check box associate to the group to remove and click Remove.
- Click Refresh to update the information on the screen.

### 4.7.23 Accounting List Configuration

Use the Accounting List Configuration page to view and configure the accounting lists for users who access the command-line interface (CLI) to manage and monitor the device. Accounting lists are used to record user activity on the device. The device is preconfigured with accounting lists. These are default lists, and they cannot be deleted. Additionally, the List Name and Accounting Type settings for the default lists cannot be changed.

To access the Authentication List Configuration page, click System > AAA > Accounting List in the navigation menu.

**Figure 47: Accounting List Configuration**



Use the buttons to perform the following tasks:

- To configure a new accounting list, click Add.
- To edit a list, select the entry to modify and click Edit. The settings that can be edited depend on the list type.
- To remove a non-default accounting list, click the – (minus) button associated with the entry. You must confirm the action before the entry is deleted.
- To reset the Method Options for a default accounting list to the factory default values, click the Reset icon associated with the entry. You must confirm the action before the entry is reset.

Figure 48: Add New Accounting List

Figure 48: "Add New Accounting List," on page 72 shows the fields on the Add New Accounting List page.

After you click Add or Edit, a window opens and allows you to configure accounting list settings. When adding an accounting list, you can configure the List Name, Accounting Type, and Record Type fields as well as the Accounting Methods. When editing an existing authentication list, only the Record Type and Accounting Methods can be configured. The following information describes how to set the Accounting Methods.

Table 39: Add New Accounting List Fields

Field	Description
Accounting Methods	This area includes the Available Methods and Selected Methods fields. If a list uses multiple accounting methods, the order in which you move the method from the Available Methods field to the Selected Methods field determines the order in which the device attempts to send accounting notifications. If the device successfully sends the accounting notifications by using the first method, the next method is not attempted.
Available Methods	The accounting methods that can be used for the accounting list. To set the accounting method, select the method in the Available Methods field and click the right arrow to move it into the Selected Methods field.
Selected Methods	The accounting methods currently configured for the list. When multiple methods are in this field, the order in which the methods are listed is the order in which the methods will be used. If the device is unable to send accounting notifications by using the first method, the device attempts to send notifications by using the second method. To remove a method from this field, select it and click the left arrow to return it to the Available Methods area.

Table 40, "Accounting List Configuration Fields," on page 73 describes the fields on the Accounting List Configuration page.

**Table 40: Accounting List Configuration Fields**

Field	Description
Accounting Type	The type of accounting list, which is one of the following: Command – Each CLI command executed by the user, along with the time the command was executed, is recorded and sent to an external AAA server. EXEC – User login and logout times are recorded and sent to an external AAA server.
List Name	The name of the accounting list. This field can be configured only when adding a new accounting list.
Record Type	Indicates when to record and send information about the user activity: <ul style="list-style-type: none"> <li>• StartStop – Accounting notifications are sent at the beginning and at the end of an exec session or a user-executed command. User activity does not wait for the accounting notification to be recorded at the AAA server.</li> <li>• StopOnly – Accounting notifications are sent at the end of an exec session or a user-executed command.</li> </ul>
Method Options	The method(s) used to record user activity. The possible methods are as follows: <ul style="list-style-type: none"> <li>• TACACS+ – Accounting notifications are sent to the configured TACACS+ server.</li> <li>• RADIUS – Accounting notifications are sent to the configured RADIUS server.</li> </ul>
List Type	The type of accounting list, which is one of the following: <ul style="list-style-type: none"> <li>• Default – The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options and Record Type settings are configurable.</li> <li>• Configured – The list has been added by a user.</li> </ul>
Access Line	The access method(s) that use the list for accounting user activity. The settings for this field are configured on the Accounting Selection page.

If you change any of the parameters, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

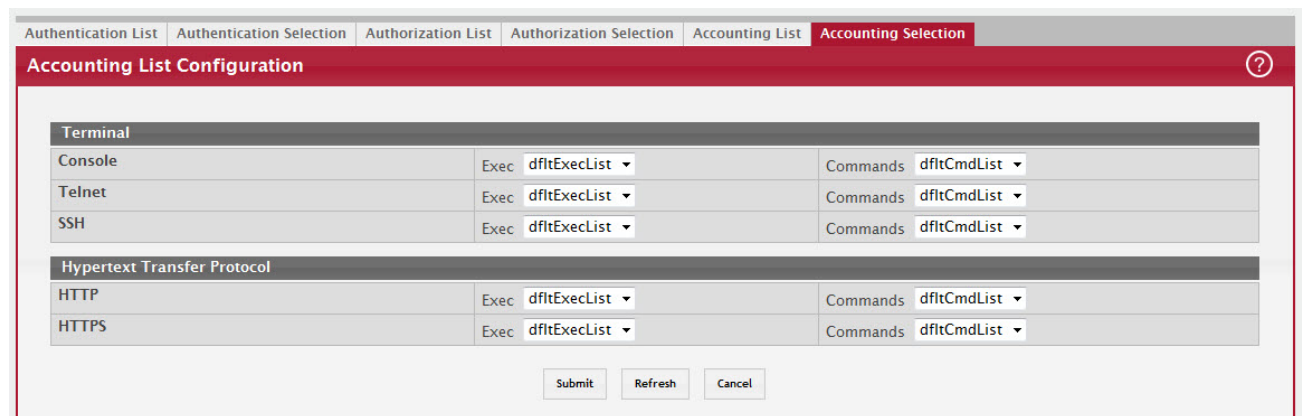
### 4.7.24 Accounting List Configuration

Use this page to associate an accounting list with each access method. For each access method, the following two accounting lists are associated:

- Exec – The accounting list to record user login and logout times.
- Commands – The accounting list to record which actions a user takes on the system, such as page views or configuration changes. This list also records the time when the action occurred. For Terminal access methods, this list records the CLI commands a user executes and when each command is issued.

To access the Accounting List Configuration page, click System > AAA > Accounting Selection in the navigation menu.

**Figure 49: Accounting List Configuration**



**Table 41: Accounting List Configuration Fields**

Field	Description
Terminal	<p>The access methods in this section are CLI-based.</p> <ul style="list-style-type: none"> <li>• Console—The Exec accounting list and the Commands accounting list to apply to users who access the CLI by using a connection to the console port.</li> <li>• Telnet — The Exec accounting list and the Commands accounting list to apply to users who access the CLI by using a Telnet session.</li> <li>• SSH —The Exec accounting list and the Commands accounting list to apply to users who access the CLI by using a secure shell (SSH) session.</li> </ul>
Hypertext Transfer Protocol	<p>The access methods in this section are through a web browser.</p> <ul style="list-style-type: none"> <li>• HTTP —The Exec accounting list and the Commands accounting list to apply to users who access the web-based management interface by using HTTP.</li> <li>• Telnet —The Exec accounting list and the Commands accounting list to apply to users who access the web-based management interface by using secure HTTP.</li> </ul>

If you change any of the parameters, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

### 4.7.25 Authentication List Summary

Use the Authentication List Summary page to view and configure the authentication lists used for management access and port-based (IEEE 802.1X) access to the system. An authentication list specifies which authentication method(s) to use to validate the credentials of a user who attempts to access the device. Several authentication lists are preconfigured on the system. These are default lists, and they cannot be deleted. Additionally, the List Name and Access Type settings for the default lists cannot be changed.

To access the Authentication List Summary page, click System > AAA > Authentication List in the navigation menu.

**Figure 50: Authentication List Configuration Summary**

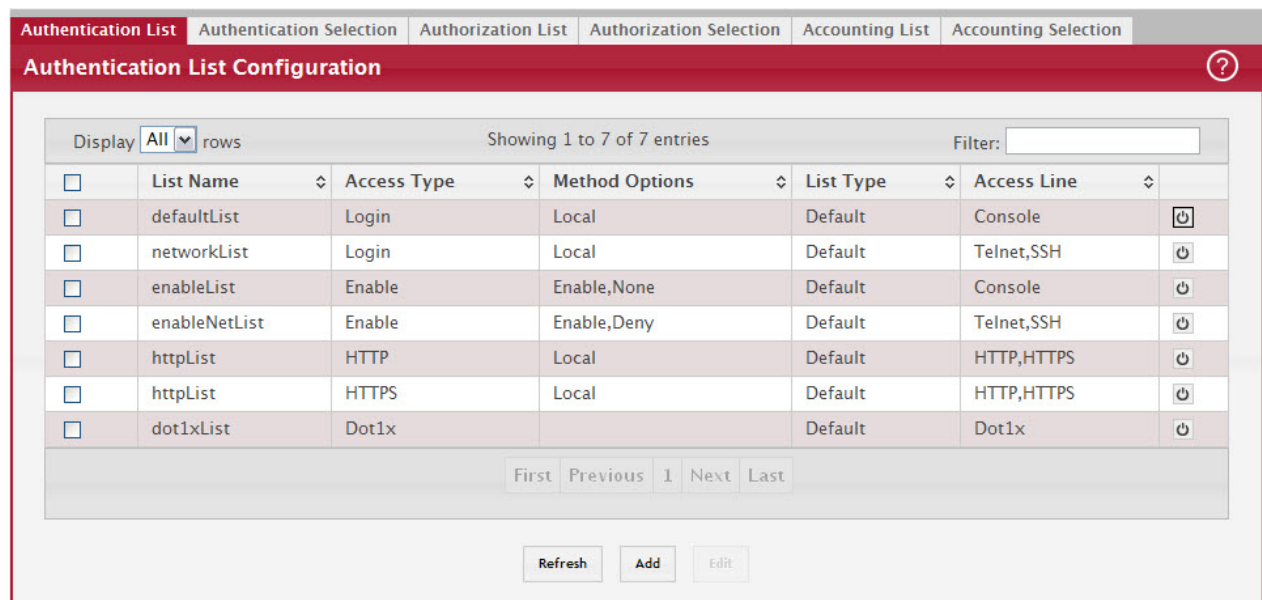


Table 42, "Authentication List Configuration Summary Fields," on page 75 describes the fields for the Authentication List Summary page.



Table 42: Authentication List Configuration Summary Fields

Field	Description
List Name	The name of the authentication list. This field can be configured only when adding a new authentication list.
Access Type	The way the user accesses the system. This field can be configured only when adding a new authentication list, and only the Login and Enable access types can be selected. The access types are as follows: <ul style="list-style-type: none"> <li>• Login – User EXEC-level management access to the command-line interface (CLI) by using a console connection or a telnet or SSH session. Access at this level has a limited number of CLI commands available to view or configure the system.</li> <li>• Enable – Privileged EXEC-level management access to the CLI by using a console connection or a telnet or SSH session. In Privileged EXEC mode, read-write users have access to all CLI commands.</li> <li>• HTTP – Management-level access to the web-based user interface by using HTTP.</li> <li>• HTTPS – Management-level access to the web-based user interface by using secure HTTP.</li> <li>• Dot1x – Port-based access to the network through a switch port that is controlled by IEEE 802.1X.</li> </ul>
Method Options	The method(s) used to authenticate a user who attempts to access the management interface or network. The possible methods are as follows: <ul style="list-style-type: none"> <li>• Enable – Uses the locally configured Enable password to verify the user's credentials.</li> <li>• Line – Uses the locally configured Line password to verify the user's credentials.</li> <li>• Local – Uses the ID and password in the Local User database to verify the user's credentials.</li> <li>• RADIUS – Sends the user's ID and password to the configured RADIUS server to verify the user's credentials.</li> <li>• TACACS+ – Sends the user's ID and password to the configured TACACS+ server to verify the user's credentials.</li> <li>• None – No authentication is used.</li> <li>• IAS – Uses the local Internal Authentication Server (IAS) database for 802.1X port-based authentication.</li> </ul>
List Type	The type of list, which is one of the following: <ul style="list-style-type: none"> <li>• Default – The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options are configurable.</li> <li>• Configured – The list has been added by a user.</li> </ul>
Access Line	The access method(s) that use the list for authentication. The settings for this field are configured on the Authentication Selection page.

- Click Refresh to update the information on the screen.
- To create a new authentication list, see [Section 4.7.21: "Authentication Server Users"](#). To assign users to a specific authentication list, see [Section 4.7.20: "User Accounts"](#). To configure the 802.1x port security users, see [Section 7.3: "RADIUS Settings"](#).

#### 4.7.26 Select Authentication List

Use the Select Authentication List Configuration page to associate an authentication list with each CLI-based access method (Console, Telnet, and SSH). Each access method has the following two authentication lists associated with it:

- Login – The authentication list to use for User EXEC-level management access to the CLI. Access at this level has a limited number of CLI commands available to view or configure the system. The options available in this menu include the default Login authentication lists as well as any user-configured Login lists.
- Enable – The authentication list to use for Privileged EXEC-level management access to the CLI. In Privileged EXEC mode, read-write users have access to all CLI commands. The options available in this menu include the default Enable authentication lists as well as any user-configured Enable lists.

To access the Select Authentication List page, click System > AAA > Authentication Selection in the navigation menu.

Figure 51: Select Authentication List

Terminal				
Console	Login	defaultList	Enable	enableList
Telnet	Login	networkList	Enable	enableNetList
SSH	Login	networkList	Enable	enableNetList

Table 43, "Select Authentication List Fields," on page 76 describes the fields for the Select Authentication List page.

**Table 43: Select Authentication List Fields**

Field	Description
Console	The Login authentication list and the Enable authentication list to apply to users who attempt to access the CLI by using a connection to the console port.
Telnet	The Login authentication list and the Enable authentication list to apply to users who attempt to access the CLI by using a Telnet session.
Secure Telnet (SSH)	The Login authentication list and the Enable authentication list to apply to users who attempt to access the CLI by using a secure shell (SSH) session.
List Name	The name of the authentication list. This field can be configured only when adding a new authentication list.
Access Type	<p>The way the user accesses the system. This field can be configured only when adding a new authentication list, and only the Login and Enable access types can be selected. The access types are as follows:</p> <ul style="list-style-type: none"> <li>• Login – User EXEC-level management access to the command-line interface (CLI) by using a console connection or a telnet or SSH session. Access at this level has a limited number of CLI commands available to view or configure the system.</li> <li>• Enable – Privileged EXEC-level management access to the CLI by using a console connection or a telnet or SSH session. In Privileged EXEC mode, read-write users have access to all CLI commands.</li> <li>• HTTP – Management-level access to the web-based user interface by using HTTP.</li> <li>• HTTPS – Management-level access to the web-based user interface by using secure HTTP.</li> <li>• Dot1x – Port-based access to the network through a switch port that is controlled by IEEE 802.1X.</li> </ul>
Method Options	<p>The method(s) used to authenticate a user who attempts to access the management interface or network. The possible methods are as follows:</p> <ul style="list-style-type: none"> <li>• Enable – Uses the locally configured Enable password to verify the user's credentials.</li> <li>• Line – Uses the locally configured Line password to verify the user's credentials.</li> <li>• Local – Uses the ID and password in the Local User database to verify the user's credentials.</li> <li>• RADIUS – Sends the user's ID and password to the configured RADIUS server to verify the user's credentials.</li> <li>• TACACS+ – Sends the user's ID and password to the configured TACACS+ server to verify the user's credentials.</li> <li>• None – No authentication is used.</li> <li>• IAS – Uses the local Internal Authentication Server (IAS) database for 802.1X port-based authentication.</li> </ul>

Table 43: Select Authentication List Fields (Continued)

Field	Description
List Type	The type of list, which is one of the following: <ul style="list-style-type: none"> <li>Default – The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options are configurable.</li> <li>Configured – The list has been added by a user.</li> </ul>
Access Line	The access method(s) that use the list for authentication. The settings for this field are configured on the Authentication Selection page.

### Command Button

The page has the following command button:

- Submit–Update the switch with the values on the screen. If you want the switch to retain the new values across a power cycle you must perform a save.

## 4.7.27 Authorization List Configuration

Use this page to view and configure the authorization lists for users who access the command-line interface (CLI) and for users who access the network through IEEE 802.1X-enabled ports. Authorization lists are used to determine whether a user is permitted to perform a given activity on the system or network. Several authorization lists are preconfigured on the system. These are default lists, and they cannot be deleted. Additionally, the List Name and Authorization Type settings for the default lists cannot be changed.

To access the Authorization List Configuration page, click System > AAA > Authorization List in the navigation menu.

Figure 52: Authorization List Configuration

List Name	Authorization Type	Method Options	List Type	Access Line
dfltCmdAuthList	Commands	None	Default	Console,Telnet,SSH
dfltExecAuthList	Exec	None	Default	Console,Telnet,SSH
networkList	Network		Default	Dot1x

Table 44: Authorization List Configuration Fields

Field	Description
List Name	The name of the authorization list. This field can be configured only when adding a new authorization list.
Authorization Type	The type of authorization list, which is one of the following: <ul style="list-style-type: none"> <li>Command – Determines which CLI commands a user is permitted to issue. When command authorization is enabled, each command a user enters must be validated before the command is executed.</li> <li>EXEC – Determines whether a user can bypass User EXEC mode and enter Privileged EXEC mode directly after a successful Login authentication.</li> <li>Network – Determines whether the user is permitted to access various network services. This authorization type applies to port-based access (IEEE 802.1X) rather than access to the CLI.</li> </ul>

**Table 44: Authorization List Configuration Fields (Continued)**

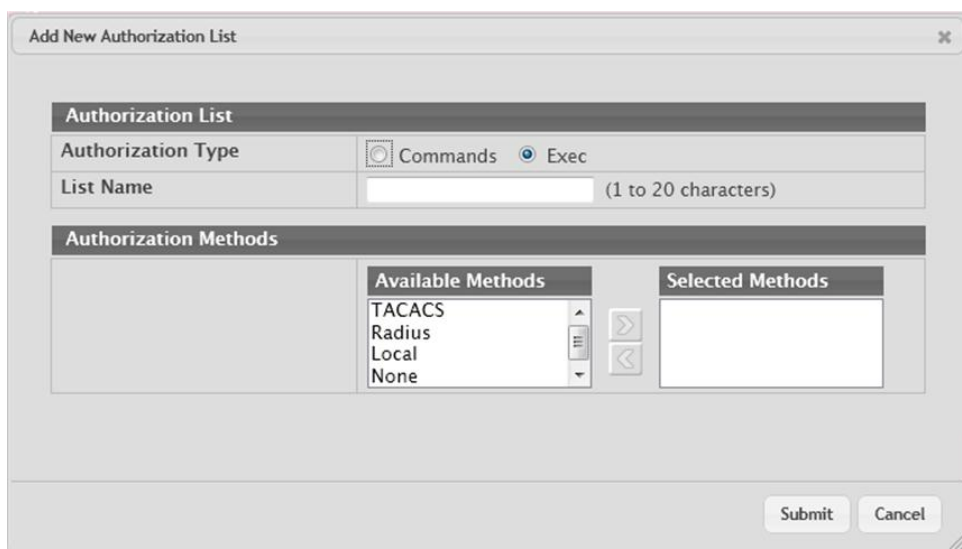
Field	Description
Method Options	<p>The method(s) used to authorize a user's access to the device or network services. The possible methods are as follows:</p> <ul style="list-style-type: none"> <li>• TACACS+ – When a user issues a CLI command, the device contacts the configured TACACS+ server to verify whether the user is allowed to issue the command. If approved, the command is executed. Otherwise, the command fails.</li> <li>• RADIUS – When a user is authenticated by the RADIUS server, the device downloads a list of permitted/denied commands from the RADIUS server. The list of authorized commands that are associated with the authenticated user is cached during the user's session. If this method is selected, the authentication method for the access type must also be RADIUS.</li> <li>• Local – Uses a list stored locally on the system to determine whether the user is authorized to access the given services.</li> <li>• None – No authorization is used. If the method is None, the authorization type is effectively disabled.</li> </ul>
List Type	<p>The type of authorization list, which is one of the following:</p> <ul style="list-style-type: none"> <li>• Default – The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options are configurable.</li> <li>• Configured – The list has been added by a user.</li> </ul>
Access Line	<p>The access method(s) that use the list for authorization. The settings for this field are configured on the Authorization Selection page.</p>

- To configure a new authorization list, click Add.
- To edit a list, select the entry to modify and click Edit. The settings that can be edited depend on the list type.
- To remove a non-default authorization list, click the – (minus) button associated with the entry. You must confirm the action before the entry is deleted.

To reset the Method Options for a default authorization list to the factory default values, click the Reset icon associated with the entry. You must confirm the action before the entry is reset.

After you click Add or Edit, a window opens and allows you to configure authorization list settings.

**Figure 53: Add New Authorization List**



When adding an authorization list, you can configure the List Name and Authorization Type fields as well as the Authorization Methods. When editing an existing authentication list, only the Authorization Methods can be configured. The following information describes how to set the Authorization Methods.

Table 45: Add New Authorization List Fields

Field	Description
Authorization Methods	This area includes the Available Methods and Selected Methods fields. For lists that allow multiple authorization methods, the order in which you move the method from the Available Methods field to the Selected Methods field determines the order in which the device attempts to authorize the user.
Available Methods	The authorization methods that can be used for the authorization list. Not all methods are available for all lists. To set the authorization method, select the method in the Available Methods field and click the right arrow to move it into the Selected Methods field.
Selected Methods	The authorization methods currently configured for the list. When multiple methods are in this field, the order in which the methods are listed is the order in which the methods will be used to authorize a user. If the user fails to be authorized using the first method, the device attempts to authorize the user by using the next method in the list. No authorization methods can be added after None. To remove a method from this field, select it and click the left arrow to return it to the Available Methods area.

### 4.7.28 Line Password

Use the Line Password page to configure line mode passwords. The administrator can input line mode passwords in encrypted format and specify the encryption type, which helps transfer passwords between devices without having to know the passwords. A password entered or copied from another switch configuration is already encrypted. Line Password Configuration

Table 46: Line Password Configuration Fields

Field	Description
Line Mode	Any or all of the following passwords may be changed on this page by checking the box that precedes it: <ul style="list-style-type: none"> <li>• Console</li> <li>• Telnet</li> <li>• SSH</li> </ul>
Password (8–64 characters)	Enter the new password for the corresponding Line Mode in this field. Be sure the password conforms to the allowed number of characters. The password characters are not displayed on the page, but are disguised in a browser-specific manner.

Table 46: Line Password Configuration Fields (Continued)

Field	Description
Confirm Password (8–64 characters)	Re-enter the new password for the corresponding Line Mode in this field. This must be the same value entered in the Password field. Be sure the password conforms to the allowed number of characters. The password characters are not displayed on the page, but are disguised in a browser-specific manner.
Encryption Type	Select the option to specify encryption type. Options are AES or MD5-Salt. The default encryption type is AES.
Encrypted Password	Select the option to indicate whether the password is specified in encrypted format.

If you change any data, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

#### 4.7.29 Enable Password

Use the Enable Password page to configure the enable password. The administrator can input enable mode passwords in encrypted format and specify the encryption type, which helps transfer passwords between devices without having to know the passwords.

Figure 54: Enable Password Configuration

Table 47: Enable Password Configuration Fields

Field	Description
Enable Password	Specify the password all users must enter after executing the enable command at the CLI prompt.
Confirm Enable Password	Confirms the new enable password. The password appears in the ***** format.
Encryption Type	Select the option to specify encryption type. Options are AES or MD5-Salt. The default encryption type is AES.
Encrypted Password	Select the option to indicate whether the password is specified in encrypted format.

If you change any data, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

### 4.7.30 Password Rules

Use this page to configure settings that apply to all user passwords.

To display the page, click System > Passwords > Password Rules in the navigation menu.

Figure 55: Password Rules

The screenshot shows the 'Password Rules' configuration page. At the top, there is a breadcrumb 'System > Passwords > Password Rules' and buttons for 'Save Configuration', 'Hide Device View', and 'Log Out'. Below this are navigation tabs for 'System', 'Switching', 'Routing', 'Security', and 'QoS'. The main content area has tabs for 'Line Password', 'Enable Password', 'Password Rules' (selected), 'Last Password', and 'Reset Passwords'. The 'Password Rules' section contains the following fields:

- Minimum Length: 8 (0 to 64)
- Aging (Days): 0 (1 to 365, 0 = Default, 0 = Disable)
- History: 0 (0 to 10)
- Lockout Attempts: 0 (0 to 5, 0 = Default, 0 = Disable)
- Strength Check:  Disable  Enable
- Minimum Number of Uppercase Letters: 2 (0 to 16, 2 = Default, 0 = Disable)
- Minimum Number of Lowercase Letters: 2 (0 to 16, 2 = Default, 0 = Disable)
- Minimum Number of Numeric Characters: 2 (0 to 16, 2 = Default, 0 = Disable)
- Minimum Number of Special Characters: 2 (0 to 16, 2 = Default, 0 = Disable)
- Maximum Number of Repeated Characters: 0 (0 to 15, 0 = Default, 0 = Disable)
- Maximum Number of Consecutive Characters: 0 (0 to 15, 0 = Default, 0 = Disable)
- Minimum Character Classes: 4 (0 to 4, 4 = Default, 0 = Disable)

Below these fields is an 'Exclude Keyword Name' field with a search icon and a table that is currently empty. At the bottom, there are 'Submit', 'Refresh', and 'Cancel' buttons.

Table 48: Password Rules Fields

Field	Description
Minimum Length	Passwords must have at least this many characters (8 to 64).
Aging (days)	The number of days that a user password is valid from the time the password is set. Once a password expires, the user is required to enter a new password at the next login.
History	The number of previous passwords that are retained to prevent password reuse. This helps to ensure that a user does not attempt to reuse the same password too often.
Lockout Attempts	After a user fails to log in this number of times, the user is locked out until the password is reset by the administrator.
Strength Check	Enable or disable the password strength check feature. Enabling this feature forces the user to configure passwords that comply with the strong password configuration specified in the following fields.
Minimum Number of Uppercase Letters	Specify the minimum number of uppercase letters a password must include.



**Table 48: Password Rules Fields (Continued)**

Field	Description
Minimum Number of Lowercase Letters	Specify the minimum number of lowercase letters a password must include.
Minimum Number of Numeric Characters	Specify the minimum number of numbers a password must include.
Minimum Number of Special Characters	Specify the minimum number of special characters (non-alphanumeric, such as # or &) a password must include.
Maximum Number of Repeated Characters	Specify the maximum number of repeated characters a password is allowed to include. An example of four repeated characters is aaaa.
Maximum Number of Consecutive Characters	Specify the maximum number of consecutive characters a password is allowed to include. An example of four consecutive characters is abcd
Minimum Character Classes	Specify the minimum number of character classes a password must contain. There are four character classes: <ul style="list-style-type: none"> <li>• Uppercase</li> <li>• Lowercase</li> <li>• Numbers</li> <li>• Special Characters</li> </ul>
Exclude Keyword	The password to be configured should not contain the keyword mentioned in this field. The valid range for the keyword is (2 to 64) characters in length.
Exclude Keyword Name	The list of keywords that a valid password must not contain. Excluded keyword checking is case-insensitive. Additionally, a password cannot contain the backwards version of an excluded keyword. For example, if pass is an excluded keyword, passwords such as 23passA2c, ssapword, and PAsSwoRD are prohibited. Use the plus and minus buttons to perform the following tasks: <ul style="list-style-type: none"> <li>• To add a keyword to the list, click the + (plus) button, type the word to exclude in the Exclude Keyword Name field, and click Submit.</li> <li>• To remove a keyword from the list, click the – (minus) button associated with the keyword to remove and confirm the action.</li> <li>• To remove all keywords from the list, click the – (minus) button in the header row and confirm the action.</li> </ul>

If you change any data, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

### 4.7.31 Last Password Result

Use the Last Password Result page view information about the last attempt to set a user password. If the password set was unsuccessful, a reason for the failure is given.

To display the page, click System > Password > Last Password in the navigation menu.

**Figure 56: Last Password Result**






Table 49: Last Password Result

Field	Description
Last Password Set Result	Displays information about the last (User/Line/Enable) password configuration result. If the field is blank, no passwords have been configured on the device. Otherwise, the field shows that the password was successfully set or provides information about the type of password configuration that failed and why it could not be set.
Strength Check	Displays Enabled if Strength Check is applied in last password change, otherwise it displays Disabled.

### 4.7.32 Denial of Service

Use the Denial of Service (DoS) page to configure DoS control. FASTPATH software provides support for classifying and blocking specific types of DoS attacks. You can configure your system to monitor and block these types of attacks:

- SIP=DIP: Source IP address = Destination IP address.
- First Fragment: TCP Header size smaller than configured value.
- TCP Fragment: IP Fragment Offset = 1.
- TCP Flag: TCP Flag SYN set and Source Port < 1024 or TCP Control Flags = 0 and TCP Sequence Number = 0 or TCP Flags FIN, URG, and PSH set and TCP Sequence Number = 0 or TCP Flags SYN and FIN set.
- L4 Port: Source TCP/UDP Port = Destination TCP/UDP Port.
- ICMP: Limiting the size of ICMP Ping packets.

#### NOTICE

Monitoring and blocking of the types of attacks listed below are supported only on the following platforms:

- BCM56514
- BCM56624
- BCM56820
- BCM56224
- BCM56634
- BCM56636

- SMAC=DMAC: Source MAC address=Destination MAC address.
- TCP Port: Source TCP Port = Destination TCP Port.
- UDP Port: Source UDP Port = Destination UDP Port.
- TCP Flag & Sequence: TCP Flag SYN set and Source Port < 1024 or TCP Control Flags = 0 and TCP Sequence Number = 0 or TCP Flags FIN, URG, and PSH set and TCP Sequence Number = 0 or TCP Flags SYN and FIN set.
- TCP Offset: TCP Header Offset = 1.
- TCP SYN: TCP Flag SYN set.
- TCP SYN & FIN: TCP Flags SYN and FIN set.
- TCP FIN & URG & PSH: TCP Flags FIN and URG and PSH set and TCP Sequence Number = 0.
- ICMP V6: Limiting the size of ICMPv6 Ping packets.
- ICMP Fragment: Checks for fragmented ICMP packets.
- Smurf Attack: A flood of spoofed broadcast ping messages are sent to the system.
- PingFlood Attack: Similar to a Smurf Attack, a flood of ping packets are sent to the system.
- SYN ACK Flood Attack: A series of SYN requests are sent to force the switch to reply with SYN-ACK messages.

To access the Denial of Service page, click System > Advanced Configuration > Protection > Denial of Service in the navigation menu.

Figure 57: Denial of Service Configuration

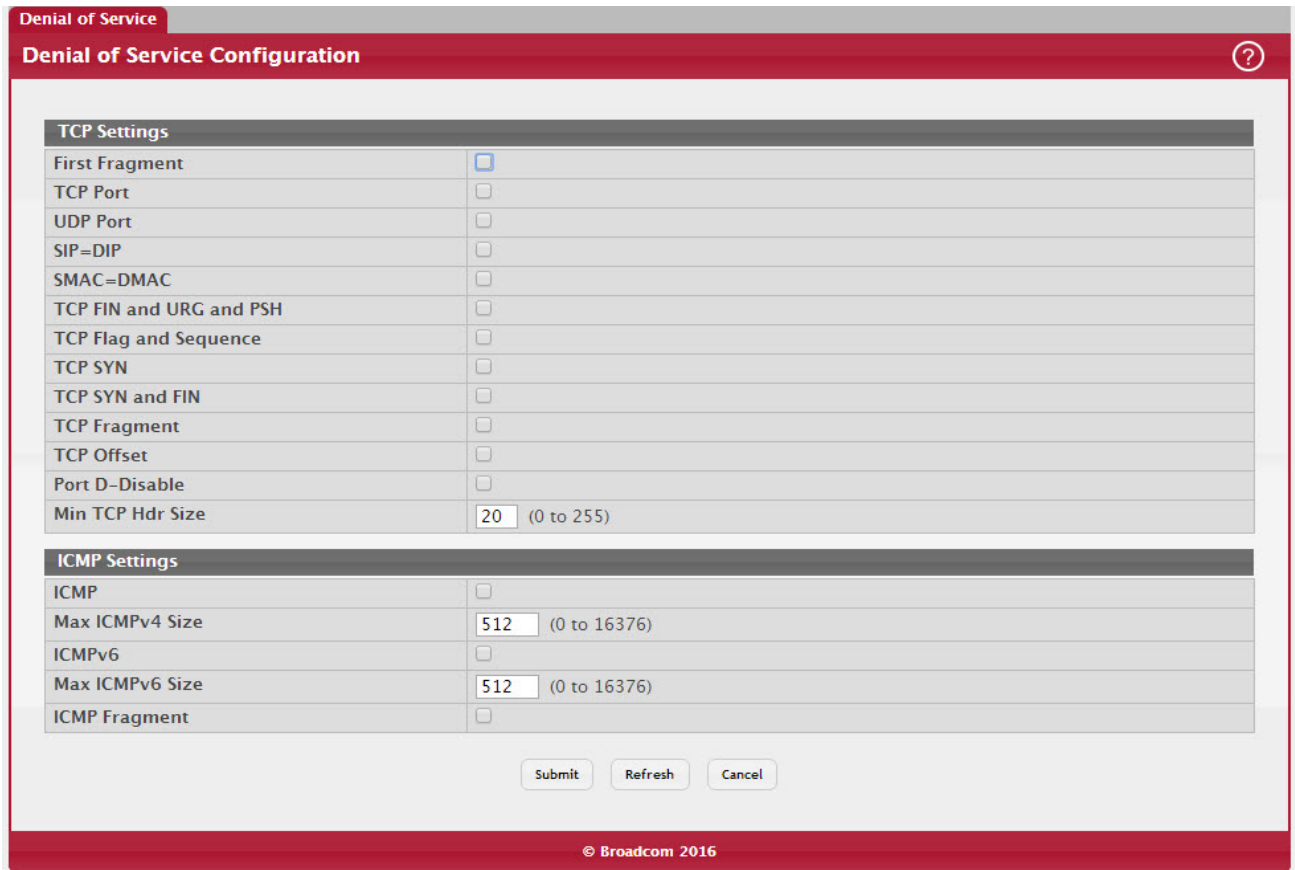


Table 50: Denial of Service Configuration Fields

Field	Description
TCP Settings	
First Fragment	Enable this option to allow the device to drop packets that have a TCP header smaller than the value configured in the Min TCP Hdr Size field.
TCP Port	Enable this option to allow the device to drop packets that have the TCP source port equal to the TCP destination port.
UDP Port	Enable this option to allow the device to drop packets that have the UDP source port equal to the UDP destination port.
SIP=DIP	Enable this option to allow the device to drop packets that have a source IP address equal to the destination IP address.
SMAC=DMAC	Enable this option to allow the device to drop packets that have a source MAC address equal to the destination MAC address.
TCP FIN and URG and PSH	Enable this option to allow the device to drop packets that have TCP Flags FIN, URG, and PSH set and a TCP Sequence Number equal to 0.
TCP Flag and Sequence	Enable this option to allow the device to drop packets that have TCP control flags set to 0 and the TCP sequence number set to 0.
TCP SYN	Enable this option to allow the device to drop packets that have TCP Flags SYN set.
TCP SYN and FIN	Enable this option to allow the device to drop packets that have TCP Flags SYN and FIN set.

Table 50: Denial of Service Configuration Fields (Continued)

Field	Description
TCP Fragment	Enable this option to allow the device to drop packets that have a TCP payload where the IP payload length minus the IP header size is less than the minimum allowed TCP header size.
TCP Offset	Enable this option to allow the device to drop packets that have a TCP header Offset set to 1.
Port D-Disable	Enable this option to allow the system to diagnostically disable an interface if a potential DoS attack has been detected on that interface. If an interface is diagnostically disabled, it remains in the disabled state until an administrator manually enables the interface.
Min TCP Hdr Size	The minimum TCP header size allowed. If First Fragment DoS prevention is enabled, the device will drop packets that have a TCP header smaller than this configured value.
ICMP Settings: These options help prevent the device and the network from attacks that involve issues with the ICMP echo request packets (pings) that the device receives.	
ICMP	Enable this option to allow the device to drop ICMP packets that have a type set to ECHO_REQ (ping) and a payload size greater than the ICMP payload size configured in the Max ICMPv4 Size or Max ICMPv6 Size fields.
ICMP Fragment	Enable this option to allow the device to drop fragmented ICMP packets.
Max ICMPv4 Size	The maximum allowed ICMPv4 packet size. If ICMP DoS prevention is enabled, the device will drop ICMPv4 ping packets that have a size greater than this configured maximum ICMPv4 packet size.
Max ICMPv6 Size	The maximum allowed IPv6 ICMP packet size. If ICMP DoS prevention is enabled, the switch will drop IPv6 ICMP ping packets that have a size greater than this configured maximum ICMPv6 packet size.

If you change any of the DoS settings, click Submit to apply the changes to the switch. To preserve the changes across a switch reboot, you must perform a save.

## 4.8 Configuring and Searching the Forwarding Database

The forwarding database maintains a list of MAC addresses after having received a packet from this MAC address. The transparent bridging function uses the forwarding database entries to determine how to forward a received frame.

### 4.8.1 Switch Configuration

Use the Switch Configuration page to set the amount of time to keep a learned MAC address entry in the forwarding database. The forwarding database contains static entries, which are never aged out, and dynamically learned entries, which are removed if they are not updated within a given time.

To access the Configuration page, click System > Basic Configuration > Switch in the navigation menu.

Figure 58: Switch Configuration

The screenshot shows the 'Switch Configuration' page. At the top, there is a red header with the text 'Switch Configuration' and a help icon. Below the header, there are two configuration fields:

- 802.3x Flow Control Mode:** This field has two radio buttons, 'Disable' (which is selected) and 'Enable'.
- MAC Address Aging Interval (Seconds):** This field has a text input box containing the value '300' and a range indicator '(10 to 1000000)'.

At the bottom of the configuration area, there are three buttons: 'Submit', 'Refresh', and 'Cancel'.

Table 51: Switch Configuration Fields

Field	Description
802.3x Flow Control Mode	Enable or disable 802.3x flow control on the switch. IEEE 802.3x flow control works by pausing a port when the port becomes oversubscribed. It also allows a port to drop all traffic for small bursts of time during the congestion condition. This can lead to high-priority and/or network control traffic loss. When enabled, flow control allows lower speed switches to communicate with higher-speed switches by requesting that the higher-speed switch refrain from sending packets. Transmissions are temporarily halted to prevent buffer overflows.
MAC Address Aging Interval	The MAC address table (forwarding database) contains static entries, which never age out, and dynamically-learned entries, which are removed if they are not updated within a given time. Specify the number of seconds a dynamic address should remain in the MAC address table after it has been learned.

---

IEEE 802.1D recommends a default of 300 seconds, which is the factory default.

### **NOTICE**

Click Submit to apply the changes to the system. You must perform a save to make the changes persist across a reboot.

## 4.9 Managing Logs

The switch may generate messages in response to events, faults, or errors occurring on the platform as well as changes in configuration or other occurrences. These messages are stored both locally on the platform and forwarded to one or more centralized points of collection for monitoring purposes as well as long term archival storage. Local and remote configuration of the logging capability includes filtering of messages logged or forwarded based on severity and generating component.

The in-memory log stores messages in memory based upon the settings for message component and severity. On stackable systems, this log exists only on the management unit. Other platforms in the stack forward their messages to the management unit log. Access to in-memory logs on other than the management unit is not supported.

### 4.9.1 Log Configuration

The Log Configuration page allows administrators with the appropriate privilege level to configure the administrative mode and various settings for logging features on the switch.

To access the Log Configuration page, click System > Logs > Configuration in the navigation menu.

Figure 59: Log Configuration

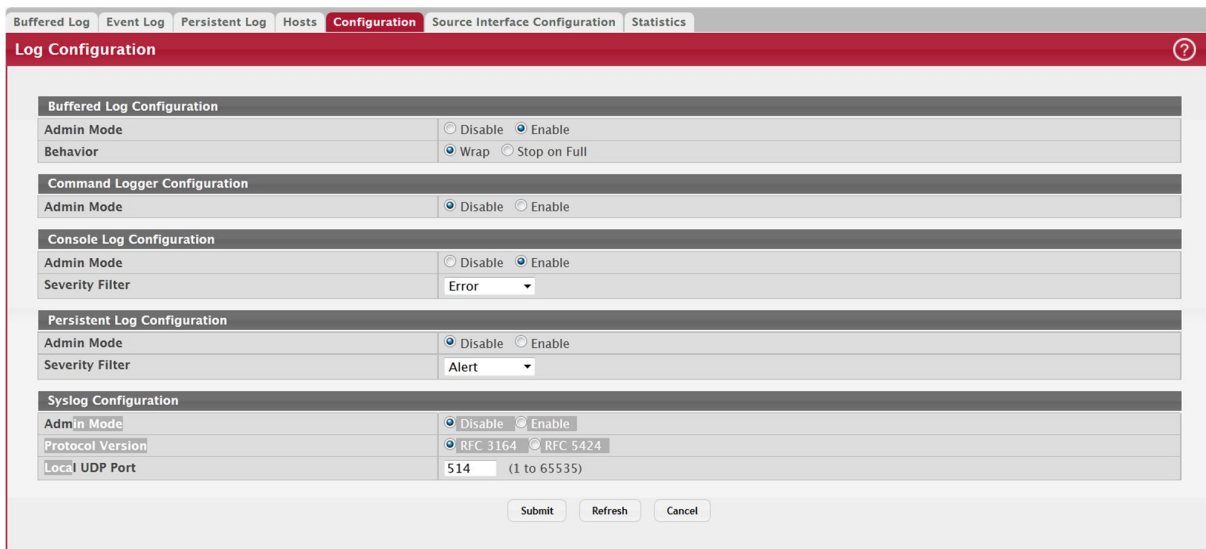


Table 52: Log Configuration Fields

Field	Description
<b>Buffered Log Configuration</b>	
Admin Mode	Enable or disable logging to the buffered (RAM) log file.
Behavior	Specify what the device should do when the buffered log is full. It can either overwrite the oldest messages (Wrap) or stop writing new messages to the buffer (Stop on Full).
<b>Command Logger Configuration</b>	
Admin Mode	Enable or disable logging of the command-line interface (CLI) commands issued on the device.
<b>Console Log Configuration</b>	
Admin Mode	Enable or disable logging to any serial device attached to the host.
Severity Filter	Select the severity of the messages to be logged. All messages at and above the selected threshold are logged to the console. The severity can be one of the following: <ul style="list-style-type: none"> <li>Emergency (0): The device is unusable.</li> <li>Alert (1): Action must be taken immediately.</li> <li>Critical (2): The device is experiencing primary system failures.</li> <li>Error (3): The device is experiencing non-urgent failures.</li> <li>Warning (4): The device is experiencing conditions that could lead to system errors if no action is taken.</li> <li>Notice (5): The device is experiencing normal but significant conditions.</li> <li>Info (6): The device is providing non-critical information.</li> <li>Debug (7): The device is providing debug-level information.</li> </ul>
<b>Persistent Log Configuration</b>	
Admin Mode	Enable or disable logging to the persistent log. These messages are not deleted when the device reboots.
Severity Filter	Select the severity of the messages to be logged. All messages at and above the selected threshold are logged to the console. See the previous severity filter description for more information about each severity level.

Table 52: Log Configuration Fields (Continued)

Field	Description
Syslog Configuration	
Admin Mode	Enable or disable logging to configured syslog hosts. When the syslog admin mode is disabled the device does not relay logs to syslog hosts, and no messages will be sent to any collector/relay. When the syslog admin mode is enabled, messages will be sent to configured collectors/relays using the values configured for each collector/relay.
Protocol Version	The RFC version of the syslog protocol.
Local UDP Port	The UDP port on the local host from which syslog messages are sent.

If you change the buffered log settings, click Submit to apply the changes to the system. To preserve the changes after a system reboot, you must perform a save.

## 4.9.2 Buffered Log

The log messages the device generates in response to events, faults, errors, and configuration changes are stored locally on the device in the RAM (cache). This collection of log files is called the RAM log or buffered log. When the buffered log file reaches the configured maximum size, the oldest message is deleted from the RAM when a new message is added. If the system restarts, all messages are cleared.

To access the Buffered Log page, click System > Logs > Buffered Log in the navigation menu.

Figure 60: Buffered Log

The screenshot shows the 'Buffered Log' page with the following data:

Log Index	Log Time	Severity	Component	Description
1	Apr 2 14:38:12	Info	USER_MGR	HTTP Session 17 started for user admin connected from 10.12.17.130
2	Apr 2 13:55:30	Info	USER_MGR	HTTP Session 16 ended for user admin connected from 10.12.17.130
3	Apr 2 10:51:35	Info	USER_MGR	HTTP Session 16 started for user admin connected from 10.12.17.130
4	Apr 2 10:51:29	Info	USER_MGR	HTTP Session 15 ended for user admin connected from 10.12.17.130
5	Apr 2 08:24:03	Info	USER_MGR	HTTP Session 15 started for user admin connected from 10.12.17.130
6	Apr 1 17:07:09	Info	USER_MGR	HTTP Session 14 ended for user admin connected from 10.12.17.130
7	Apr 1 15:15:59	Info	USER_MGR	HTTP Session 14 started for user admin connected from 10.12.17.130
8	Apr 1 13:38:35	Info	USER_MGR	HTTP Session 13 ended for user admin connected from 10.12.17.130
9	Apr 1 09:25:20	Info	USER_MGR	HTTP Session 13 started for user admin connected from 10.12.17.130
10	Apr 1 09:06:38	Info	USER_MGR	HTTP Session 12 ended for user admin connected from 10.27.65.54

**Table 53: Buffered Log Fields**

Field	Description
Log Index	The position of the entry within the buffered log file. The most recent log message always has a Log Index value of 1.
Log Time	The time the entry was added to the log.
Severity	The severity level associated with the log entry. The severity can be one of the following: <ul style="list-style-type: none"> <li>Emergency (0): The device is unusable.</li> <li>Alert (1): Action must be taken immediately.</li> <li>Critical (2): The device is experiencing primary system failures.</li> <li>Error (3): The device is experiencing non-urgent failures.</li> <li>Warning (4): The device is experiencing conditions that could lead to system errors if no action is taken.</li> <li>Notice (5): The device is experiencing normal but significant conditions.</li> <li>Info (6): The device is providing non-critical information.</li> <li>Debug (7): The device is providing debug-level information.</li> </ul>
Component	The component that issued the log entry.
Description	The text description for the log entry.

Click Refresh to update the screen and associated messages.

### 4.9.3 Event Log

Use the Event Log page to display the event log, which is used to hold error messages for catastrophic events. After the event is logged and the updated log is saved in flash memory, the switch will be reset. The log can hold at least 2,000 entries (the actual number depends on the platform and OS), and is erased when an attempt is made to add an entry after it is full. The event log is preserved across system resets.

To access the Event Log page, click System > Logs > Event Log in the navigation menu.

**Figure 61: Event Log**

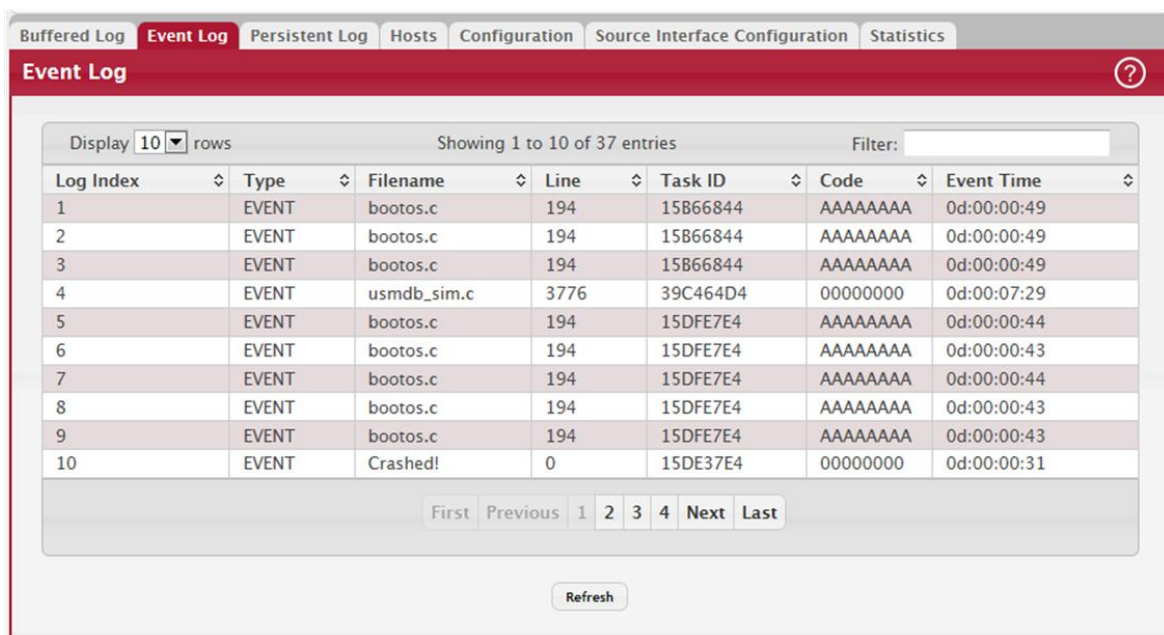




Table 54: Event Log Fields

Field	Description
Entry	The number of the entry within the event log. The most recent entry is first.
Type	The incident category that indicates the cause of the log entry: EVENT, ERROR, etc.
Filename	The FASTPATH source code file name identifying the code that detected the event.
Line	The line number within the source file of the code that detected the event.
Task ID	The OS-assigned ID of the task reporting the event.
Code	The event code passed to the event log handler by the code reporting the event.
Time	The time the event occurred, measured from the previous reset.

Click Refresh to update the screen and associated messages.

#### 4.9.4 Hosts Log Configuration

Use the Host Log Configuration page to configure remote logging hosts where the switch can send logs.

To access the Host Log Configuration page, click System > Logs > Hosts in the navigation menu.

Figure 62: Logging Hosts

Host	Status	Port	Severity Filter	Transport Mode	Authentication Mode	Certificate Index
1.2.3.4	Active	100	Debug	TLS	anon	
2.2.2.2	Active	200	Notice	UDP		
8.8.8.8	Active	514	Debug	TLS	anon	
9.9.9.9	Active	6514	Debug	TLS	x509name	4

Table 55: Logging Hosts Fields

Field	Description
Host (IP Address/Host Name)	The IP address or DNS-resolvable host name of the remote host to receive log messages.
Status	Indicates whether the host has been configured to be actively logging or not.
Port	The UDP port on the logging host to which syslog messages are sent.
Severity Filter	Severity level threshold for log messages. All log messages with a severity level at and above the configured level are forwarded to the logging host.
Transport Mode	Transport mode used while sending messages to syslog servers. Supported modes are UDP and TLS. If TLS is not configured, default transport mode is UDP.
Authentication Mode	Using TLS security user can configure anonymous authentication mode, in which no client authentication is done by the syslog server. For x509/name authentication mode, two-way authentication is done both by syslog client and client authentication by syslog server side.
Certificate Index	The index used for identifying corresponding certificate files.



Use the buttons to perform the following tasks:

- To add a logging host, click Add and configure the desired settings.
- To change information for an existing logging host, select the check box associated with the entry and click Edit. You cannot edit the host name or address of a host that has been added.
- To delete a configured logging host from the list, select the check box associated with each entry to delete and click Remove.

**Figure 63: Add Host**

After you add a logging host, the screen displays additional fields.

**Table 56: Add Host Fields**

Field	Description
IP Address/Host Name	The IP address or DNS-resolvable host name of the remote host to receive log messages.
Transport Mode	Transport mode used while sending messages to syslog servers. Supported modes are UDP and TLS. If TLS is not configured then default transport mode is UDP.
Authentication Mode	Using TLS security user can configure anonymous authentication mode, in which no client authentication is done by the syslog server. For x509/name authentication mode, two way authentication is done both by syslog client and client authentication by syslog server side.
Certificate Index	The index used for identifying corresponding certificate files.
Port	The UDP port on the logging host to which syslog messages are sent.
Severity Filter	Severity level threshold for log messages. All log messages with a severity level at and above the configured level are forwarded to the logging host.

#### 4.9.4.1 Adding a Remote Logging Host

Use the following procedures to add, configure, or delete a remote logging host.

1. From the Host field, select Add to add a new host, or select the IP address of an existing host to configure the host.

If you are adding a new host, enter the IP address of the host in the IP Address field and click Submit. The screen refreshes, and additional fields appear.

2. In the Port field, type the port number on the remote host to which logs should be sent.
3. Select the severity level of the logs to send to the remote host.
4. Click Submit to apply the changes to the system.

### 4.9.4.2 Deleting a Remote Logging Host

To delete a remote logging host from the configured list, select the IP address of the host from the Host field, and then click Delete.

### 4.9.5 Syslog Source Interface Configuration

Use this page to specify the physical or logical interface to use as the logging (Syslog) client source interface. When an IP address is configured on the source interface, this address is used for all Syslog communications between the local logging client and the remote Syslog server. The IP address of the designated source interface is used in the IP header of Syslog management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access the Syslog Source Interface Configuration page, click System > Logs > Source Interface Configuration in the navigation menu.

Figure 64: Syslog Source Interface Configuration

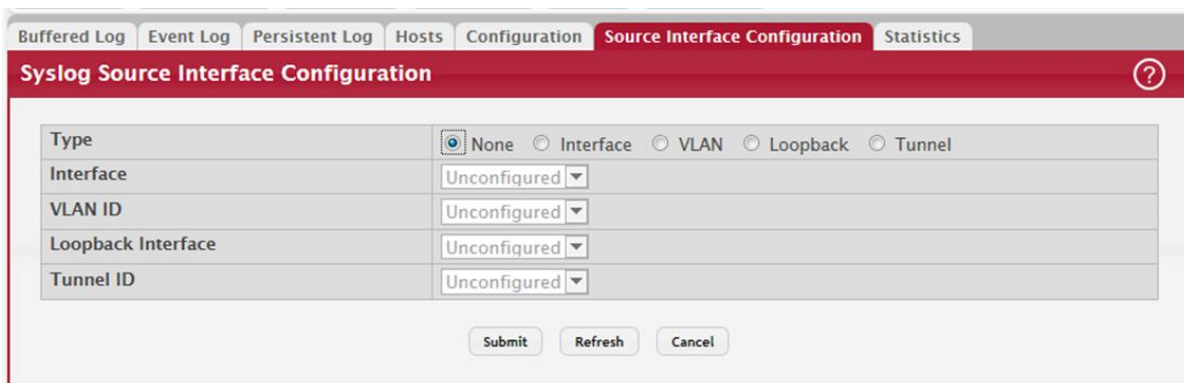


Table 57: Syslog Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> <li>• None – The primary IP address of the originating (outbound) interface is used as the source address.</li> <li>• Interface – The primary IP address of a physical port is used as the source address.</li> <li>• Loopback – The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up.</li> <li>• VLAN – The primary IP address of a VLAN routing interface is used as the source address.</li> </ul>
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.

If you change any of the settings on the page, click Submit to apply the changes to system.

### 4.9.6 Persistent Log

Use the Persistent Log page to view the persistent log messages.

To access the Persistent Log page, click System > Log > Persistent Log in the navigation menu.

Figure 65: Persistent Log

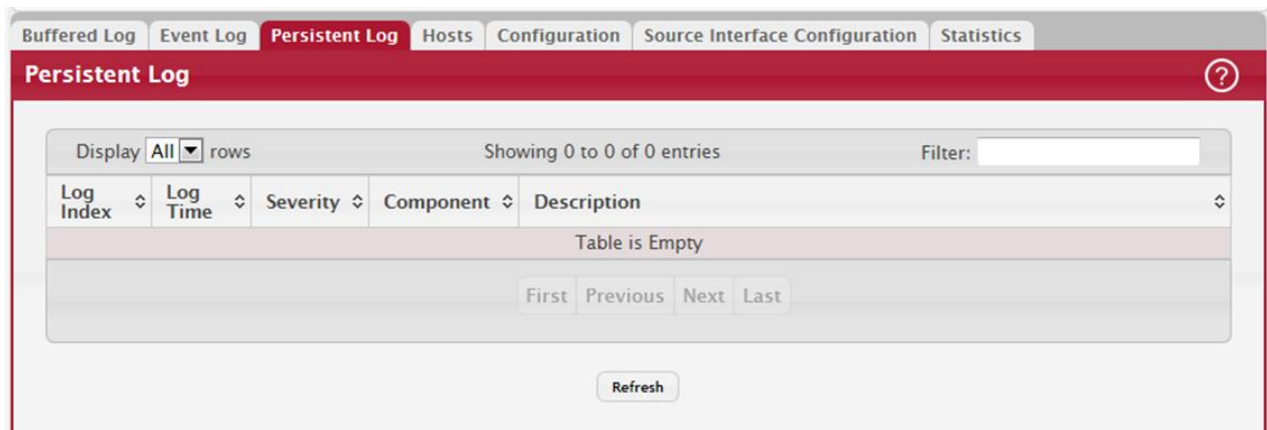


Table 58: Persistent Log Fields

Field	Description
Log Index	The position of the entry within the buffered log file. The most recent log message always has a Log Index value of 1.
Log Time	The time the entry was added to the log.
Severity	The severity level associated with the log entry. The severity can be one of the following: <ul style="list-style-type: none"> <li>Emergency (0): The device is unusable.</li> <li>Alert (1): Action must be taken immediately.</li> <li>Critical (2): The device is experiencing primary system failures.</li> <li>Error (3): The device is experiencing non-urgent failures.</li> <li>Warning (4): The device is experiencing conditions that could lead to system errors if no action is taken.</li> <li>Notice (5): The device is experiencing normal but significant conditions.</li> <li>Info (6): The device is providing non-critical information.</li> <li>Debug (7): The device is providing debug-level information.</li> </ul>
Component	The component that has issued the log entry.
Description	The text description for the log entry.

## 4.10 Configuring Email Alerts

With the email alerting feature, log messages can be sent to one or more email addresses. You must configure information about the network Simple Mail Transport Protocol (SMTP) server for email to be successfully sent from the switch.

The pages available from the Email Alerting folder allow you to configure information about what type of log message are sent via email and to what address(es) the messages are emailed.

### 4.10.1 Email Alert Global Configuration

Use the Email Alert Global Configuration page to configure the common settings for log messages emailed by the switch.

To access the Email Alert Global Configuration page, click System > Advanced Configuration > Email Alerts > Global in the navigation menu.

Figure 66: Email Alert Global Configuration

Table 59: Email Alert Global Configuration Fields

Field	Description
Admin Mode	Sets the administrative mode of the feature. <ul style="list-style-type: none"> <li>• Enable – The device can send email alerts to the configured SMTP server.</li> <li>• Disable – The device will not send email alerts.</li> </ul>
From Address	Specifies the email address of the sender (the switch).
Log Duration	This duration in minutes determines how frequently the non critical messages are sent to the SMTP Server.
Urgent Messages Severity	Configures the severity level for log messages that are considered to be urgent. Messages in this category are sent immediately. The security level you select and all higher levels are urgent: <ul style="list-style-type: none"> <li>• Emergency indicates system is unusable. It is the highest level of severity.</li> <li>• Alert indicates action must be taken immediately</li> <li>• Critical indicates critical conditions</li> <li>• Error indicates error conditions</li> <li>• Warning indicates warning conditions</li> <li>• Notice indicates normal but significant conditions</li> <li>• Informational indicates informational messages</li> <li>• Debug indicates debug-level messages</li> </ul>
Non Urgent Messages Severity	Configures the severity level for log messages that are considered to be nonurgent. Messages in this category are collected and sent in a digest form at the time interval specified by the Log Duration field. The security level you select and all levels up to, but not including the lowest Urgent level are considered nonurgent. Messages below the security level you specify are not sent via email. See the Urgent Message field description for information about the security levels.
Traps Severity	Configures the severity level for trap log messages. See the Urgent Message field description for information about the security levels.

If you make any changes to the page, click Submit to apply the change to the system.

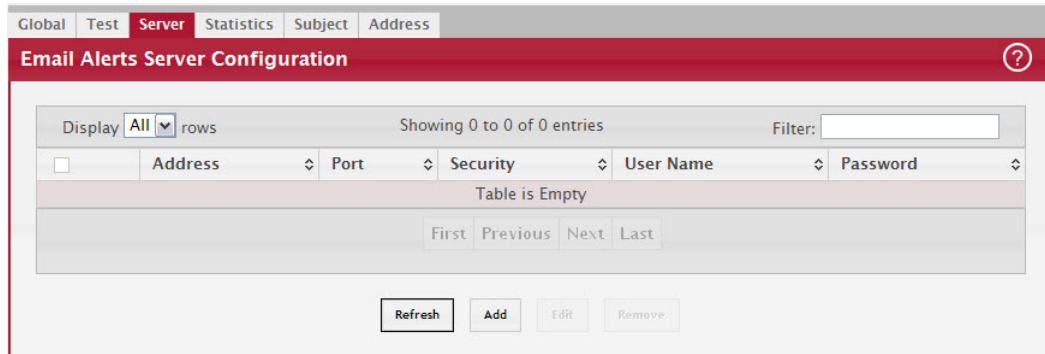
After configuring all email alert settings, click Test to send a test message to the configured address(es).

### 4.10.2 Email Alerts Server Configuration

Use the Email Alerts Server Configuration page to configure information about up to three SMTP (mail) servers on the network that can handle email alerts sent from the switch.

To access the Email Alerts Server Configuration page, click System > Advanced Configuration > Email Alerts > Server in the navigation menu.

Figure 67: Email Alerts Server Configuration



Use the buttons to perform the following tasks:

- To add an SMTP server, click Add and configure the desired settings.
- To change information for an existing SMTP server, select the check box associated with the entry and click Edit. You cannot edit the host name or address of a server that has been added.
- To delete a configured SMTP server from the list, select the check box associated with the entry to delete and click Remove.

Figure 68: Add New Email Server

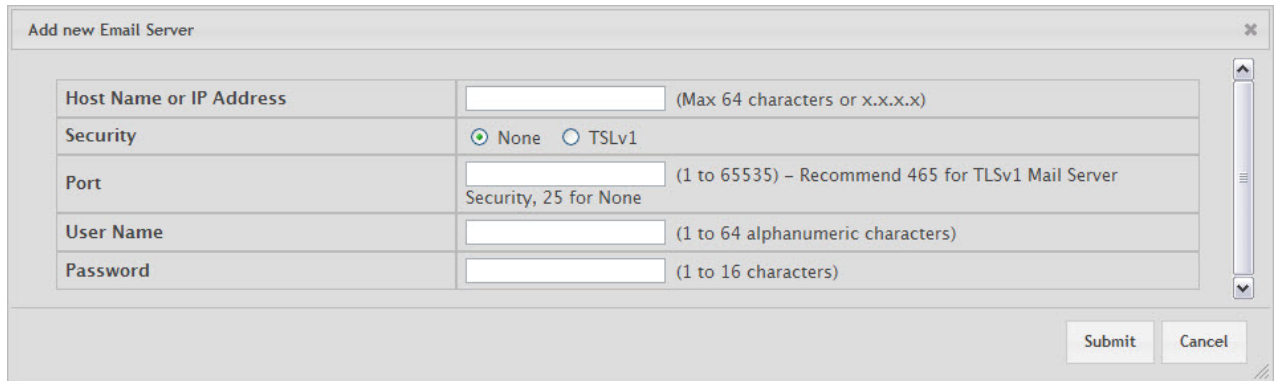


Table 60: Add New Email Server Fields

Field	Description
Host Name or IP Address	Shows the address or host name of the SMTP server that handles email alerts that the device sends.
Security	Specifies the type of authentication to use with the mail server, which can be TLSv1 (SMTP over SSL) or None (no authentication is required).
Port	Specifies the TCP port that email alerts are sent to on the SMTP server.

Table 60: Add New Email Server Fields (Continued)

Field	Description
User Name	If the Security is TLSv1, this field specifies the user name required to access the mail server.
Password	If the Security is TLSv1, this field specifies the password associated with the configured user name for mail server access. When adding or editing the server, you must retype the password to confirm that it is entered correctly.

If you make any changes to the page, click Submit to apply the change to the system. To remove a configured SMTP server, select the Remove check box and click Delete.

### 4.10.3 Email Alert Statistics

Use the Email Alert Statistics page to view information about email alerts sent from the switch.

To access the Email Alert Statistics page, click System > Advanced Configuration > Email Alerts > Statistics in the navigation menu.

Figure 69: Email Alert Statistics

Field	Description
Number of Emails Sent	0
Number of Emails Failed	0
Time Since Last Email Sent	0 days, 0 hours, 0 mins, 0 secs

Table 61: Email Alert Statistics Fields

Field	Description
Number of Emails Sent	Displays the number of email alert messages sent since last reset.
Number of Emails Failed	Displays the number of email alert messages that were unable to be sent since last reset.
Time Since Last Email Sent	Time that has passed since the last email alert message was sent successfully.

To update the page with the most current information, click Refresh. To reset the values on the page to zero, click Clear Counters.

### 4.10.4 Email Alert Subject Configuration

Use the Email Alert Subject Configuration page to configure the subject line of the email alert messages sent from the switch.

To access the Email Alert Subject Configuration page, click System > Advanced Configuration > Email Alerts > Subject in the navigation menu.

Figure 70: Email Alert Subject Configuration

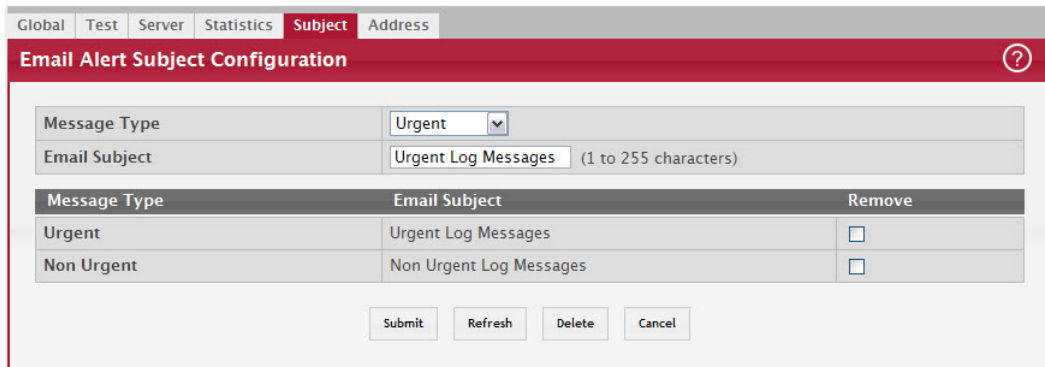


Table 62: Email Alert Subject Configuration Fields

Field	Description
Message Type	Select the appropriate option to configure the subject line of Urgent messages or Nonurgent messages.
Email Subject	Specify the text to be displayed in the subject of the email alert message.
Remove	To reset the email alert subject to the default value, select the Remove option associated with the message type to reset, and click Delete.

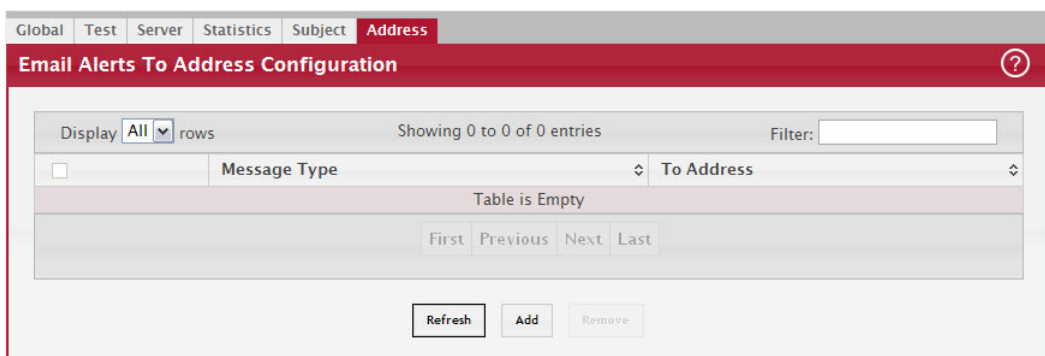
If you make any changes to the page, click Submit to apply the change to the system. To remove a configured Email Subject, select the Remove check box associated with the entry and click Delete.

#### 4.10.5 Email Alerts To Address Configuration

Use the Email Alerts To Address Configuration page to configure the email addresses to which alert messages sent.

To access the Email Alerts To Address Configuration page, click System > Advanced Configuration > Email Alerts > Address in the navigation menu.

Figure 71: Email Alerts To Address Configuration



Use the buttons to perform the following tasks:

- To add an email address to the list of email alert message recipients, click Add and configure the desired settings.
- To delete an entry from the list, select the check box associated with each entry to delete and click Remove.

Table 63: Email Alerts To Address Configuration Fields

Field	Description
Message Type	Select the appropriate option to configure email address where Urgent messages or Nonurgent messages are sent.
To Address	Specify the email address to which the selected type of messages are sent.

If you make any changes to the page, click Submit to apply the change to the system. To remove a configured email address, select the Remove check box associated with the entry and click Delete.

## 4.11 Configuring and Viewing Device Slot Information

The pages in the Slot folder provide information about the cards installed in the slots on the switch. The physical location of the slots depends on the hardware on which FASTPATH software is running. From the Configuration page, you can also manually configure information about cards on some platforms.

### 4.11.1 Slot Card Configuration

Use the Card Configuration page to view information about the cards installed in a switch. On some platforms, you can manually configure information about slots.

To access the Card Configuration page, click System > Slot > Configuration in the navigation menu.

Figure 72: Slot Configuration

The screenshot displays the 'Slot Configuration' page within the FASTPATH software. At the top, there is a breadcrumb trail 'System > Slot > Configuration' and several utility buttons: 'Save Configuration', 'Hide Device View', and 'Log Out'. Below this, there are navigation tabs for 'System', 'Switching', 'Routing', 'Security', and 'QoS'. The main content area is titled 'Slot Configuration' and features a table with the following data:

Slot	Status	Administrative State	Power State	Card Model	Card Description
0	Full	Enable	Enable	BCM56634-48GIG-4TENCE	Broadcom BCM56634 - 48 Port 4 Ten-Gigabit Ethernet Line Card

Below the table, there are navigation buttons: 'Refresh', 'Add', 'Edit', and 'Remove'. The page also includes a 'Filter' input field and a 'Showing 1 to 1 of 1 entries' indicator.

Table 64: Slot Configuration Fields

Field	Description
Slot	Identifies the slot number.
Status	Indicates whether the slot is empty or full.
Administrative State	Indicates whether the slot is administratively enabled or disabled. For some devices, you can change the Administrative State when you add or edit slot information.
Power State	Indicates whether the device is providing power to the slot. For some devices, you can change the Power State when you add or edit slot information.



Table 64: Slot Configuration Fields (Continued)

Field	Description
Card Model	The model ID of the card configured for the slot.
Card Description	The description of the card configured for the slot.

Figure 73: "Edit Existing Card," on page 99 shows the fields that display when the slot contains a card.

Figure 73: Edit Existing Card

Edit Existing Card	
Unit	1
Slot	0
Card Index	5
Card Description	BCM56624-48GIG-4TENGE
Status	Full
Administrative State	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
Power State	<input checked="" type="radio"/> Enable <input type="radio"/> Disable

Use the buttons to perform the following tasks:

- To preconfigure a card before adding it to a slot, click Add and configure the desired settings.
- To change slot or card settings, select the check box associated with the entry and click Edit.
- To delete a slot configuration entry from the list, select the check box associated with each entry to delete and click Remove.

Table 65: Edit Existing Card Fields

Field	Description
Unit	Indicates the unit in the stack for which data is to be displayed or configured.
Slot	Indicates the slot in the selected unit for which data is to be displayed or configured.
Card Index	Identifies the index number assigned to the card. This value is helpful when configuring the system by using SNMP.
Card Description	The description of the card configured for the slot.
Status	Indicates whether the slot is empty or full.
Administrative State	Indicates whether the card can be administratively enabled or disabled. If the value is Disable, the Administrative State cannot be configured.
Power State	Indicates whether the Power State can be administratively enabled or disabled. If the value is Disable, the Power State cannot be configured.

- If you make any changes to the page, click Submit to apply the changes to the system.
- Click Refresh to redisplay the page with the current data from the switch.

## 4.11.2 Slot Supported Cards

The Slot Supported Cards page provides information about the cards that your platform supports.

To access the Slot Supported Cards page, click System > Slot > Supported Cards in the navigation menu.

Figure 74: Slot Supported Cards

Card Index	Supported Cards	Card Type	Card Model	Card Description
5	BCM56624-48GIG-4TENGE	0x56624101	BCM56624-48GIG-4TENGE	Broadcom BCM56624 - 48 Port 4Ten-Gigabit Ethernet Line Card
6	BCM56680-24GIG-4TENGE	0x56680101	BCM56680-24GIG-4TENGE	Broadcom BCM56680 - 24 Port 4Ten-Gigabit Ethernet Line Card
7	BCM56820-24TENGE-4GE	0x56820001	BCM56820-24TENGE-4GE	Broadcom BCM56820 - 24 Port 10GB + 4 Port 1GB Ethernet Line Card
8	BCM56634-48GIG-4TENGE	0x56634101	BCM56634-48GIG-4TENGE	Broadcom BCM56634 - 48 Port 4 Ten-Gigabit Ethernet Line Card
9	BCM56524-24GIG-4TENGE	0x56524001	BCM56524-24GIG-4TENGE	Broadcom BCM56524 - 24 Port 4Ten-Gigabit Ethernet Line Card
10	BCM56143-52GIG	0x56143001	BCM56143-52GIG	Broadcom BCM56143 - 52 Port Ethernet Line Card
11	BCM56524-24GIG-4TENGE-EEE	0x565240AE	BCM56524-24GIG-4TENGE-EEE	Broadcom BCM56524 - 24 EEE Port 4Ten-Gigabit Ethernet Line Card
12	BCM56636-25GIG-6TENGE	0x56636001	BCM56636-25GIG-6TENGE	Broadcom BCM56636 - 25 Port 6Ten-Gigabit Ethernet Line Card
13	BCM56538-48GIG-4TENGE	0x56538101	BCM56538-48GIG-4TENGE	Broadcom BCM56538 - 48 Port 4 Ten-Gigabit Ethernet Line Card
14	BCM56334-24GIG-4TENGE	0x56334001	BCM56334-24GIG-4TENGE	Broadcom BCM56334 - 24 Port 4Ten-Gigabit Ethernet Line Card

Table 66: Slot Supported Card Fields

Field	Description
Card Index	Displays the index assigned to the selected card type.
Supported Cards	The menu contains the list of all cards that the system can support. To view information about a card, select it from the drop-down list. The screen refreshes, and the information about that card appears in the other fields on the page.
Card Type	Displays the hardware type of this supported card. This is a 32-bit data field.
Card Model	Displays the string to identify the model of the supported card.
Card Descriptor	Displays a data field used to identify the supported card.

Click Refresh to redisplay the most current information from the router.

## 4.12 Configuring Power Over Ethernet (PoE) and PoE Statistics

Use these pages to view Power over Ethernet (PoE) status information, configure global PoE settings, configure PoE settings on interfaces and view PoE interface statistical information.

### 4.12.1 PoE Configuration

Use this page to view Power over Ethernet (PoE) status information and configure global PoE settings.

To access the PoE Configuration page, click System > PoE > Configuration in the navigation menu.

Figure 75: PoE Configuration

Field	Value
Firmware Version	1.3.2.3
Operational Status	Off
Total Power Available (mWatts)	380000
Threshold Power (mWatts)	342000
Consumed Power (mWatts)	0
System Usage Threshold (%)	90 (1 to 99)
Power Management Mode	<input type="radio"/> Static <input checked="" type="radio"/> Dynamic
Port Auto Reset Mode	<input type="checkbox"/>
Traps	<input checked="" type="radio"/> Enable <input type="radio"/> Disable

Table 67: PoE Configuration Fields

Field	Description
Firmware Version	The firmware version of the PoE software component.
Operational Status	The current status of the switch PoE functionality, which can be one of the following: <ul style="list-style-type: none"> <li>On – At least one port on the switch is delivering power to a connected device.</li> <li>Off – The PoE functionality is operational but no ports are delivering power.</li> <li>Faulty – The PoE functionality is not operational.</li> </ul>
Total Power Available	The total power in mWatts that can be provided by the switch.
Threshold Power	When the PoE power being used exceeds this threshold, a trap is generated to the system log to alert the system administrator of high power usage. This value is determined by the configurable System Usage Threshold percent.
Consumed Power	The amount of power in mWatts currently being consumed by connected PoE devices.
System Usage Threshold	A percentage of the total power available. This percentage determines the Threshold Power.
Power Management Mode	The method by which the PoE controller determines supplied power, which can be one of the following: <ul style="list-style-type: none"> <li>Static – The power allocated to each port is reserved and is not available to any other port, even when less than the maximum allocation is being used.</li> <li>Dynamic – The power allocated to each port is not reserved. Unused power may be allocated from one port to another as needed, up to the power limit defined for each port.</li> </ul>
Port Auto Reset Mode	When enabled, the switch automatically resets a PoE port if an error condition occurs. When disabled, the administrator must reset the port manually.
Traps	When enabled, SNMP traps will be generated when certain events occur. Trap events include a change in whether power is being delivered on a port and when the power usage threshold is exceeded.

- If you make any changes to the page, click Submit to apply the changes to the system.
- Click Refresh to redisplay the page with the current data from the switch.

### 4.12.2 PoE Port Configuration

Use this page to configure PoE settings on interfaces.

To access the PoE Configuration page, click System > PoE > Port Configuration in the navigation menu.

Figure 76: PoE Port Configuration

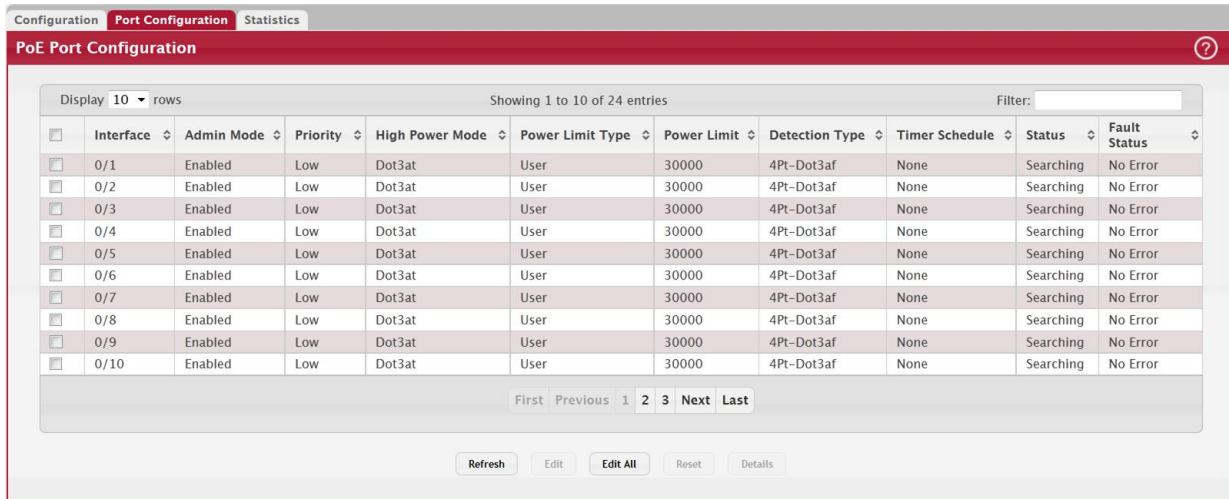


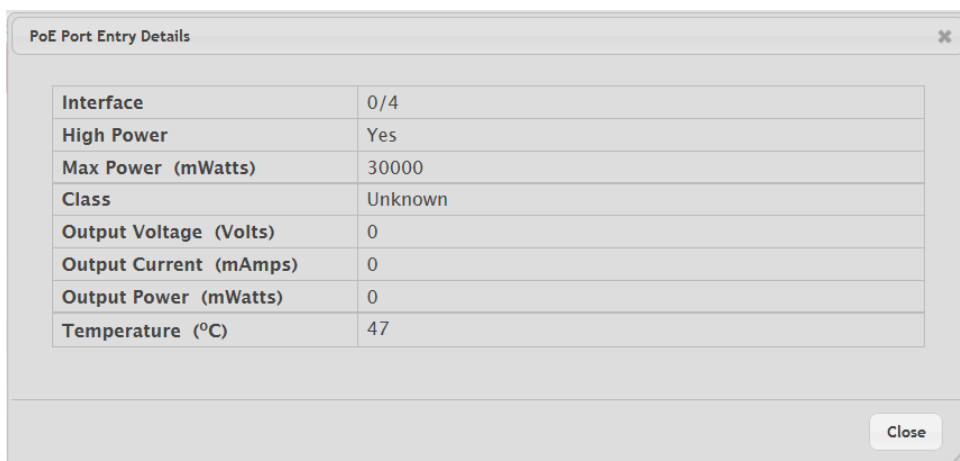
Table 68: PoE Port Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring PoE settings, this field identifies the interface(s) being configured.
Admin Mode	Indicates whether PoE is administratively enabled or disabled on the interface.
Priority	The priority of the port when allocating available power. Power is delivered to the higher priority ports when needed before providing it to the lower priority ports. Possible values are <i>Critical</i> , <i>High</i> , and <i>Low</i> .
High Power Mode	When enabled, the port supports the PoE+ power standard, which allows for providing up to 30W of power. When disabled, the port supports the original PoE standard only, which allows for providing up to 15.4W of power.
Power Limit Type	The type of power limiting used for the port, which can be one of the following: <ul style="list-style-type: none"> <li>• <i>Class</i> – The device class determines the power limit. The switch learns the class of the device through the receipt of Link Layer Discovery Protocol (LLDP) messages.</li> <li>• <i>User</i> – The power limit is user defined, overriding the LLDP information. When set to <i>User</i>, the Power Limit field is enabled.</li> </ul>
Power Limit	The power limit for the port, which can be specified. This field displays only when Power Limit Type is set to <i>User</i> .

**Table 68: PoE Port Configuration Fields (Continued)**

Field	Description
Detection Type	<p>The protocol(s) that can be used to detect the presence of a PD when connected to a PoE port. The IEEE specification 802.3af (Dot3af) specifies various detection algorithms. Some PDs use legacy detection algorithms that were in place prior to the 802.3af standard, which can be one of the following:</p> <ul style="list-style-type: none"> <li>• Legacy – The switch uses a legacy detection scheme not defined in 802.3af.</li> <li>• 4Pt-Dot3af – The switch uses the 802.3af 4-point detection scheme only.</li> <li>• 4Pt-Dot3af + Legacy – The switch uses the 802.3af 4-point detection scheme, followed by the legacy detection scheme.</li> <li>• 2Pt-Dot3af – The switch uses the 802.3af 2-point detection scheme.</li> <li>• 2Pt-Dot3af + Legacy – The switch uses the 802.3af 2-point detection scheme, followed by the legacy detection scheme.</li> <li>• None – No detection is performed.</li> </ul>
Timer Schedule	The time range from the list of time ranges configured on the system.
Status	<p>The status of the port as a provider of PoE. Such devices are referred to as PSE. The status can be one of the following:</p> <ul style="list-style-type: none"> <li>• Disabled – The PSE is disabled.</li> <li>• Delivering Power – The PSE is delivering power.</li> <li>• Fault – The PSE has experienced a fault condition.</li> <li>• Test – The PSE is in test mode.</li> <li>• Other Fault – The PSE has experienced a variable error condition.</li> <li>• Searching – The PSE is transitioning between states.</li> <li>• Requesting Power – The PSE is currently not able to deliver power because power is unavailable to the port.</li> </ul>
Fault Status	<p>The error when PSE port is in fault status, which can be one of the following:</p> <ul style="list-style-type: none"> <li>• None – PSE port is not in any error state.</li> <li>• MPS Absent – PSE port has detected absence of main power supply.</li> <li>• Short – PSE port has detected a short circuit condition.</li> <li>• Overload – PD connected to PSE port tried to draw more power than permissible by the hardware.</li> <li>• Power Denied – PSE port has been denied power due to administrative action or shortage of power.</li> </ul>

To display additional PoE interface information, select an entry and click Details.



The following information describes the fields in the Details window.

High Power	Indicates whether high power mode is enabled or disabled.
Max Power	If Power Limit Type for the port is set to User (user defined), this field displays the configured power limit. If Power Limit Type is set to Class, this field is blank.

Table 68: PoE Port Configuration Fields (Continued)

Field	Description
Class	If Power Limit Type is set to Class, this field displays the class of the connected device, as learned in LLDP messages. Possible values are Unknown and Class 0 through Class 4. A higher class value indicates that the device requires higher power.
Output Voltage	The voltage being applied to the connected device.
Output Current	The current in milliamps being drawn by the powered device.
Output Power	The power in mWatts being drawn by the connected device.
Temperature	The temperature measured at the PoE port.

Use the buttons to perform the following tasks:

- To configure the settings for one or more interfaces, select each entry to modify and click Edit.
- To apply the same settings to all interfaces, click Edit All.

### 4.12.3 PoE Port Statistics

Use this page to view PoE interface statistical information.

To access the PoE Port Statistics page, click System > PoE > Statistics in the navigation menu.

Figure 77: PoE Port Statistics

Interface	0/1
Overload Counter	0
Short Counter	0
Power Denied Counter	0
MPS Absent Counter	0
Invalid Signature Counter	0

Refresh

Table 69: PoE Port Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data on the page.
Overload Counter	Number of times there has been a power overload. Power overload occurs when a powered device connected to a port tries to draw more power than permissible by the hardware.
Short Counter	Number of times there has been a short circuit condition.
Power Denied Counter	Number of times the powered device has been denied power. Power is denied due to administrative action or shortage of power.
MPS Absent Counter	Number of times power has stopped because the powered device was not detected.
Invalid Signature Counter	Number of times an invalid signature was received. Signature detection is a stage in detecting the presence of a powered device, where a resistance value on the powered device is expected to be found within a particular range.

Click Refresh to redisplay the page with the current data from the switch.

## 4.13 Viewing Device Port Information

The pages in the Port folder allow you to view and monitor the physical port information for the ports available on the switch. The Port folder has links to the following pages.

### 4.13.1 Port Summary

Use the Port Summary page to view the settings for all physical ports on the platform.

To access the Port Summary page, click System > Port > Summary in the navigation menu.

Figure 78: Port Summary

Interface	Interface Index	Type	Admin Mode	Physical Mode	Physical Status	Auto Negotiate Capabilities	STP Mode	LACP Mode	Link Status
1/0/1	1	Normal	Enabled	Auto	Unknown	10f   100f   1000f	Enabled	Enabled	Link Down
1/0/2	2	Normal	Enabled	Auto	1000 Mbps Full Duplex	10f   100f   1000f	Enabled	Enabled	Link Up
1/0/3	3	Normal	Enabled	Auto	Unknown	10f   100f   1000f	Enabled	Enabled	Link Down
1/0/4	4	Normal	Enabled	Auto	1000 Mbps Full Duplex	10f   100f   1000f	Enabled	Enabled	Link Up
1/0/5	5	Normal	Enabled	Auto	Unknown	10f   100f   1000f	Enabled	Enabled	Link Down
1/0/6	6	Normal	Enabled	Auto	Unknown	10f   100f   1000f	Enabled	Enabled	Link Down
1/0/7	7	Normal	Enabled	Auto	Unknown	10f   100f   1000f	Enabled	Enabled	Link Down
1/0/8	8	Normal	Enabled	Auto	Unknown	10f   100f   1000f	Enabled	Enabled	Link Down
1/0/9	9	Normal	Enabled	Auto	Unknown	10f   100f   1000f	Enabled	Enabled	Link Down
1/0/10	10	Normal	Enabled	Auto	Unknown	10f   100f   1000f	Enabled	Enabled	Link Down

Table 70: Port Summary Fields

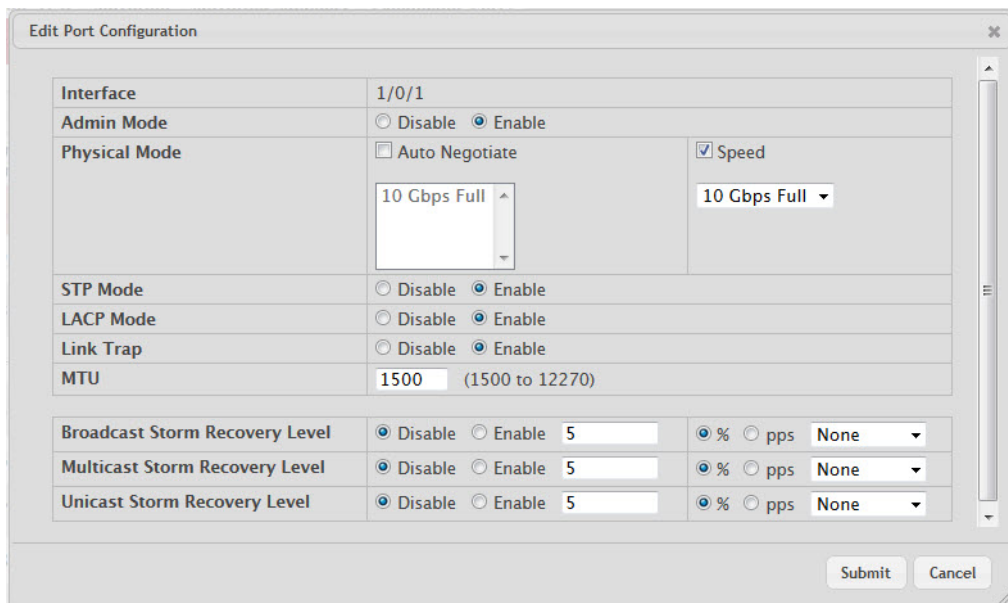
Field	Description
Interface	Identifies the port that the information in the rest of the row is associated with.
Interface Index	The interface index object value assigned by the IF-MIB. This value is used to identify the interface when managing the device by using SNMP.
Type	For most ports this field is blank. Otherwise, the possible values are: <ul style="list-style-type: none"> <li>Normal - The port is a normal port, which means it is not a LAG member or configured for port mirroring.</li> <li>Trunk Member - The port is a member of a LAG.</li> <li>Mirrored - Indicates that the port has been configured as a monitoring port and is the source port in a port mirroring session. For more information about port monitoring and probe ports, see <a href="#">Section 4.13.4: "Mirroring"</a>.</li> <li>Probe - Indicates that the port has been configured as a monitoring port and is the destination port in a port mirroring session. For more information about port monitoring and probe ports, see <a href="#">Section 4.13.4: "Mirroring"</a>.</li> </ul>
Admin Mode	Shows the port control administration state, which can be one of the following: <ul style="list-style-type: none"> <li>Enabled: The port can participate in the network (default).</li> <li>Disabled: The port is administratively down and does not participate in the network.</li> </ul>



**Table 70: Port Summary Fields (Continued)**

Field	Description
Physical Mode	Shows the speed and duplex mode at which the port is configured: <ul style="list-style-type: none"> <li>• Auto: The duplex mode and speed will be set by the auto-negotiation process. The port's maximum capability will be advertised. The option to enable auto-negotiation</li> <li>• &lt;Speed&gt; Half Duplex: The port speeds available from the menu depend on the platform on which the FASTPATH software is running and which port you select. In half-duplex mode, the transmissions are one-way. In other words, the port does not send and receive traffic at the same time.</li> <li>• &lt;Speed&gt; Full Duplex: The port speeds available from the menu depend on the platform on which the FASTPATH software is running and which port you select. In half-duplex mode, the transmissions are two-way. In other words, the port can send and receive traffic at the same time.</li> </ul> The physical mode for a LAG is reported as LAG.
Physical Status	Indicates the port speed and duplex mode at which the port is operating. The physical status for LAGs is not reported. When a port is down, the physical status is unknown.
Auto Negotiate Capabilities	Indicates the list of configured capabilities for a port when Auto Negotiate is on. The Capability status for LAGs is not reported.
STP Mode	The Spanning Tree Protocol (STP) Administrative Mode associated with the port or LAG. STP is a Layer 2 protocol that provides a tree topology for switches on a bridged LAN. STP allows a network to have redundant paths without the risk of network loops, by providing a single path between end stations on a network. The possible values for STP mode are: <ul style="list-style-type: none"> <li>• Enable - Spanning tree is enabled for this port.</li> <li>• Disable - Spanning tree is disabled for this port.</li> </ul>
LACP Mode	Indicates the Link Aggregation Control Protocol administration state. The mode must be enabled for the port to participate in Link Aggregation. This field can have the following values: <ul style="list-style-type: none"> <li>• Enable: Specifies that the port is allowed to participate in a port channel (LAG), which is the default mode.</li> <li>• Disable: Specifies that the port cannot participate in a port channel (LAG).</li> <li>• N/A: For LAG ports.</li> </ul>
Link Status	Indicates whether the Link is up or down.

The following fields can be accessed by selecting a port and clicking Edit:



Auto Negotiate	Select this option to enable auto negotiation on the port.
----------------	--



Table 70: Port Summary Fields (Continued)

Field	Description
Speed	Select this option to manually configure the physical mode for the port (speed and duplex mode).
Link Trap	This object determines whether or not to send a trap when link status changes. The factory default is enabled. <ul style="list-style-type: none"> <li>• Enable: Specifies that the system sends a trap when the link status changes.</li> <li>• Disable: Specifies that the system does not send a trap when the link status changes.</li> </ul>
Maximum Frame Size	The maximum Ethernet frame size the interface supports or is configured to support. The maximum frame size includes the Ethernet header, CRC, and payload.
Broadcast Storm Recovery Mode	Specifies the broadcast storm control threshold for the port. Broadcast storm control limits the amount of broadcast frames accepted and forwarded by the port. If the broadcast traffic on the Ethernet port exceeds the configured threshold, the system blocks (discards) the broadcast traffic. Specifies the broadcast storm recovery action to either Shutdown or Trap for specific interface. If configured to Shutdown, the interface which receives broadcast packets at a rate which is above threshold is diagnostically disabled. The <code>Trap</code> option sends trap messages at approximately every 30 seconds until broadcast storm control recovers.
Multicast Storm Recovery Level	Specifies the multicast storm control threshold for the port. Multicast storm control limits the amount of multicast frames accepted and forwarded by the port. If the multicast traffic on the Ethernet port exceeds the configured threshold, the system blocks (discards) the multicast traffic. Specifies the multicast storm recovery action to either Shutdown or Trap for specific interface. If configured to Shutdown, the interface which receives multicast packets at a rate which is above threshold is diagnostically disabled. The option <code>Trap</code> sends trap messages at approximately every 30 seconds until multicast storm control recovers.
Unicast Storm Recovery Level	Specifies the unicast storm control threshold for the port. Unicast storm control limits the amount of unicast frames accepted and forwarded by the switch. If the unicast traffic on the Ethernet port exceeds the configured threshold, the system blocks (discards) the unicast traffic. Specifies the unicast storm recovery action to either Shutdown or Trap for specific interface. If configured to Shutdown, the interface which receives unicast packets at a rate which is above threshold is diagnostically disabled. The <code>Trap</code> option sends trap messages at approximately every 30 seconds until unicast storm control recovers.

Click Refresh to redisplay the most current information from the router.

### 4.13.2 Port Description

Use the Port Description page to configure a human-readable description of the port.

To access the Port Description page, click System > Port > Description in the navigation menu.

Figure 79: Port Description

Interface	Physical Address	PortList Bit Offset	Interface Index	Port Description
1/0/1	00:10:18:7F:F4:4D	1	1	
1/0/2	00:10:18:7F:F4:4D	2	2	
1/0/3	00:10:18:7F:F4:4D	3	3	
1/0/4	00:10:18:7F:F4:4D	4	4	
1/0/5	00:10:18:7F:F4:4D	5	5	
1/0/6	00:10:18:7F:F4:4D	6	6	
1/0/7	00:10:18:7F:F4:4D	7	7	
1/0/8	00:10:18:7F:F4:4D	8	8	
1/0/9	00:10:18:7F:F4:4D	9	9	
1/0/10	00:10:18:7F:F4:4D	10	10	

Table 71: Port Description Fields

Field	Description
Interface	Select the interface for which data is to be displayed or configured.
Port Description (Input field)	A user-configurable description to help identify the port. To add a description to a port, select the port or LAG from the Interface drop-down menu, type a description in the Port Description field, and then click Submit.
Physical Address	Displays the physical address of the specified interface.
PortList Bit Offset	Displays the bit offset value which corresponds to the port when the MIB object type PortList is used to manage in SNMP.
Interface Index	The interface index object value assigned by the IF-MIB. This value is used to identify the interface when managing the device by using SNMP.
Port Description	Shows the configured port description. By default, the port does not have an associated description.

- If you change a port description, click Submit to apply the change to the system.
- Click Refresh to redisplay the page with the latest information from the router.

### 4.13.3 Port Cable Test

The Port Cable Test feature enables you to determine the cable connection status on a selected port. You can also obtain an estimate of the length of the cable connected to the port, if the PHY on the ports supports this functionality.

#### NOTICE

The Port Cable Test feature is supported only for copper cable. It is not supported for optical fiber cable.

To access the Port Cable Test feature, click System > Port > Cable Test.

The page displays with additional fields when you click Test Cable. The fields that display depend on the cable test results.

Figure 80: Port Cable Test

Table 72: Port Cable Test Fields

Field	Description
Interface	If the test has not been performed, this is the only field that displays. Select the interface to test. After the test has been performed, this field shows the interface that was tested.
Failure Location Distance	The estimated distance from the end of the cable to the failure location. <b>Note:</b> This field displays a value only when the Cable Status is Open or Short; otherwise, this field is blank.
Cable Length	The estimated length of the cable. If the cable length cannot be determined, Unknown is displayed. This field shows the range between the shortest estimated length and the longest estimated length. <b>Note:</b> This field displays a value only when the Cable Status is Normal; otherwise, this field is blank.
Cable Status	This displays the cable status as Normal, Open, or Short. <ul style="list-style-type: none"> <li>• Normal: The cable is working correctly.</li> <li>• Open: The cable is disconnected or there is a faulty connector.</li> <li>• Open and Short: There is an electrical short in the cable.</li> <li>• Cable status Test Failed: The cable status could not be determined. The cable may in fact be working. This field is displayed after you click Test Cable and results are available.</li> </ul>

Select a port and click Test Cable to display its status.

If the port has an active link while the cable test is run, the link can go down for the duration of the test. The test may take several seconds to run.

The command returns a cable length estimate if this feature is supported by the PHY for the current link speed.

### NOTICE

If the link is down and a cable is attached to a 10/100 Ethernet adapter, then the cable status may display as Open or Short because some Ethernet adapters leave unused wire pairs unterminated or grounded.

## 4.13.4 Mirroring

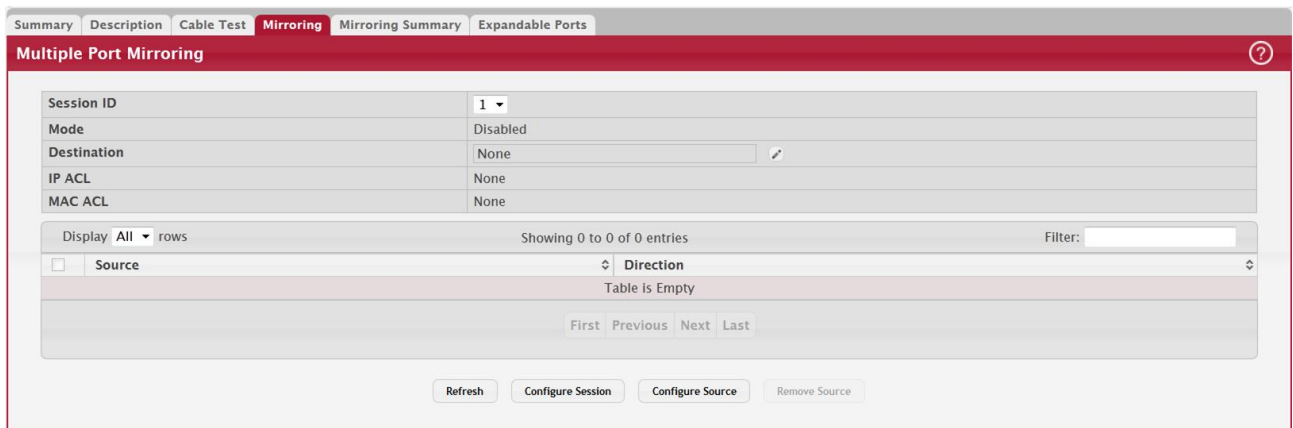
Port mirroring selects the network traffic for analysis by a network analyzer. This is done for specific ports of the switch. As such, many switch ports are configured as source ports and one switch port is configured as a destination port. You have the ability to configure how traffic is mirrored on a source port. Packets that are received on the source port, that are transmitted on a port, or are both received and transmitted, can be mirrored to the destination port.

The packet that is copied to the destination port is in the same format as the original packet on the wire. This means that if the mirror is copying a received packet, the copied packet is VLAN tagged or untagged as it was received on the source port. If the mirror is copying a transmitted packet, the copied packet is VLAN tagged or untagged as it is being transmitted on the source port.

Use the Multiple Port Mirroring page to define port mirroring sessions.

To access the Multiple Port Mirroring page, click System > Port > Mirroring in the navigation menu.

Figure 81: Multiple Port Mirroring



Use the buttons to perform the following tasks:

- To configure the administrative mode for a port mirroring session, click **Configure Session** and configure the desired settings.
- To configure the port mirroring destination, click the **Edit icon** in the **Destination** field and configure the desired settings.
- To configure one or more source ports for the mirroring session and to determine which traffic is mirrored (Tx, Rx, or both), click **Configure Source** and configure the desired settings.
- To remove one or more source ports from the port mirroring session, select the check box associated with each source port to remove and click **Remove Source**.

Table 73: Multiple Port Mirroring Fields

Field	Description
Session ID	The port mirroring session ID. The number of sessions allowed is platform specific.
Mode	The administrative mode for the selected port mirroring session. If the mode is disabled, the configured source is not mirroring traffic to the destination.
Destination	The interface to which traffic is mirrored, which is one of the following: <ul style="list-style-type: none"> <li>• Remote VLAN – Traffic is mirrored to the VLAN on the system that is configured as the RSPAN VLAN. In an RSPAN configuration, the destination should be the Remote VLAN on any device that does not have a port connected to the network traffic analyzer. This destination has to be configured with RSPAN VLAN membership.</li> <li>• Interface – Traffic is mirrored to a physical port on the local device. The interface is the probe port that is connected to a network traffic analyzer.</li> <li>• None – The destination is not configured.</li> </ul> <p><b>Note:</b> This field also identifies the status of the <b>Remove RSPAN Tag</b> option, which can be configured in the <b>Destination Configuration</b> window. When this option is set as <b>False</b>, packets received at the RSPAN destination port are double tagged. When the <b>Remove RSPAN Tag</b> option is <b>True</b>, the RSPAN VLAN ID tag is removed for the mirroring session.</p>
IP ACL	The IP access-list ID or name attached to the port mirroring session.
MAC ACL	The MAC access-list name attached to the port mirroring session.
Source	The ports or VLAN configured to mirror traffic to the destination. You can configure multiple source ports or one source VLAN per session. The source VLAN can also be a remote VLAN.
Direction	The direction of traffic on the source port(s) that is sent to the probe port. Possible values are: <ul style="list-style-type: none"> <li>• Tx and Rx – Both ingress and egress traffic.</li> <li>• Rx – Ingress traffic only.</li> <li>• Tx – Egress traffic only.</li> </ul>

### 4.13.4.1 Configuring a Port Mirroring Session

#### NOTICE

If an interface participates in some VLAN and is a LAG member, this VLAN cannot be assigned as a source VLAN for a Monitor session. At the same time, if an interface participates in some VLAN and this VLAN is assigned as a source VLAN for a Monitor session, the interface can be assigned as a LAG member.

1. From the Multiple Port Mirroring page, select the Session ID of the port mirroring session to configure and click Configure Session to display the Session Configuration page.

2. Configure the following fields:

Table 74: Multiple Port Mirroring—Session Configuration

Field	Description
Session ID	The port mirroring session ID. The number of sessions allowed is platform specific.
Mode	The administrative mode for the selected port mirroring session. If the mode is disabled, the configured source is not mirroring traffic to the destination.

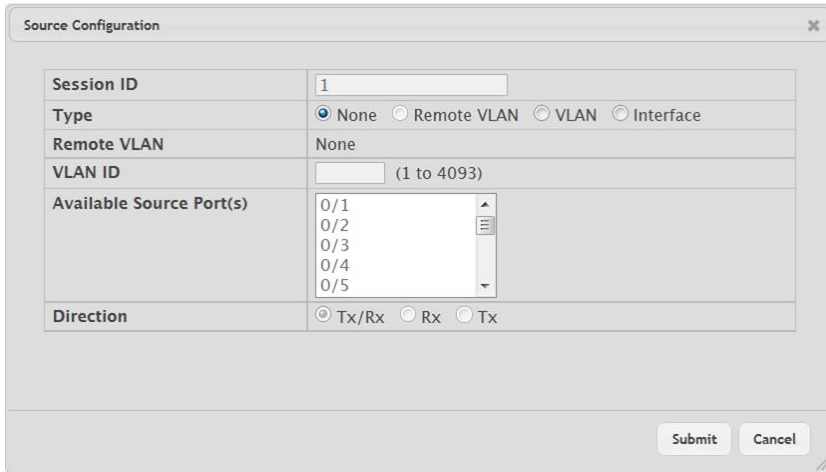
3. Click Submit to apply the changes to the system.

### 4.13.4.2 Configuring a Port Mirroring Source

#### NOTICE

If an interface participates in some VLAN and is a LAG member, this VLAN cannot be assigned as a source VLAN for a Monitor session. At the same time, if an interface participates in some VLAN and this VLAN is assigned as a source VLAN for a Monitor session, the interface can be assigned as a LAG member.

1. From the Multiple Port Mirroring page, click Configure Source to display the Source Configuration page.



2. Configure the following fields:

**Table 75: Multiple Port Mirroring—Source Configuration**

Field	Description
Session ID	The port mirroring session ID. The number of sessions allowed is platform specific.
Type	The type of interface to use as the source: <ul style="list-style-type: none"> <li>• None – The source is not configured.</li> <li>• Remote VLAN – The VLAN configured as the RSPAN VLAN is the source. In an RSPAN configuration, the remote VLAN is the source on the destination device that has a physical port connected to the network traffic analyzer.</li> <li>• VLAN – Traffic to and from a configured VLAN is mirrored. In other words, all the packets sent and received on all the physical ports that are members of the VLAN are mirrored.</li> <li>• Interface – Traffic is mirrored from one or more physical ports on the device.</li> </ul>
Remote VLAN	The VLAN that is configured as the RSPAN VLAN.
VLAN ID	The VLAN to use as the source. Traffic from all physical ports that are members of this VLAN is mirrored. This field is available only when the selected Type is VLAN.
Available Source Ports	The physical port or ports to use as the source. To select multiple ports, CTRL + click each port. This field is available only when the selected Type is Interface.
Direction	The direction of traffic on the source port (or source ports) or VLAN that is sent to the specified destination. A source VLAN mirrors all received and transmitted packets to the destination. Possible values for source ports are: <ul style="list-style-type: none"> <li>Tx and Rx – Both ingress and egress traffic.</li> <li>Rx – Ingress traffic only.</li> <li>Tx – Egress traffic only.</li> </ul>

3. Click Submit to apply the changes to the system.

#### 4.13.4.3 Configuring the Destination Port for a Port Mirroring Session

A port will be removed from a VLAN or LAG when it becomes a destination mirror.

**NOTICE**

1. From the Multiple Port Mirroring page, select the Session ID of the port mirroring session to configure and click the Edit icon in the Destination field.

2. Configure the following fields:

**Table 76: Multiple Port Mirroring—Destination Configuration**

Field	Description
Session ID	The port mirroring session ID. The number of sessions allowed is platform specific.
Type	The type of interface to use as the destination: <ul style="list-style-type: none"> <li>• None – The destination is not configured.</li> <li>• Remote VLAN – Traffic is mirrored to the VLAN on the system that is configured as the RSPAN VLAN. In an RSPAN configuration, the destination should be the Remote VLAN on any device that does not have a port connected to the network traffic analyzer.</li> <li>• Interface – Traffic is mirrored to a physical port on the local device. The interface is the probe port that is connected to a network traffic analyzer.</li> </ul>
Remote VLAN	The VLAN that is configured as the RSPAN VLAN.
Port	The port to which traffic is mirrored. If the Type is Remote VLAN, the selected port is a reflector port. The reflector port is a trunk port that carries the mirrored traffic towards the destination device. If the Type is Interface, the selected port is the probe port that is connected to a network traffic analyzer.
Remove RSPAN Tag	The packets received at RSPAN destination port are double tagged. Enable this option to remove RSPAN VLAN ID tag for mirroring session.

3. Click Submit to apply the changes to the system.

#### 4.13.4.4 Removing or Modifying a Port Mirroring Session

1. From the Port Mirroring page, click Remove Source Port.
2. Select one or more source ports to remove from the session.
3. Use the CTRL key to select multiple ports to remove.
4. Click Remove.

The source ports are removed from the port mirroring session, and the device is updated.

#### 4.13.5 Mirroring Summary

Use the Multiple Port Mirroring Summary page to view the port mirroring summary.

To access the Multiple Port Mirroring Summary page, click System > Port > Mirroring Summary in the navigation menu.

Figure 82: Multiple Port Mirroring Summary

Session ID	Admin Mode	Probe Port	Remove RSPAN Tag	Src VLAN	Mirrored Port	Reflector Port	Src RVLAN	Dst RVLAN	Type	IP ACL	MAC ACL
1	Enabled	1/0/16	True				100				
2	Disabled	1/0/1	False		1/0/3				Tx/Rx		macACL
2	Disabled	1/0/1	False		1/0/4				Tx/Rx		macACL
2	Disabled	1/0/1	False		1/0/5				Tx/Rx		macACL
3	Disabled										
4	Disabled				1/0/5				Tx	10	

Table 77: Multiple Port Mirroring Summary Fields

Field	Description
Session ID	The port mirroring session ID. The number of sessions allowed is platform specific.
Admin Mode	The administrative mode for the selected port mirroring session. If the mode is disabled, the configured source is not mirroring traffic to the destination.
Probe Port	The interface that receives traffic from all configured source ports.
Remove RSPAN Tag	The packets received at an RSPAN destination port are double tagged. If this option is True, the RSPAN VLAN ID tag is removed for the mirroring session.
Src VLAN	The VLAN configured to mirror traffic to the destination. You can configure one source VLAN per session. The source VLAN can also be a remote VLAN.
Mirrored Port	The ports configured to mirror of traffic to the destination. You can configure multiple source ports per session.
Reflector Port	This port carries all the mirrored traffic at source switch.
Src RVLAN	The VLAN configured as the RSPAN VLAN is the source. In an RSPAN configuration, the remote VLAN is the source on the destination device that has a physical port connected to the network traffic analyzer.
Dst RVLAN	Traffic is mirrored to the VLAN on the system that is configured as the RSPAN VLAN. In an RSPAN configuration, the destination should be the Remote VLAN on any device that does not have a port connected to the network traffic analyzer.
Type	The type of traffic on the source port (or source ports) or VLAN that is sent to the specified destination. A source VLAN mirrors all received and transmitted packets to the destination. Possible values for source ports are: <ul style="list-style-type: none"> <li>Tx and Rx – Both ingress and egress traffic.</li> <li>Rx – Ingress traffic only.</li> <li>Tx – Egress traffic only.</li> </ul>
IP ACL	The ID of the IP ACL to apply to traffic from the source. Only traffic that matches the rules in the ACL is mirrored to the destination.
MAC ACL	The ID of the MAC ACL to apply to traffic from the source. Only traffic that matches the rules in the ACL is mirrored to the destination.

Click Refresh to redisplay the page with the latest information from the router.



### 4.13.6 Expandable Ports

Use this page to view and configure settings on expandable ports present on the device. The Expandable ports feature is used to support physical ports capable of operating at multiple port speeds. It is implemented in two flavors due to platform restrictions. In the Dynamic version, the operational mode of the expandable port changes immediately after user configuration. In the Static version, any user configuration is stored as part of persistent data and is applied only on the next reboot.

To change the settings for one or more ports, select each entry to modify and click Edit.

To access the Expandable Port Summary page, click System > Port > Expandable Ports in the navigation menu.

Figure 83: Expandable Port Summary

Interface	Physical Type	Expansion Interfaces	Configured Mode	Operational Mode
1/0/81	40G EP	1/0/93, 1/0/94, 1/0/95, 1/0/96	1x40G	1x40G
1/0/82	40G EP	1/0/97, 1/0/98, 1/0/99, 1/0/100	1x40G	1x40G
1/0/83	40G EP	1/0/101, 1/0/102, 1/0/103, 1/0/104	1x40G	1x40G
1/0/84	40G EP	1/0/105, 1/0/106, 1/0/107, 1/0/108	1x40G	1x40G
1/0/85	100G EP	1/0/109, 1/0/110, 1/0/111, 1/0/112	1x100G	1x100G
1/0/86	100G EP	1/0/113, 1/0/114, 1/0/115, 1/0/116	1x100G	1x100G
1/0/87	100G EP	1/0/117, 1/0/118, 1/0/119, 1/0/120	1x100G	1x100G
1/0/88	100G EP	1/0/121, 1/0/122, 1/0/123, 1/0/124	1x100G	1x100G

Table 78: Expandable Port Fields

Field	Description
Interface	The interface operating as a 1x40G or 1x100G port and associated with the rest of the data in the row.
Expansion Interfaces	Interfaces operating as independent 4x10G or 4x25G ports.
Configured Mode	The port mode configured as either a 1x40G port or independent 4x10G ports for 40G EP interface and 1x100G port or independent 4x25G or 2x50G ports for 100G EP interface. In the static version of the expandable ports feature, since any user configuration is effective only after a switch reboot, the configured mode may be different than the operational mode on the interface.
Operational Mode	The operational port mode as either a 1x40G port or independent 4x10G ports for 40G EP interface and 1x100G port or independent 4x25G ports for 100G EP interface. In the dynamic version of the expandable ports feature, the configured and operational modes are same on the interface.
After you click Edit, the Edit Expandable Port Configuration window opens and allows you to configure the port mode. The following information describes the field in this window.	
40G Interfaces	The interface list operating as a 1x40G port.
100G Interfaces	The interface list operating as a 1x100G port.
Port Mode	The port mode, which is one of the following for 40G EP interface: <ul style="list-style-type: none"> <li>• 1x40G – A single 40G port using 4 lanes.</li> <li>• 4x10G – Four 10G ports, each on a separate lane.</li> </ul> And one of the following for 100G EP interface: <ul style="list-style-type: none"> <li>• 1x100G – A single 100G port using 4 lanes.</li> <li>• 4x25G – Four 25G ports, each on a separate lane.</li> </ul>

### 4.13.7 Green Mode Statistics

For platforms that include Green Energy features, the Green Mode Statistics page shows information about the amount of energy saved for each port. This page also allows you to enable or disable the green mode features that the switch supports. The green mode features can be controlled on a port-by-port basis.

To access the Green Mode Statistics page, click System > Advanced Configuration > Green Ethernet > Statistics.

Figure 84: Green Ethernet Statistics

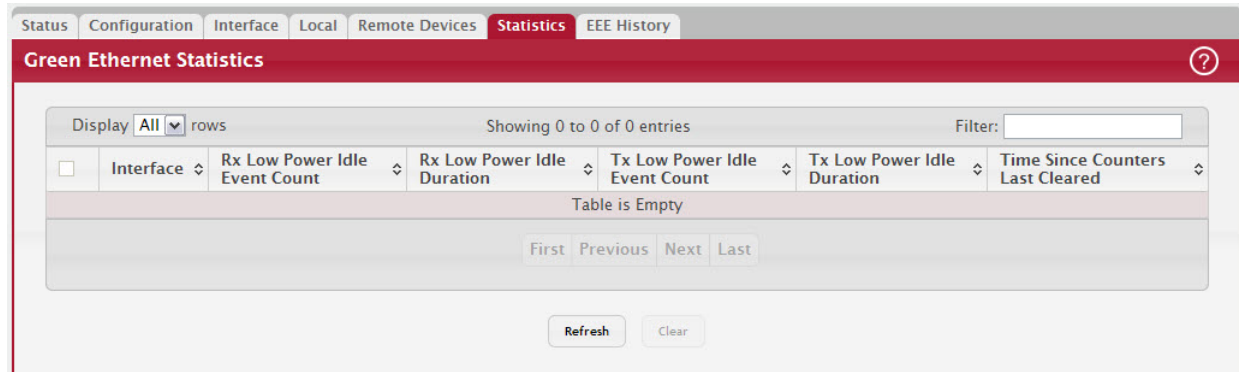


Table 79: Green Ethernet Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. The table includes all interfaces that are enabled for EEE.
Rx Low Power Idle Event Count	The number of times the local interface has entered a low-power idle state.
Rx Low Power Idle Duration	The amount of time (in 10 microsecond increments) the local interface has spent in a low-power idle state.
Tx Low Power Idle Event Count	The number of times the link partner has entered a low-power idle state.
Tx Low Power Idle Duration	The amount of time (in 10 microsecond increments) the link partner has spent in a low-power idle state.
Time Since Counters Last Cleared	The amount of time since the statistics on this page were reset to zero.

#### Command Buttons

This page has the following command buttons:

- Clear—Resets all Green Ethernet statistics counters on this page to 0.
- Refresh—Refresh the data on the screen with the present state of the data in the switch.

### 4.13.8 Green Ethernet EEE Interface History Table

Use the Green Mode EEE History page to set the sampling interval for EEE LPI data and to specify the number of samples to keep. From this page, you can also view per-port EEE LPI data.

To access the Green Ethernet EEE Interface History Table page, click System > Advanced Configuration > Green Ethernet > EEE History.

Figure 85: Green Ethernet EEE Interface History Table

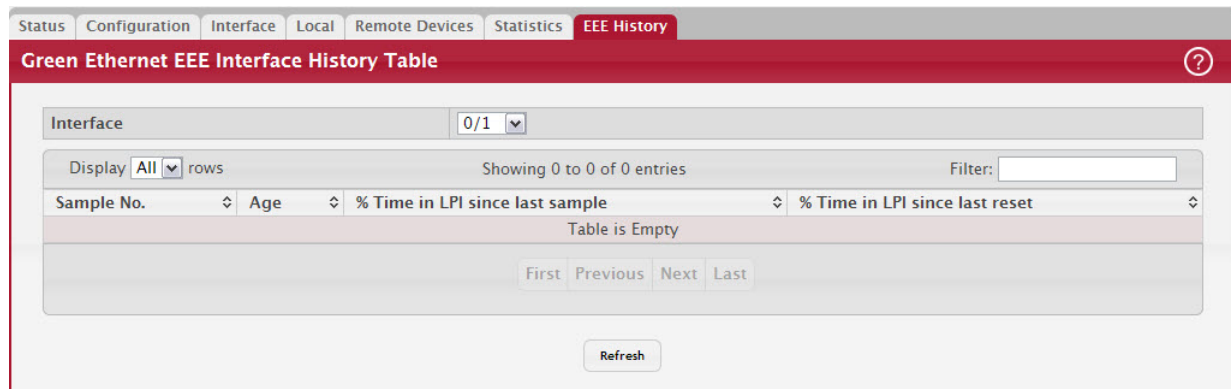


Table 80: Green Ethernet EEE Interface History Table Fields

Field	Description
Interface	Select the interface with the green mode information to view or configure.
Sample No.	A unique number that identifies the sample.
Age	The amount of time that has passed since the sample was recorded.
% Time in LPI since last sample	The percentage of time the interface has spent in LPI mode since the last sample was taken.
% Time in LPI since last reset	The percentage of time the interface has spent in LPI mode since the last time the EEE statistics were cleared.

## 4.14 Configuring sFlow

sFlow is the standard for monitoring high-speed switched and routed networks. sFlow technology is built into network equipment and gives complete visibility into network activity, enabling effective management and control of network resources.

The sFlow monitoring system consists of an sFlow Agent (embedded in a switch or router or in a standalone probe) and a central sFlow Collector. The sFlow Agent uses sampling technology to capture traffic statistics from the device it is monitoring. sFlow datagrams are used to immediately forward the sampled traffic statistics to an sFlow Collector for analysis.

The sFlow Agent uses two forms of sampling: statistical packet-based sampling of switched or routed Packet Flows, and time-based sampling of counters.

### 4.14.1 sFlow Agent Summary

Packet Flow Sampling and Counter Sampling are performed by sFlow Instances associated with individual Data Sources within the sFlow Agent. Packet Flow Sampling and Counter Sampling are designed as part of an integrated system. Both types of samples are combined in sFlow datagrams. Packet Flow Sampling will cause a steady, but random, stream of sFlow datagrams to be sent to the sFlow Collector. Counter samples may be taken opportunistically to fill these datagrams.

To perform Packet Flow Sampling, an sFlow Sampler Instance is configured with a Sampling Rate. The Packet Flow sampling process results in the generation of Packet Flow Records. To perform Counter Sampling, the sFlow Poller Instance is configured with a Polling Interval. The Counter Sampling process results in the generation of Counter Records. The sFlow Agent collects Counter Records and Packet Flow Records and sends them in the form of sFlow datagrams to sFlow Collectors.

To access the sFlow Agent Summary page, click System > Advanced Configuration > sFlow > Agent in the navigation menu.

Figure 86: sFlow Agent Summary




Table 81: sFlow Agent Summary Fields

Field	Description
Version	Uniquely identifies the version and implementation of this MIB. The version string must have the following structure: MIB Version; Organization; Software Revision where: <ul style="list-style-type: none"> <li>• MIB Version: '1.3', the version of this MIB.</li> <li>• Organization: Broadcom Corp.</li> <li>• Revision: 1.0</li> </ul>
Agent Address	The IP address associated with this agent.

Use the Refresh button to refresh the page with the most current data from the switch.

#### 4.14.2 sFlow Receiver Configuration

Use the sFlow Receiver Configuration page to configure the sFlow Receiver.

To access the sFlow Receiver Configuration page, click System > Advanced Configuration > sFlow > Receiver in the navigation menu.

Figure 87: sFlow Receiver Configuration

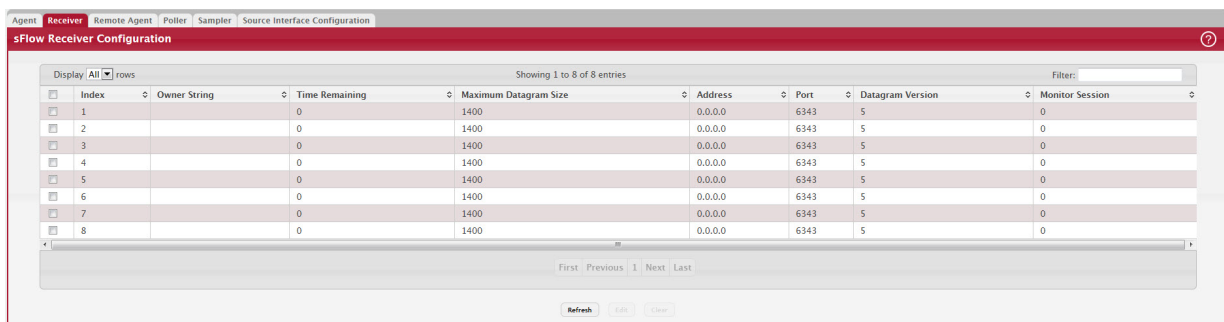


Table 82: sFlow Receiver Configuration Fields

Field	Description
Index	Selects the receiver for which data is to be displayed or configured. The allowed range is 1 to 8.
Owner String	The entity making use of this sFlowRcvrTable entry. The empty string indicates that the entry is currently unclaimed and the receiver configuration is reset to the default values. An entry wishing to claim an sFlowRcvrTable entry must ensure that the entry is unclaimed before trying to claim it. The entry is claimed by setting the owner string. The entry must be claimed before any changes can be made to other sampler objects.

**Table 82: sFlow Receiver Configuration Fields**

Field	Description
Time Remaining	The time (in seconds) remaining before the sampler is released and stops sampling. A value of 0 essentially means the receiver is not configured. When configuring the sFlow receiver settings, you must select the Timeout Mode option before you can configure a Timeout Value.
Maximum Datagram Size	The maximum number of data bytes that can be sent in a single sample datagram. The manager should set this value to avoid fragmentation of the sFlow datagrams. The default value is 1400. The allowed range is 200 to 9116.
Address	The IP address of the sFlow collector. If set to 0.0.0.0 no sFlow datagrams will be sent.
Port	The destination port for sFlow datagrams. The allowed range is 1 to 65535.
Datagram Version	The version of sFlow datagrams that should be sent.
Monitor Session	Monitor session to enable sFlow hardware feature.

- Use the Submit button to sent updated data to the switch and cause the changes to take effect on the switch.
- Use the Refresh button to refresh the page with the most current data from the switch.

Use the Edit button to configure the monitor session for a specific receiver (only for IPv4). After successful configuration, the sFlow packet processing will be done in hardware.

**Figure 88: Edit Receiver Configuration**

Edit Receiver Configuration	
Index	3
Owner String	<input type="text"/> (Max 127 characters)
Timeout Mode	<input checked="" type="checkbox"/>
Timeout Value (Seconds)	0 (0 to 2147483647)
Maximum Datagram Size	1400 (200 to 9316)
Host IP Address	0.0.0.0 (x.x.x.x or x:x:x:x:x:x:x:x)
Port	6343 (1 to 65535)
Datagram Version	5
Monitor Session	0

Use the Submit button to sent updated data to the switch and cause the changes to take effect on the switch.

### 4.14.3 sFlow Poller Configuration

The sFlow agent collects time-based sampling of network interface statistics and sends them to the configured sFlow receivers. A data source configured to collect counter samples is called a poller.

#### 4.14.3.1 Counter Sampling

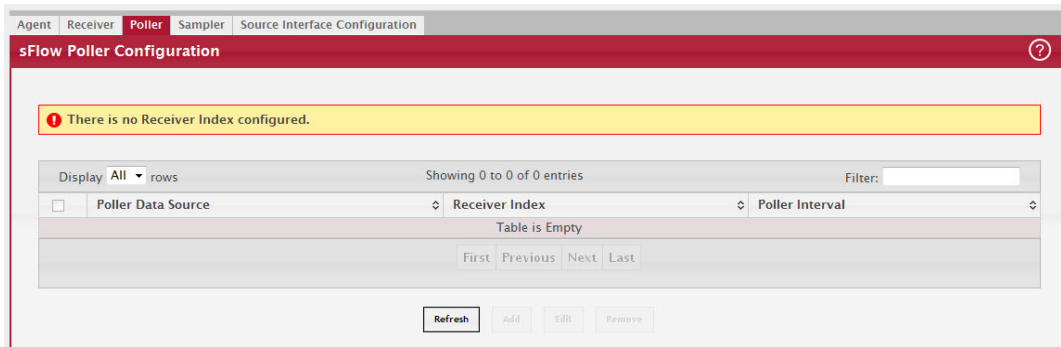
The primary objective of Counter Sampling is to efficiently, periodically export counters associated with Data Sources. A maximum Sampling Interval is assigned to each sFlow instance associated with a Data Source.

Counter Sampling is accomplished as follows:

The sFlow Agent keeps a list of counter sources being sampled. When a Packet Flow Sample is generated, the sFlow Agent examines the list and adds counters to the sample datagram, least recently sampled first. Counters are only added to the datagram if the sources are within a short period, i.e. five seconds, of failing to meet the required Sampling Interval. Periodically, i.e. every second, the sFlow Agent examines the list of counter sources and sends any counters that need to be sent to meet the sampling interval requirement.

To access the sFlow Poller Configuration page, click System > Advanced Configuration > sFlow > Poller in the navigation menu.

Figure 89: sFlow Poller Configuration



Use the buttons to perform the following tasks:

- To add an sFlow poller instance, click Add and complete the required information.
- To edit an existing sFlow poller instance, select the appropriate check box or click the row to select the sFlow poller instance and click Edit. Modify the sFlow poller configuration information as needed.
- To delete an sFlow poller instance, select one or more table entries and click Remove.

Figure 90: Add Poller

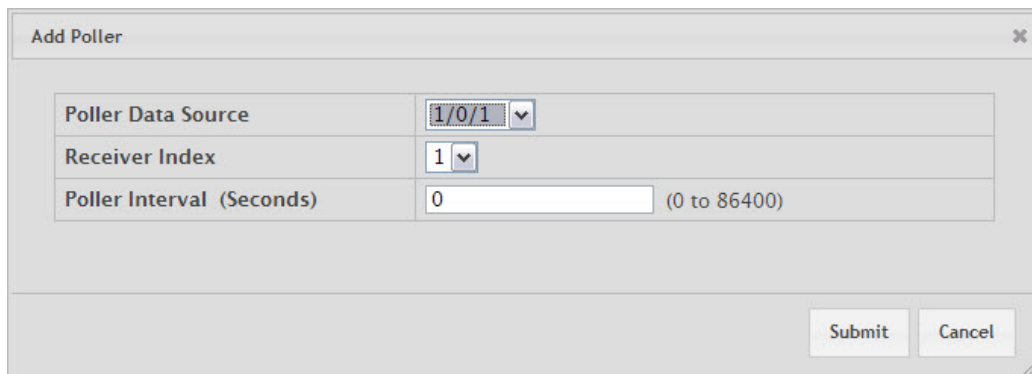


Table 83: Add Poller Fields

Field	Description
Poller Data-Source	The sFlow Sampler Datasource for this flow sampler. This Agent will support Physical ports only.
Receiver Index	The sFlowReceiver for this sFlow Counter Poller. If set to zero, the poller configuration is set to the default and the poller is deleted. Only active receivers can be set. If a receiver expires, then all pollers associated with the receiver will also expire. The allowed range is 1 to 8.
Poller Interval	The maximum number of seconds between successive samples of the counters associated with this data source

Click Refresh to refresh the page with the most current data from the switch.

### 4.14.4 sFlow Sampler Configuration

The sFlow Agent collects a statistical packet-based sampling of the switched flows and sends them to the configured receivers. A data source configured to collect flow samples is called a sampler.

#### 4.14.4.1 Packet Flow Sampling

The Packet Flow Sampling mechanism carried out by each sFlow instance ensures that any packet observed at a Data Source has an equal chance of being sampled, irrespective of the Packet Flow(s) to which it belongs.

Packet Flow Sampling is accomplished as follows:

- When a packet arrives on an interface, the Network Device makes a filtering decision to determine whether the packet should be dropped.
- If the packet is not filtered (dropped), a destination interface is assigned by the switching/routing function.
- At this point, a decision is made on whether or not to sample the packet. The mechanism involves a counter that is decremented with each packet. When the counter reaches zero, a sample is taken. When a sample is taken, the counter that indicates how many packets to skip before taking the next sample is reset. The value of the counter is set to a random integer where the sequence of random integers used over time is the Sampling Rate.

To access the sFlow Sampler Configuration page, click System > Advanced Configuration > sFlow > Sampler in the navigation menu.

Figure 91: sFlow Sampler Configuration



Table 84: sFlow Sampler Configuration Fields

Field	Description
Sampler Data Source	The sFlowDataSource for this sFlow sampler. The sFlow agent supports physical ports as sFlow data sources.
Receiver Index	The sFlowReceiver for this sFlow sampler. The specified Receiver Index must be associated with an active sFlow receiver. If a receiver expires, all samplers associated with the receiver will also expire.
Remote Agent Index	The remote agent index.
Ingress Sampling Rate	sFlow instance packet Sampling Rate for Ingress sampling.
Egress Sampling Rate	sFlow instance egress packet Sampling Rate.
Flow-based Sampling Rate	sFlow instance flow based packet Sampling Rate.
Maximum Header Size	The maximum number of bytes that should be copied from a sampled packet.
IP ACL	The ID of the IP ACL to apply to traffic from the sampler.
IP MAC	The ID of the MAC ACL to apply to traffic from the sampler.

Click Refresh to refresh the page with the most current data from the switch.

Use the buttons to perform the following tasks:

- To add an sFlow sampler instance, click Add and complete the required information.
- To edit an existing sFlow sampler instance, select the appropriate check box or click the row to select the sFlow sampler instance and click Edit. Modify the sFlow sampler configuration information as needed.
- To delete an sFlow sampler instance, select one or more table entries and click Remove.

The Add Sampler page lets you configure the sampling rate for ingress/egress/flow based sampling. After successful configuration, the sFlow packet sampling is performed based on sampling rate.

**Figure 92: Add Sampler**

Sampler Data Source	1/0/1	
Receiver Index	1	
Remote Agent Index	0	
Ingress Sampling Rate	<input type="text"/>	(1024 to 65536)
Egress Sampling Rate	<input type="text"/>	(1024 to 65536)
Flow Based Sampling Rate	<input type="text"/>	(1024 to 65536)
Maximum Header Size	128	(20 to 256, 128 = Default)
IP ACL	0	
IP MAC	0	

**Table 85: Add Sampler Fields**

Field	Description
Sampler Data Source	The sFlowDataSource for this sFlow sampler. The sFlow agent supports physical ports as sFlow data sources.
Receiver Index	The sFlow Receiver for this sFlow sampler. If set to zero, no packets will be sampled. Only active receivers can be set. If a receiver expires, then all samplers associated with the receiver will also expire. The allowed range is 1 to 8.
Ingress Sampling Rate	sFlow instance packet Sampling Rate for Ingress sampling.
Egress Sampling Rate	sFlow instance egress packet Sampling Rate.
Flow Based Sampling Rate	sFlow instance flow based packet Sampling Rate.
Maximum Header Size	The maximum number of bytes that should be copied from a sampled packet. The allowed range is 20 to 256.

#### 4.14.5 sFlow Source Interface Configuration

Use this page to specify the physical or logical interface to use as the sFlow client source interface. When an IP address is configured on the source interface, this address is used for all sFlow communications between the local sFlow client and the remote sFlow server. The IP address of the designated source interface is used in the IP header of sFlow management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access the sFlow Source Interface Configuration page, click System > Advanced Configuration > sFlow > Source Interface Configuration in the navigation menu.



Figure 93: sFlow Source Interface Configuration

Table 86: sFlow Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> <li>• None – The primary IP address of the originating (outbound) interface is used as the source address.</li> <li>• Interface – The primary IP address of a physical port is used as the source address.</li> <li>• Loopback – The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up.</li> <li>• VLAN – The primary IP address of a VLAN routing interface is used as the source address.</li> </ul>
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Tunnel ID	The primary IP address of a tunnel interface is used as the source address.

If you make any changes to the page, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

## 4.15 Defining SNMP Parameters

Simple Network Management Protocol (SNMP) provides a method for managing network devices. The device supports SNMP version 1, SNMP version 2, and SNMP version 3.

### 4.15.1 SNMP v1 and v2

The SNMP agent maintains a list of variables, which are used to manage the device. The variables are defined in the Management Information Base (MIB). The MIB presents the variables controlled by the agent. The SNMP agent defines the MIB specification format, as well as the format used to access the information over the network. Access rights to the SNMP agent are controlled by access strings.

## 4.15.2 SNMP v3

SNMP v3 also applies access control and a new traps mechanism to SNMPv1 and SNMPv2 PDUs. In addition, the User Security Model (USM) is defined for SNMPv3 and includes:

- **Authentication:** Provides data integrity and data origin authentication.
- **Privacy:** Protects against disclosure of message content. Cipher-Block-Chaining (CBC) is used for encryption. Either authentication is enabled on an SNMP message, or both authentication and privacy are enabled on an SNMP message. However privacy cannot be enabled without authentication.
- **Timeliness:** Protects against message delay or message redundancy. The SNMP agent compares incoming message to the message time information.
- **Key Management:** Defines key generation, key updates, and key use.

The device supports SNMP notification filters based on Object IDs (OID). OIDs are used by the system to manage device features. SNMP v3 supports the following features:

- Security
- Feature Access Control
- Traps

Authentication or Privacy Keys are modified in the SNMPv3 User Security Model (USM).

Use the SNMP page to define SNMP parameters. To display the SNMP page, click System > Advanced Configuration > SNMP in the navigation menu.

## 4.15.3 SNMP Community Configuration

Access rights are managed by defining communities on the SNMPv1, 2 Community page. When the community names are changed, access rights are also changed. SNMP Communities are defined only for SNMP v1 and SNMP v2.

---

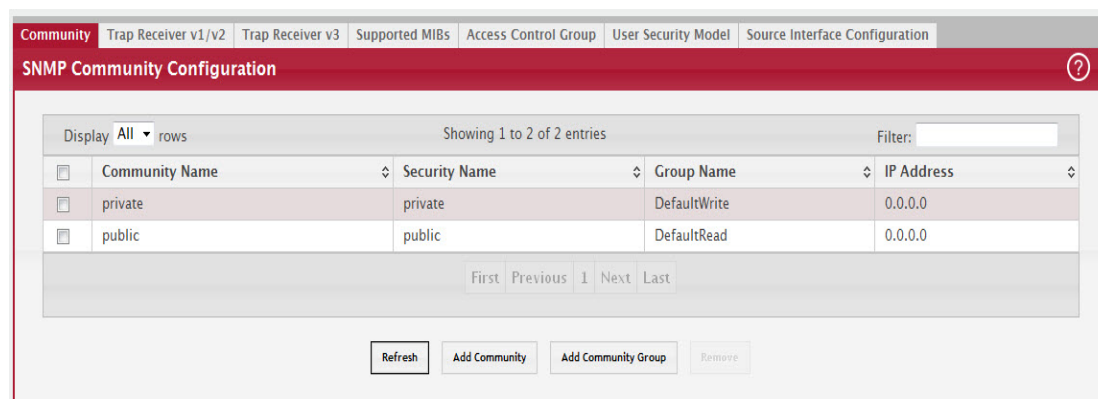
Starting with FASTPATH 8.2, no SNMP communities exist by default.

### NOTICE

Use the SNMP Community Configuration page to enable SNMP and Authentication notifications.

To display the SNMP Community Configuration page, click System > Advanced Configuration > SNMP > Community in the navigation menu.

**Figure 94:** SNMP Community Configuration



Use the buttons to perform the following tasks:

- To add a community, click Add and configure the desired settings.
- To change information for an existing community, select the check box associated with the entry and click Edit.
- To delete a configured community from the list, select the check box associated with each entry to delete and click Remove.

Figure 95: Add New Community

The screenshot shows a dialog box titled "Add New Community" with a close button (X) in the top right corner. The dialog contains four rows of configuration fields:

- Community Name:** A text input field with a placeholder and the constraint "(1 to 20 characters)".
- Community Access:** Three radio button options: "DefaultRead" (selected), "DefaultWrite", and "DefaultSuper".
- Community View:** A text input field with a placeholder and the constraint "(1 to 30 characters)".
- IP Address:** A text input field with the value "0.0.0.0" and the constraint "(x.x.x.x)".

At the bottom right of the dialog, there are two buttons: "Submit" and "Cancel".

Table 87: Community Configuration Fields

Field	Description
Community Name	Contains the user-defined community strings that act as a password and are used to authenticate the SNMP management station to the device. A community string can contain a maximum of 20 characters. By default, the options available in the menu are as follows: <ul style="list-style-type: none"> <li>public: This SNMP community has Read Only privileges and its status set to enable.</li> <li>private: This SNMP community has Read/Write privileges and its status set to enable.</li> </ul>
Community Access	Specifies the access control policy for the community.
Client IP Address	Taken together, the Client IP Address and Client IP Mask denote a range of IP addresses from which SNMP clients may use that community to access this device. If either (IP Address or IP Mask) value is 0.0.0.0, access is allowed from any IP address. Otherwise, every client's IP address is ANDed with the mask, as is the Client IP Address, and, if the values are equal, access is allowed. For example, if the Client IP Address and Client IP Mask parameters are 192.168.1.0/255.255.255.0, then any client whose IP address is 192.168.1.0 through 192.168.1.255 (inclusive) will be allowed access. To allow access from only one station, use a Client IP Mask value of 255.255.255.255, and use that machine's IP address for Client IP Address.
Client IP Mask	Along with the Client IP Address, the Client IP Mask denotes a range of IP addresses from which SNMP clients may use that community to access this device.
Access Mode	Specify the access level for this community: <ul style="list-style-type: none"> <li>Read-Only: The Community has read only access to the MIB objects configured in the view.</li> <li>Read-Write: The Community has read/modify access to the MIB objects configured in the view.</li> </ul>
Status	Specify the status of this community: <ul style="list-style-type: none"> <li>Enable: The community is enabled, and the Community Name must be unique among all valid Community Names or the set request will be rejected.</li> <li>Disable: The Community is disabled and the Community Name becomes invalid.</li> </ul>

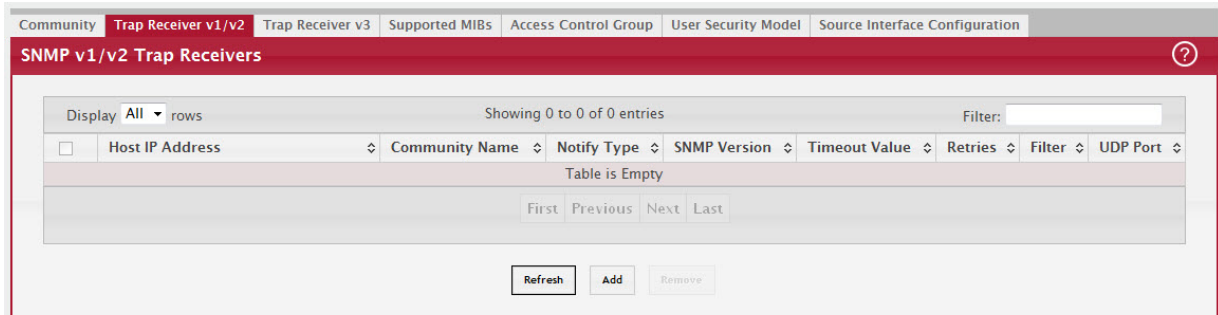
- If you make any changes to the page, click Submit to apply the changes to the system. If you create a new Community, it is added to the table below the Submit button.
- Click Remove to delete the selected SNMP Community.

### 4.15.4 Trap Receiver v1/v2 Configuration

Use the Trap Receiver v1/v2 Configuration page to configure settings for each SNMPv1 or SNMPv2 management host that will receive notifications about traps generated by the device. The SNMP management host is also known as the SNMP trap receiver.

To access the Trap Receiver v1/v3 Configuration page, click System > Advanced Configuration > SNMP > Trap Receiver V1/V2 from the navigation menu.

Figure 96: SNMP v1/v2 Trap Receivers



Use the buttons to perform the following tasks:

- To add an SNMP trap receiver and configure its settings, click Add and complete the required information.
- To delete one or more SNMP trap receivers from the list, select each entry to delete and click Remove.

Figure 97: Add SNMP v1/v2 Host

Table 88: Add SNMP v1/v2 Host Fields

Field	Description
Host IP Address	The IP address of the SNMP management host that will receive traps generated by the device.
Community Name	The name of the SNMP community that includes the SNMP management host and the SNMP agent on the device.
Notify Type	The type of SNMP notification to send the SNMP management host: <ul style="list-style-type: none"> <li>• Inform – An SNMP message that notifies the host when a certain event has occurred on the device. The message is acknowledged by the SNMP management host. This type of notification is not available for SNMPv1.</li> <li>• Trap – An SNMP message that notifies the host when a certain event has occurred on the device. The message is not acknowledged by the SNMP management host.</li> </ul>

Table 88: Add SNMP v1/v2 Host Fields (Continued)

Field	Description
SNMP Version	The version of SNMP to use, which is either SNMPv1 or SNMPv2.
Timeout Value	The number of seconds to wait for an acknowledgment from the SNMP management host before resending an inform message.
Retries	The number of times to resend an inform message that is not acknowledged by the SNMP management host.
Filter	The name of the filter for the SNMP management host. The filter is configured by using the CLI and defines which MIB objects to include or exclude from the view. This field is optional.
UDP Port	The UDP port on the SNMP management host that will receive the SNMP notifications. If no value is specified when configuring a receiver, the default UDP port value is used.

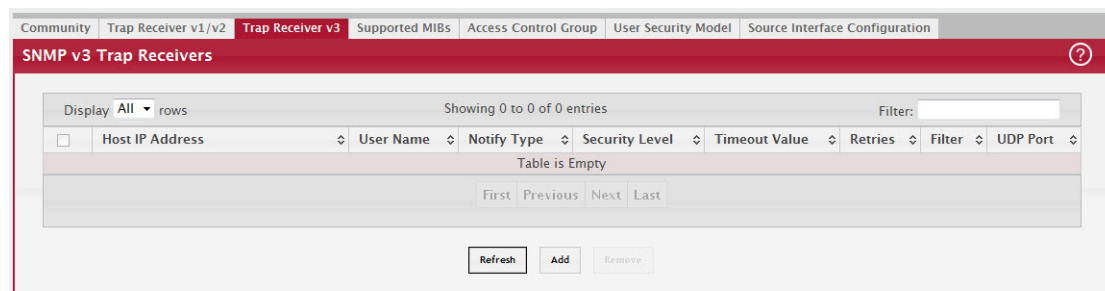
If you make any changes to the page, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

#### 4.15.5 Trap Receiver v3 Configuration

Use the Trap Receiver V3 Configuration v3 page to configure settings for each SNMPv3 management host that will receive notifications about traps generated by the device. The SNMP management host is also known as the SNMP trap receiver.

To access the Trap Receiver v3 Configuration page, click System > Advanced Configuration > SNMP > Trap Receiver V3 from the navigation menu.

Figure 98: SNMP v3 Trap Receivers



Use the buttons to perform the following tasks:

- To add an SNMP trap receiver and configure its settings, click Add and complete the required information.
- To delete one or more SNMP trap receivers from the list, select each entry to delete and click Remove.

Figure 99: Add SNMP v3 Host

Table 89: Add SNMP v3 Host Fields

Field	Description
Host IP Address	The IP address of the SNMP management host that will receive traps generated by the device.
User Name	The name of the SNMP user that is authorized to receive the SNMP notification.
Notify Type	The type of SNMP notification to send the SNMP management host: <ul style="list-style-type: none"> <li>• Inform – An SNMP message that notifies the host when a certain event has occurred on the device. The message is acknowledged by the SNMP management host. This type of notification is not available for SNMPv1.</li> <li>• Trap – An SNMP message that notifies the host when a certain event has occurred on the device. The message is not acknowledged by the SNMP management host.</li> </ul>
Security Level	The security level associated with the SNMP user, which is one of the following: <ul style="list-style-type: none"> <li>• No Auth No Priv – No authentication and no data encryption (no security).</li> <li>• Auth No Priv – Authentication, but no data encryption. With this security level, users send SNMP messages that use an MD5 key/password for authentication, but not a DES key/password for encryption.</li> <li>• Auth Priv – Authentication and data encryption. With this security level, users send an MD5 key/password for authentication and a DES key/password for encryption.</li> </ul>
Retries	The number of times to resend an inform message that is not acknowledged by the SNMP management host.
Timeout Value	The number of seconds to wait for an acknowledgment from the SNMP management host before resending an inform message.
Filter	The name of the filter for the SNMP management host. The filter is configured by using the CLI and defines which MIB objects to include or exclude from the view. This field is optional.
UDP Port	The UDP port on the SNMP management host that will receive the SNMP notifications. If no value is specified when configuring a receiver, the default UDP port value is used.

If you make any changes to the page, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

#### 4.15.6 SNMP Supported MIBs

The SNMP Supported MIBs page lists the MIBs that the system currently supports.

To access the SNMP Supported MIBs page, click System > Advanced Configuration > SNMP > Supported MIBs in the navigation menu.

Figure 100: SNMP Supported MIBs

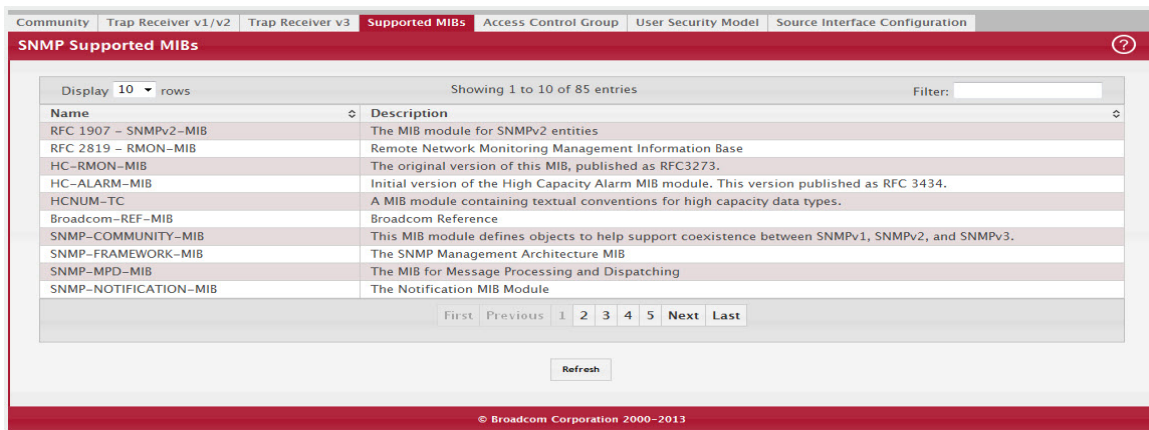


Table 90: SNMP Supported MIBs Fields

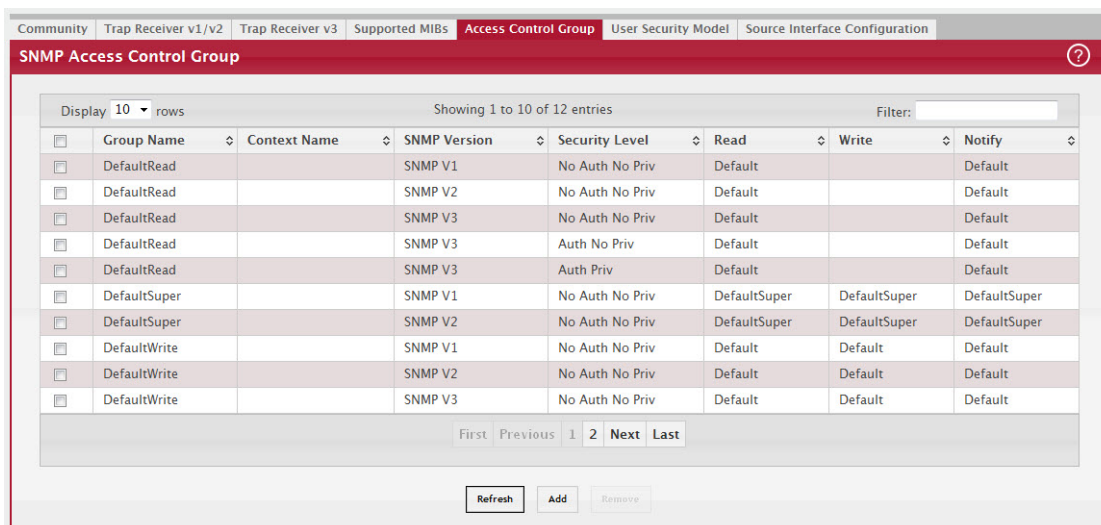
Field	Description
Name	The RFC number if applicable and the name of the MIB.
Description	The RFC title or MIB description.

### 4.15.7 SNMP Access Control Group

Use this page to configure SNMP access control groups. These SNMP groups allow network managers to assign different levels of authorization and access rights to specific device features and their attributes. The SNMP group can be referenced by the SNMP community to provide security and context for agents receiving requests and initiating traps as well as for management systems and their tasks. An SNMP agent will not respond to a request from a management system outside of its configured group, but an agent can be a member of multiple groups at the same time to allow communication with SNMP managers from different groups. Several default SNMP groups are preconfigured on the system.

To access the SNMP Access Control Group page, click System > Advanced Configuration > SNMP > Access Control Group in the navigation menu. A portion of the web screen is shown [Figure 101: "SNMP Access Control Group," on page 129.](#)

Figure 101: SNMP Access Control Group



Use the buttons to perform the following tasks:

- To add an SNMP group, click Add and specify the desired setting.
- To remove one or more SNMP groups, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

**Figure 102: Add New Access Control Group**

**Table 91: Add New Access Control Group Fields**

Field	Description
Group Name	The name that identifies the SNMP group.
Context Name	The SNMP context associated with the SNMP group and its views. A user or a management application specifies the context name to get the performance information from the MIB objects associated with that context name. The Context EngineID identifies the SNMP entity that should process the request (the physical router), and the Context Name tells the agent in which context it should search for the objects requested by the user or the management application.
SNMP Version	The SNMP version associated with the group.
Security Level	The security level associated with the group, which is one of the following: <ul style="list-style-type: none"> <li>• No Auth No Priv – No authentication and no data encryption (no security). This is the only Security Level available for SNMPv1 and SNMPv2 groups.</li> <li>• Auth No Priv – Authentication, but no data encryption. With this security level, users send SNMP messages that use an MD5 key/password for authentication, but not a DES key/password for encryption.</li> <li>• Auth Priv – Authentication and data encryption. With this security level, users send an MD5 key/password for authentication and a DES key/password for encryption.</li> </ul>
Read	The level of read access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that restricts management access to viewing the contents of the agent.
Write	The level of write access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that permits management read-write access to the contents of the agent but not to the community.
Notify	The level of notify access rights for the group. The menu includes the available SNMP views. When adding a group, select the check box to allow the field to be configured, then select the desired view that permits sending SNMP traps or informs.



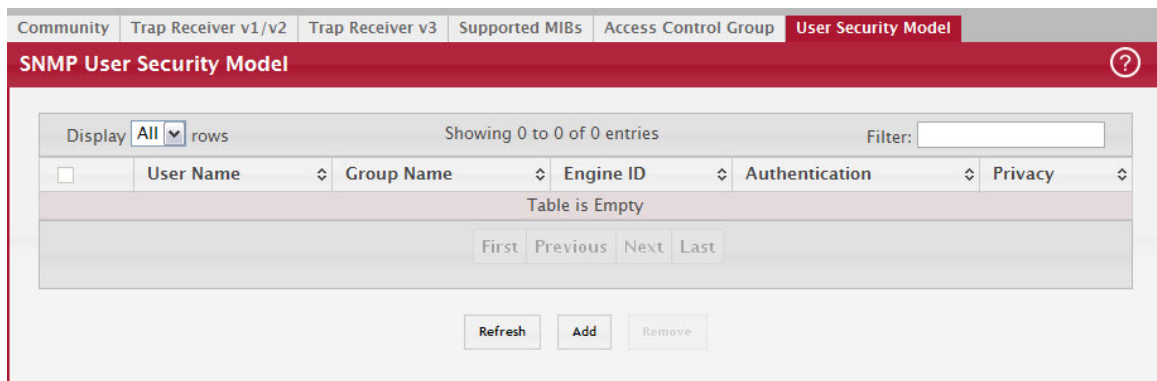
If you make any changes to the page, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

### 4.15.8 SNMP User Security Model

The SNMP User Security Model page provides the capability to configure the SNMP V3 user accounts.

To access the SNMP User Security Model page, click System > Advanced Configuration > SNMP > User Security Model in the navigation menu.

Figure 103: SNMP User Security Model



Use the buttons to perform the following tasks:

- To add a user, click Add. The Add New SNMP User dialog box opens. Specify the new account information in the available fields.
- To remove a user, select one or more table entries and click Remove to delete the selected entries.

Figure 104: Add New SNMP User

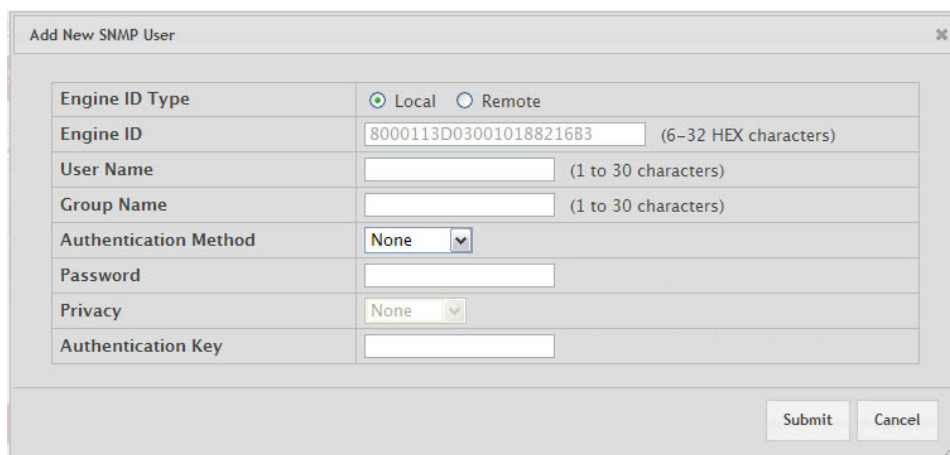


Table 92: Add New SNMP User Fields

Field	Description
Engine ID Type	Select the option for a local or remote engine ID type.
Engine ID	Each SNMPv3 agent has an engine ID that uniquely identifies the agent in the device. If given this entry will be used only for packets whose engine id is this. This field takes a hexadecimal string in the form 0102030405.

Table 92: Add New SNMP User Fields (Continued)

Field	Description
User Name	Specifies the name of the SNMP user being added for the User-based Security Model (USM). Each user name must be unique within the SNMP agent user list. A user name cannot contain any leading or embedded blanks.
Group Name	A SNMP group is a group to which hosts running the SNMP service belong. A group name parameter is simply the name of that group by which SNMP communities are identified. The use of a group name provides some security and context for agents receiving requests and initiating traps and does the same for management systems and their tasks. An SNMP agent won't respond to a request from a management system outside its configured group, but an agent can be a member of multiple groups at the same time. This allows for communications with SNMP managers from different groups.
Authentication Method	Specifies the authentication protocol to be used on authenticated messages on behalf of the specified user. <ul style="list-style-type: none"> <li>• SHA - SHA protocol will be used.</li> <li>• MD5 - MD5 protocol will be used.</li> <li>• None - No authentication will be used for this user.</li> </ul>
Password	Specifies the password used to generate the key to be used in authenticating messages on behalf of this user. This parameter must be specified if the Authentication method parameter is not NONE.
Privacy	Specifies the privacy protocol to be used on encrypted messages on behalf of the specified user. This parameter is only valid if the Authentication method parameter is not NONE. <ul style="list-style-type: none"> <li>• DES - DES protocol will be used.</li> <li>• None - No privacy protocol will be used.</li> </ul>
Authentication Key	Specifies the password used to generate the key to be used in encrypting messages to and from this user. This parameter must be specified if the Privacy parameter is not NONE.

If you make any changes to the page, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

#### 4.15.9 SNMP Source Interface Configuration

Use this page to specify the physical or logical interface to use as the SNMP client source interface. When an IP address is configured on the source interface, this address is used for all SNMP communications between the local SNMP client and the remote SNMP server. The IP address of the designated source interface is used in the IP header of SNMP management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access the SNMP Trap Source Interface Configuration page, click System > Advanced Configuration > SNMP > Source Interface Configuration in the navigation menu.

Figure 105: SNMP Trap Source Interface Configuration

The screenshot shows the 'SNMP Trap Source Interface Configuration' page. The navigation bar includes tabs for 'Community', 'Trap Receiver v1/v2', 'Trap Receiver v3', 'Supported MIBs', 'Access Control Group', 'User Security Model', and 'Source Interface Configuration'. The main content area is titled 'SNMP Trap Source Interface Configuration' and contains a form with the following fields:

- Type:** Radio buttons for None (selected), Interface, VLAN, Loopback, and Tunnel.
- Interface:** A dropdown menu currently set to 'Unconfigured'.
- VLAN ID:** A dropdown menu currently set to 'Unconfigured'.
- Loopback Interface:** A dropdown menu currently set to 'Unconfigured'.
- Tunnel ID:** A dropdown menu currently set to 'Unconfigured'.

At the bottom of the form are three buttons: 'Submit', 'Refresh', and 'Cancel'.

Table 93: SNMP Trap Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> <li>None – The primary IP address of the originating (outbound) interface is used as the source address.</li> <li>Interface – The primary IP address of a physical port is used as the source address.</li> <li>Loopback – The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up.</li> <li>VLAN – The primary IP address of a VLAN routing interface is used as the source address.</li> </ul>
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.

If you make any changes to the page, click Submit to apply the changes to the system. If you want the switch to retain the new values across a power cycle, you must perform a save.

## 4.16 Viewing System Statistics

The pages in the Statistics folder contain a variety of information about the number and type of traffic transmitted from and received on the switch.

### 4.16.1 Switch Statistics

The Switch Statistics page shows detailed statistical information about the traffic the switch handles. To access the Switch Statistics page, click System > Statistics > System > Switch in the navigation menu.

Figure 106: Switch Statistics

Statistics	Transmit	Receive
Octets Without Error	2391364	275872
Packets Without Errors	2039	1555
Packets Discarded	0	0
Unicast Packets	2027	1551
Multicast Packets	6	0
Broadcast Packets	6	4

Status	FDB Entries	VLANs
Current Usage	5	1
Peak Usage	5	1
Maximum Allowed	8192	255
Static Entries	1	1
Dynamic Entries	4	0
Total Entries Deleted	N/A	0

System	
Interface	209
Time Since Counters Last Cleared	0d:00:38:36

Table 94: Switch Statistics Fields

Field	Description
Statistics	
Octets Without Error	The total number of octets (bytes) of data successfully transmitted or received by the processor (excluding framing bits but including FCS octets).
Packets Without Errors	The total number of packets including unicast, broadcast, and multicast packets, successfully transmitted or received by the processor.
Packets Discarded	The number of outbound (Transmit column) or inbound (Receive column) packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. A possible reason for discarding a packet could be to free up buffer space.
Unicast Packets	The number of subnetwork-unicast packets delivered to or received from a higher-layer protocol.
Multicast Packets	The total number of packets transmitted or received by the device that were directed to a multicast address. <b>Note:</b> This number does not include packets directed to the broadcast address.
Broadcast Packets	The total number of packets transmitted or received by the device that were directed to the broadcast address. <b>Note:</b> This number does not include multicast packets.
Status	
Current Usage	In the FDB Entries column, the value shows the number of learned and static entries in the MAC address table. In the VLANs column, the value shows the total number of static and dynamic VLANs that currently exist in the VLAN database.
Peak Usage	The highest number of entries that have existed in the MAC address table or VLAN database since the most recent reboot.
Maximum Allowed	The maximum number of statically configured or dynamically learned entries allowed in the MAC address table or VLAN database.
Static Entries	The current number of entries in the MAC address table or VLAN database that an administrator has statically configured.
Dynamic Entries	The current number of entries in the MAC address table or VLAN database that have been dynamically learned by the device.
Total Entries Deleted	The number of VLANs that have been created and then deleted since the last reboot. This field does not apply to the MAC address table entries.
System	
Interface	The interface index object value of the interface table entry associated with the Processor of this switch. This value is used to identify the interface when managing the device by using SNMP.
Time Since Counters Last Cleared	The amount of time in days, hours, minutes, and seconds, that has passed since the statistics for this device were last reset.

- Click Refresh to refresh the data on the screen with the present state of the data in the switch.
- Click Clear Counters to clear all the statistics counters, resetting all switch summary and detailed statistics to default values. The discarded packets count cannot be cleared.

#### 4.16.2 Port Summary

This page shows statistical information about the packets received and transmitted by each port and LAG.

To access the Port Summary page, click System > Statistics > System > Port Summary in the navigation menu.

Figure 107: Port Summary Statistics

Note: All entries in this table indicate packet counts.

Display 10 rows Showing 1 to 10 of 116 entries Filter:

<input type="checkbox"/>	Interface	Rx Good	Rx Errors	Rx Bcast	Tx Good	Tx Errors	Tx Collisions
<input type="checkbox"/>	2/0/1	0	0	0	0	0	0
<input type="checkbox"/>	2/0/2	0	0	0	0	0	0
<input type="checkbox"/>	2/0/3	0	0	0	0	0	0
<input checked="" type="checkbox"/>	2/0/4	0	0	0	0	0	0
<input type="checkbox"/>	2/0/5	0	0	0	0	0	0
<input type="checkbox"/>	2/0/6	0	0	0	0	0	0
<input type="checkbox"/>	2/0/7	0	0	0	0	0	0
<input type="checkbox"/>	2/0/8	0	0	0	0	0	0
<input type="checkbox"/>	2/0/9	0	0	0	0	0	0
<input type="checkbox"/>	2/0/10	0	0	0	0	0	0

First Previous 1 2 3 4 5 Next Last

Refresh Clear Counters Clear All Counters

Table 95: Port Summary Statistics Fields

Field	Description
Interface	Identifies the port or LAG.
Rx Good	The total number of inbound packets received by the interface without errors.
Rx Errors	The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.
Rx Bcast	The total number of good packets received that were directed to the broadcast address. <b>Note:</b> This number does not include multicast packets.
Tx Good	The total number of outbound packets transmitted by the interface to its Ethernet segment without errors.
Tx Errors	The number of outbound packets that could not be transmitted because of errors.
Tx Collisions	The best estimate of the total number of collisions on this Ethernet segment.

- Click Refresh to refresh the data on the screen with the present state of the data in the switch.
- Click Clear Counters to clear all the statistics counters, resetting all summary and detailed statistics for this switch to default values. The discarded packets count cannot be cleared.
- Click Clear All Counters to clear counters for all switches in the stack.

### 4.16.3 Port Detailed Statistics

The Port Detailed page displays a variety of per-port traffic statistics.

To access the Port Detailed page, click System > Statistics > System > Port Detailed in the navigation menu.

Figure 108: "Port Detailed Statistics," on page 136 shows some, but not all, of the fields on the Port Detailed page.

Figure 108: Port Detailed Statistics

Interface	2/0/1	
Maximum Frame Size	1518	
Packet Lengths Received and Transmitted		
64 Octets	0	
65-127 Octets	0	
128-255 Octets	0	
256-511 Octets	0	
512-1023 Octets	0	
1024-1518 Octets	0	
1519-1522 Octets	0	
1523-2047 Octets	0	
2048-4095 Octets	0	
4096-9216 Octets	0	
Basic		
	Transmit	Receive
Unicast Packets	0	0
Multicast Packets	0	0
Broadcast Packets	0	0
Total Packets (Octets)	0	0
Packets > 1518 Octets	0	0
802.3x Pause Frames	0	0
FCS Errors		0
Protocol		
	Transmit	Receive
STP BPDUs	0	0
RSTP BPDUs	0	0
MSTP BPDUs	0	0
GVRP PDUs	0	0
GMRP PDUs	0	0

Table 96: Port Detailed Statistics Fields

Field	Description
Interface	Use the drop-down menu to select the interface for which data is to be displayed or configured. For non-stacking systems, this field is Slot/Port.
Maximum Frame Size	The maximum Ethernet frame size the interface supports or is configured to support. The maximum frame size includes the Ethernet header, CRC, and payload.
Packet Lengths Received and Transmitted	
64 Octets	The total number of packets (including bad packets) received or transmitted that were 64 octets in length (excluding framing bits but including FCS octets).
65-127 Octets	The total number of packets (including bad packets) received or transmitted that were between 65 and 127 octets in length inclusive (excluding framing bits but including FCS octets).
128-255 Octets	The total number of packets (including bad packets) received or transmitted that were between 128 and 255 octets in length inclusive (excluding framing bits but including FCS octets).
256-511 Octets	The total number of packets (including bad packets) received or transmitted that were between 256 and 511 octets in length inclusive (excluding framing bits but including FCS octets).
512-1023 Octets	The total number of packets (including bad packets) received or transmitted that were between 512 and 1023 octets in length inclusive (excluding framing bits but including FCS octets).
1024-1518 Octets	The total number of packets (including bad packets) received or transmitted that were between 1024 and 1518 octets in length inclusive (excluding framing bits but including FCS octets).

Table 96: Port Detailed Statistics Fields (Continued)

Field	Description
1519-1522 Octets	The total number of packets (including bad packets) received or transmitted that were between 1519 and 1522 octets in length inclusive (excluding framing bits but including FCS octets).
1523-2047 Octets	The total number of packets (including bad packets) received or transmitted that were between 1523 and 2047 octets in length inclusive (excluding framing bits but including FCS octets).
2048-4095 Octets	The total number of packets (including bad packets) received or transmitted that were between 2048 and 4095 octets in length inclusive (excluding framing bits but including FCS octets).
4096-9216 Octets	The total number of packets (including bad packets) received or transmitted that were between 4096 and 9216 octets in length inclusive (excluding framing bits but including FCS octets).
Basic	
Unicast Packets	The Transmit column shows the total number of packets that higher-level protocols requested be transmitted to a subnetwork unicast address, including those that were discarded or not sent. The Receive column shows the number of subnetwork unicast packets delivered to a higher-layer protocol.
Multicast Packets	The Transmit column shows the total number of packets that higher-level protocols requested be transmitted to a multicast address, including those that were discarded or not sent. The Receive column shows the number of multicast packets delivered to a higher-layer protocol.
Broadcast Packets	The Transmit column shows the total number of packets that higher-level protocols requested be transmitted to a broadcast address, including those that were discarded or not sent. The Receive column shows the number of broadcast packets delivered to a higher-layer protocol.
Total Packets (Octets)	The total number of octets of data (including those in bad packets) transmitted or received on the interface (excluding framing bits but including FCS octets). This object can be used as a reasonable estimate of Ethernet utilization. If greater precision is desired, the etherStatsPkts and etherStatsOctets objects should be sampled before and after a common interval.
Packets > 1518 Octets	The total number of packets transmitted or received by this interface that were longer than 1518 octets (excluding framing bits, but including FCS octets) and were otherwise well formed. This counter has a maximum increment rate of 815 counts per sec at 10 Mb/s.
802.3x Pause Frames	The number of MAC Control frames transmitted or received by this interface with an opcode indicating the PAUSE operation. This counter does not increment when the interface is operating in half-duplex mode.
FCS Errors	The total number of packets transmitted or received by this interface that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with an integral number of octets.
Protocol	
STP BPDUs	The number of Spanning Tree Protocol (STP) Bridge Protocol Data Units (BPDUs) transmitted or received by the interface.
RSTP BPDUs	The number of Rapid STP BPDUs transmitted or received by the interface.
MSTP BPDUs	The number of Multiple STP BPDUs transmitted or received by the interface.
SSTP BPDUs	The number of Shared STP BPDUs transmitted or received by the interface.
GVRP PDUs	The number of Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP) PDUs transmitted or received by the interface.
GMRP PDUs	The number of GARP Multicast Registration Protocol (GMRP) PDUs transmitted or received by the interface.

Table 96: Port Detailed Statistics Fields (Continued)

Field	Description
EAPOL Frames	The number of Extensible Authentication Protocol (EAP) over LAN (EAPOL) frames transmitted or received by the interface for IEEE 802.1X port-based network access control.
Advanced - Transmit	
Total Transmit Packets Discarded	The sum of single collision frames discarded, multiple collision frames discarded, and excessive frames discarded.
Single Collision Frames	A count of the number of successfully transmitted frames on a particular interface for which transmission is inhibited by exactly one collision.
Multiple Collision Frames	A count of the number of successfully transmitted frames on a particular interface for which transmission is inhibited by more than one collision.
Excessive Collision Frames	A count of frames for which transmission on a particular interface fails due to excessive collisions.
Underrun Errors	The total number of frames discarded because the transmit FIFO buffer became empty during frame transmission.
GMRP Failed Registrations	The number of times attempted GMRP registrations could not be completed.
GVRP Failed Registrations	The number of times attempted GVRP registrations could not be completed.
Percent Utilization Transmitted (%)	The amount of link utilization, represented as a percentage of total link bandwidth, for the TX direction.
Advanced - Receive	
Total Packets Received Not Forwarded	The number of inbound packets which were chosen to be discarded to prevent them from being delivered to a higher-layer protocol, even though no errors had been detected. One possible reason for discarding such a packet is to free up buffer space.
Total Packets Received With MAC Errors	The total number of inbound packets that contained errors preventing them from being delivered to a higher-layer protocol.
Overruns	The total number of frames discarded as this port was overloaded with incoming packets, and could not keep up with the inflow.
Alignment Errors	The total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, but had a bad Frame Check Sequence (FCS) with a non-integral number of octets.
Jabbers Received	The total number of packets received that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either a bad Frame Check Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). <b>Note:</b> This definition of jabber is different than the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition where any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.
Fragments Received	The total number of packets received that were less than 64 octets in length with ERROR CRC (excluding framing bits but including FCS octets).
Undersize Received	The total number of packets received that were less than 64 octets in length with GOOD CRC (excluding framing bits but including FCS octets).
Unacceptable Frame Type	The number of frames discarded from this interface due to being a frame type that the interface cannot accept.
Percent Utilization Received (%)	The amount of link utilization, represented as a percentage of total link bandwidth, for the RX direction.
Time Since Counters Last Cleared	The amount of time in days, hours, minutes, and seconds, that has passed since the statistics for this interface were last reset.



- Click Clear Counters to clear all the counters. This resets all statistics for this port to the default values.
- Click Clear All Counters to clear all the counters for all ports on the switch. The button resets all statistics for all ports to default values.
- Click Refresh to refresh the data on the screen and display the most current statistics.

#### 4.16.4 Port DHCPv6 Client Statistics

This page displays the DHCPv6 client statistics values for the selected interface. The DHCPv6 client on the device exchanges several different types of UDP messages with one or more network DHCPv6 servers during the process of acquiring address, prefix, or other relevant network configuration information from the server. The values indicate the various counts that have accumulated since they were last cleared.

To display the Port DHCPv6 Client Statistics page, click System > Statistics > System > DHCPv6.

Figure 109: Port DHCPv6 Client Statistics

Interface	1/0/1
Advertisement Packets Received	0
Reply Packets Received	0
Received Advertisement Packets Discarded	0
Received Reply Packets Discarded	0
Malformed Packets Received	0
Total Packets Received	0
Solicit Packets Transmitted	0
Request Packets Transmitted	0
Renew Packets Transmitted	0
Rebind Packets Transmitted	0
Release Packets Transmitted	0
Decline Packets Transmitted	0
Confirm Packets Transmitted	0
Total Packets Transmitted	0

Table 97: Port DHCPv6 Client Statistics Fields

Field	Description
Interface	Select the interface to view the DHCPv6 client statistics associated with it.
Advertisement Packets Received	Number of DHCPv6 advertisement messages received from one or more DHCPv6 servers in response to the client's solicit message.
Reply Packets Received	Number of DHCPv6 reply messages received from one or more DHCPv6 servers in response to the client's request message.
Received Advertisement Packets Discarded	Number of DHCPv6 advertisement messages received from one or more DHCPv6 servers to which the client did not respond.
Received Reply Packets Discarded	Number of DHCPv6 reply messages received from one or more DHCPv6 servers to which the client did not respond.
Malformed Packets Received	Number of messages received from one or more DHCPv6 servers that were improperly formatted.
Total Packets Received	Total number of messages received from all DHCPv6 servers.
Solicit Packets Transmitted	Number of DHCPv6 solicit messages the client sent to begin the process of acquiring network information from a DHCPv6 server.
Request Packets Transmitted	Number of DHCPv6 request messages the client sent in response to a DHCPv6 server's advertisement message.

Table 97: Port DHCPv6 Client Statistics Fields (Continued)

Field	Description
Renew Packets Transmitted	Number of renew messages the DHCPv6 client has sent to the server to request an extension of the lifetime of the information provided by the server. This message is sent to the DHCPv6 server that originally assigned the addresses and configuration information.
Rebind Packets Transmitted	Number of rebind messages the DHCPv6 client has sent to any available DHCPv6 server to request an extension of its addresses and an update to any other relevant information. This message is sent only if the client does not receive a response to the renew message.
Release Packets Transmitted	Number of release messages the DHCPv6 client has sent to the server to indicate that it no longer needs one or more of the assigned addresses.
Decline Packets Transmitted	Number of decline messages the DHCPv6 client has sent to the server to indicate that one or more addresses assigned by the server are already in use on the connected link.
Confirm Packets Transmitted	Number of confirm messages the DHCPv6 client has sent to any available DHCPv6 server to determine whether the addresses it is assigned are still valid for the connected link.
Total Packets Transmitted	Total number of messages sent to all DHCPv6 servers.

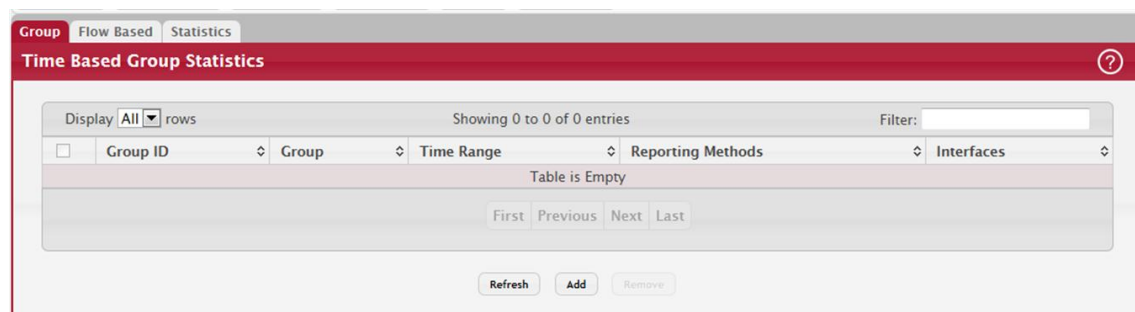
- Click Clear Counters to clear all the statistics counters, resetting all switch summary and detailed statistics to default values.

#### 4.16.5 Time Based Group Statistics

Use this page to define criteria for collecting time-based statistics for interface traffic. The time-based statistics can be useful for troubleshooting and diagnostics purposes. The statistics application uses the system clock for time-based reporting, so it is important to configure the system clock (manually or through SNTP) before using this feature.

To access the Time Based Group Statistics page, click System > Statistics > Time Based > Group in the navigation menu.

Figure 110: Time Based Group Statistics



Use the buttons to perform the following tasks:

- To add a set of time-based traffic group statistics to collect, click Add and configure the desired settings.
- To delete one or more time-based statistics groups, select each entry to delete and click Remove.

Table 98: Time Based Group Statistics Fields

Field	Description
Group	The type of traffic statistics to collect for the group, which is one of the following: <ul style="list-style-type: none"> <li>Received – The number of packets received on the interfaces within the group.</li> <li>Received Errors – The number of packets received with errors on the interfaces within the group.</li> <li>Transmitted – The number of packets transmitted by the interfaces within the group.</li> <li>Received Transmitted – The number of packets received and transmitted by the interfaces within the group.</li> <li>Port Utilization – The percentage of total bandwidth used by the port within the specified time period.</li> <li>Congestion – The percentage of time within the specified time range that the ports experienced congestion.</li> </ul>
Time Range	The name of the periodic or absolute time range to use for data collection. The time range is configured by using the Time Range Summary and Time Range Entry Summary pages. The time range must be configured on the system before the time-based statistics can be collected.
Reporting Methods	The methods for reporting the collected statistics at the end of every configured time range interval. The available options are: <ul style="list-style-type: none"> <li>None – The statistics are not reported to the console or an external server. They can be viewed only by using the web interface or by issuing a CLI command.</li> <li>Console – The statistics are displayed on the console.</li> <li>E-Mail – The statistics are sent to an e-mail address. The SMTP server and e-mail address information is configured by using the appropriate Email Alerts pages.</li> <li>Syslog – The statistics are sent to a remote syslog server. The syslog server information is configured on the Logging Hosts page.</li> </ul>
Interfaces	The interface or interfaces on which data is collected. To select multiple interfaces when adding a new group, CTRL + click each interface to include in the group.

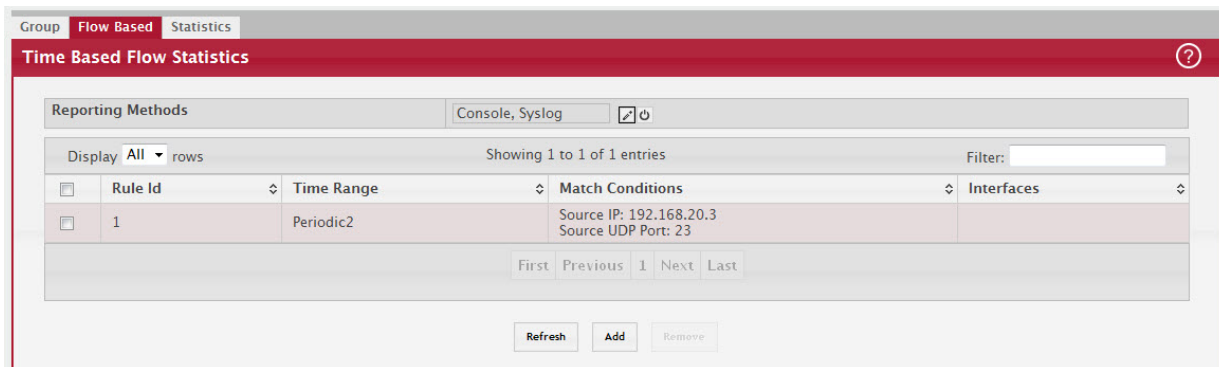
- Click Refresh to refresh the data on the screen with the present state of the data in the switch.
- To add a set of time-based traffic group statistics to collect, click Add and configure the desired settings.
- To delete one or more time-based statistics groups, select each entry to delete and click Remove.

#### 4.16.6 Time Based Flow Statistics

Use this page to define criteria for collecting time-based statistics for specific traffic flows. The statistics include a per-interface hit count based on traffic that meets the match criteria configured in a rule for the interfaces included in the rule. The hit count statistics are collected only during the specified time range. The statistics application uses the system clock for time-based reporting. Configure the system clock (manually or through SMTP) before using the time-based statistics feature.

To access the Time Based Flow Statistics page, click System > Statistics > Time Based > Flow Based in the navigation menu.

Figure 111: Time Based Flow Statistics



Use the buttons to perform the following tasks:

- To add a rule and define criteria for flow-based statistics that are collected within a time range, click Add and configure the desired settings.
- To delete one or more flow-based rules for time-based statistics, select each entry to delete and click Remove.

Table 99: Time Based Flow Statistics Fields

Field	Description
Reporting Methods	The methods for reporting the collected statistics at the end of every configured interval. To change the reporting methods for all flow-based statistics rules, click the Edit icon and select one or more methods. To reset the field to the default value, click the Reset icon. The available reporting methods are: <ul style="list-style-type: none"> <li>• None – The statistics are not reported to the console or an external server. They can be viewed only by using the web interface or by issuing a CLI command.</li> <li>• Console – The statistics are displayed on the console.</li> <li>• E-Mail – The statistics are sent to an e-mail address. The SMTP server and e-mail address information is configured by using the appropriate Email Alerts pages.</li> <li>• Syslog – The statistics are sent to a remote syslog server. The syslog server information is configured on the Logging Hosts page.</li> </ul>
Rule Id	The number that identifies the flow-based statistics collection rule.
Time Range	The name of the periodic or absolute time range to use for data collection. The time range is configured by using the Time Range Summary and Time Range Entry Summary pages. The time range must be configured on the system before the time-based statistics can be collected.
Match Conditions	The criteria a packet must meet to match the rule.
Interfaces	The interface or interfaces on which the flow-based rule is applied. Only traffic on the specified interfaces is checked against the rule.
After you click Add, the Time Based Flow Configuration window opens and allows you to configure a rule for traffic flow statistics. The match conditions are optional, but the rule must specify at least one match condition. The following information describes the match criteria fields that are available in this window.	
Match All	Select this option to indicate that all traffic matches the rule and is counted in the statistics. This option is exclusive to all other match criteria, so if Match All is selected, no other match criteria can be configured.
Source IP	The source IP address to match in the IPv4 packet header.
Destination IP	The destination IP address to match in the IPv4 packet header.
Source MAC	The source MAC address to match in the ingress frame header.
Destination MAC	The destination MAC address to match in the ingress frame header.
Source TCP Port	The TCP source port to match in the TCP header.
Destination TCP Port	The TCP destination port to match in the TCP header.

**Table 99: Time Based Flow Statistics Fields (Continued)**

Field	Description
Source UDP Port	The UDP source port to match in the UDP header.
Destination UDP Port	The UDP destination port to match in the UDP header.

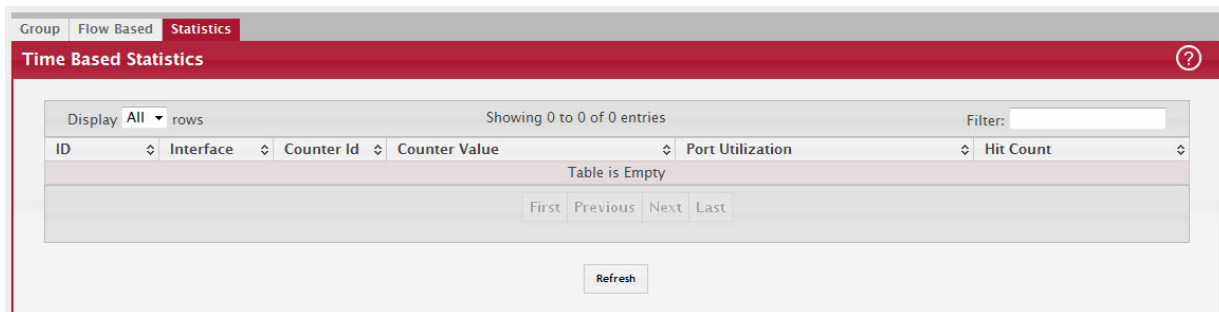
- Click Refresh to refresh the data on the screen with the present state of the data in the switch.
- To add a set of time-based traffic group statistics to collect, click Add and configure the desired settings.
- To delete one or more time-based statistics groups, select each entry to delete and click Remove.

### 4.16.7 Time Based Statistics

Use this page to view time-based statistics collected for the configured traffic groups and flow-based rules.

To access the Time Based Statistics page, click System > Statistics > Time Based > Statistics in the navigation menu.

**Figure 112: Time Based Statistics**



**Table 100: Time Based Statistics Fields**

Field	Description
ID	The traffic group name or flow-based rule ID associated with the rest of the statistics in the row.
Interface	The interface on which the statistics were reported.
Counter ID	For traffic group statistics, this field identifies the type of traffic.
Counter Value	For traffic group statistics, this field shows the number of packets of the type identified by the Counter Id field that were reported on the interface during the time range.
Port Utilization	For a port utilization traffic group, this field reports the percentage of the total available bandwidth used on the interface during the time range.
Hit Count	For flow-based statistics, this field reports the number of packets that matched the flow-based rule criteria during the time range.

- Click Refresh to refresh the data on the screen with the present state of the data in the switch.

## 4.17 Using System Utilities

The System Utilities feature menu contains links to Web pages that help you configure features that help you manage the switch.

### 4.17.1 System Reset

Use the System Reset page to reboot the system. If the platform supports stacking, you can reset any of the switches in the stack, or all switches in the stack from this page.

To access the System Reset page, click System > Utilities > System Reset in the navigation menu.

Figure 113: System Reset

Table 101: System Reset Fields

Field	Description
Generate Core Dump before reset	Generates core dump file on demand.
Switch ID	Select the specific switch unit to be reset, or specify All to reset all units in the stack.
Reset (Button)	Initiates the system reset action after displaying a confirmation message. <b>Note:</b> Any configuration changes made since the last successful save are lost whenever a switch is reset. It is possible that the IP address of the switch will change. If this occurs you will need to determine the new IP address to access the device using the web.

For Stacking platforms, you can select one or all switches in the stack to reset from the drop-down menu. For platforms that do not support stacking, this field is not present.

Click Reset to initiate the system reset. If you have not saved the changes that you submitted since the last system reset, the changes will not be applied to the system after the reset.

### 4.17.2 Ping

Use the Ping page to tell the switch to send a Ping request to a specified IP address. You can use this feature to check whether the switch can communicate with a particular network host.

To access the Ping page, click System > Utilities > Ping in the navigation menu.

Figure 114: Ping

Table 102: Ping Fields

Field	Description
Hostname/IP Address	Enter the IP address or the host name of the station you want the switch to ping. The initial value is blank. This information is not retained across a power cycle.
Count	The number of ICMP echo request packets to send to the host.
Interval	The number of Seconds to wait between sending ping packets.
Size	The size of the ping packet, in bytes. Changing the size allows you to troubleshoot connectivity issues with a variety of packet sizes, such as large or very large packets.
Source	The source IP address or interface to use when sending the echo request packets. If source is not required, select None as source option.
IP Address	The source IP address to use when sending the Echo requests packets. This field is enabled when IP Address is selected as source option.
Interface	The interface to use when sending the Echo requests packets. This field is enabled when Interface is selected as source option.
Status	Displays the results of the ping.
Results	The results of the ping test, which includes information about the reply (if any) received from the host.
Start (Button)	Starts the ping test. The device sends the specified number of ping packets to the host.
Stop (Button)	Interrupts the current ping test.

### 4.17.3 Ping IPv6

Use the Ping IPv6 page to tell the device to send one or more ping requests to a specified IPv6 host. You can use the ping request to check whether the device can communicate with a particular host on an IPv6 network. A ping request is an Internet Control Message Protocol version 6 (ICMPv6) echo request packet. The information you enter on this page is not saved as part of the device configuration.

To access the Ping IPv6 page, click System > Utilities > Ping IPv6 in the navigation menu.

Figure 115: Ping IPv6

Table 103: Ping IPv6 Fields

Field	Description
Ping	Select either a global IPv6 address or a link local address to ping. A global address is routable over the Internet, while a link-local address is intended for communication only within the local network. Link local addresses have a prefix of fe80::/64.
Interface	This field displays only when <code>Link Local</code> is selected. Select an IPv6 interface to initiate the ping.
Host Name or IPv6 Address	Enter the global or link-local IPv6 address, or the DNS-resolvable host name of the station to ping. If the ping type is <code>Link Local</code> , you must enter a link-local address and cannot enter a host name.
Count	Enter the number of ICMP echo request packets to send to the host.
Interval	Enter the number of seconds to wait between sending ping packets.
Size	The size of the ping packet, in bytes. Changing the size allows you to troubleshoot connectivity issues with a variety of packet sizes, such as large or very large packets.
Source	The source IP address or interface to use when sending the echo request packets. If source is not required, select <code>None</code> as source option.
IPv6 Address	The source IPv6 address to use when sending the Echo requests packets. This field is enabled when <code>IP Address</code> is selected as source option.
Interface	The interface to use when sending the Echo requests packets. This field is enabled when <code>Interface</code> is selected as source option.
Results	The results of the ping test, which includes information about the reply (if any) received from the host.

Click `Submit` to send the specified number of pings. The results display in the `Ping Output` box.



## 4.17.4 TraceRoute

Use this page to determine the Layer 3 path a packet takes from the device to a specific IP address or hostname. When you initiate the TraceRoute command by clicking the Start button, the device sends a series of TraceRoute probes toward the destination. The results list the IP address of each Layer 3 device a probe passes through until it reaches its destination - or fails to reach its destination and is discarded. The information you enter on this page is not saved as part of the device configuration.

To access the TraceRoute page, click System > Utilities > TraceRoute in the navigation menu.

Figure 116: TraceRoute

Table 104: TraceRoute Fields

Field	Description
Host Name or IP Address	The DNS-resolvable hostname or IP address of the system to attempt to reach.
Probes Per Hop	TraceRoute works by sending UDP packets with increasing Time-To-Live (TTL) values. Specify the number of probes sent with each TTL.
MaxTTL	The maximum Time-To-Live (TTL). The TraceRoute terminates after sending probes that can be Layer 3 forwarded this number of times. If the destination is further away, the TraceRoute will not reach it.
InitTTL	The initial Time-To-Live (TTL). This value controls the maximum number of Layer 3 hops that the first set of probes may travel.
MaxFail	The number of consecutive failures that terminate the TraceRoute. If the device fails to receive a response for this number of consecutive probes, the TraceRoute terminates.
Interval	The number of Seconds to wait between sending probes.
Port	The UDP destination port number to be used in probe packets. The port number should be a port that the target host is not listening on, so that when the probe reaches the destination, it responds with an ICMP Port Unreachable message.
Size	The size of probe payload in bytes.
Source	Select None, IP Address, Interface, or Loopback as a source.

Table 104: TraceRoute Fields (Continued)

Field	Description
Status	The current status of the TraceRoute, which can be: <ul style="list-style-type: none"> <li>Not Started – The TraceRoute has not been initiated since viewing the page.</li> <li>In Progress – The TraceRoute has been initiated and is running.</li> <li>Stopped – The TraceRoute was interrupted by clicking the Stop button.</li> <li>Done – The TraceRoute has completed, and information about the TraceRoute is displayed in the Results area.</li> </ul>
Results	The results of the TraceRoute are displayed
Start (Button)	Initiates the TraceRoute.
Stop (Button)	Interrupts the running TraceRoute.

#### 4.17.5 IP Address Conflict Detection

Use the IP Address Conflict Detection page to determine whether the IP address configured on the device is the same as the IP address of another device on the same LAN (or on the Internet, for a routable IP address) and to help you resolve any existing conflicts. An IP address conflict can make both this system and the system with the same IP address unusable for network operation.

To access the IP Address Conflict Detection page, click System > Utilities > IP Address Conflict in the navigation menu.

Figure 117: IP Address Conflict Detection



Table 105: IP Address Conflict Detection Fields

Field	Description
IP Address Conflict Currently Exists	Indicates whether a conflicting IP address has been detected since this status was last reset. <ul style="list-style-type: none"> <li>False – No conflict detected (the subsequent fields on this page display as N/A).</li> <li>True – Conflict was detected (the subsequent fields on this page show the relevant information).</li> </ul>
Last Conflicting IP Address	The device interface IP address that is in conflict. If multiple conflicts were detected, only the most recent occurrence is displayed.
Last Conflicting MAC Address	The MAC address of the remote host associated with the IP address that is in conflict. If multiple conflicts are detected, only the most recent occurrence is displayed.
Time Since Conflict Detected	The elapsed time (displayed in days, hours, minutes, and seconds) since the last address conflict was detected, provided the Clear History button has not yet been pressed.
Run Detection (Button)	Activates the IP address conflict detection operation in the system.
Clear History (Button)	Resets the IP address conflict detection status information that was last seen by the device.

## 4.17.6 File Transfer

Use the File Transfer page to upload files from the device to a remote system and to download files from a remote system to the device.

To access the File Transfer page, click System > Utilities > Transfer in the navigation menu.

Figure 118: File Transfer

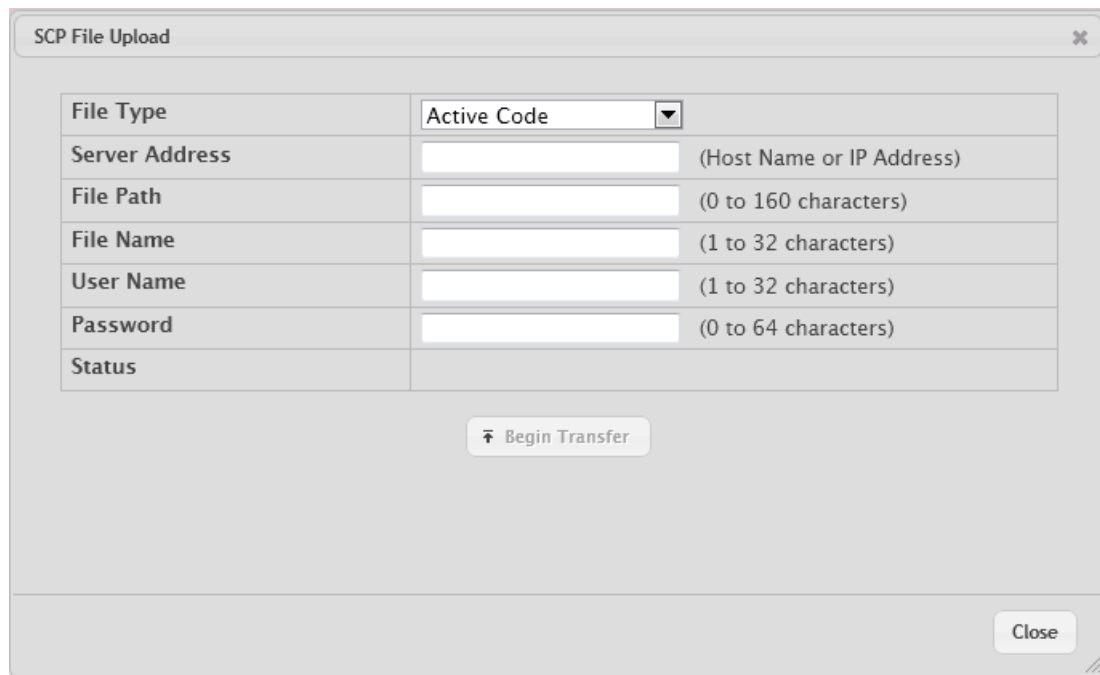


Table 106: File Transfer Fields

Field	Description
Transfer Protocol	The protocol to use to transfer the file. Files can be transferred from the device to a remote system using TFTP, FTP, SCP or SFTP. Files can be transferred from a remote system to the device using HTTP, TFTP, FTP, SCP or SFTP.
Upload	To transfer a file from the device to a remote system using TFTP, FTP, SCP, or SFTP, click the upload icon in the same row as the desired transfer protocol. The File Upload window appears. Configure the information for the file transfer (described below), and click the upload icon to the right of the Progress field to begin the transfer.
Download	To transfer a file from a remote system to the device using HTTP, TFTP, FTP, SCP, or SFTP, click the download icon in the same row as the desired transfer protocol. The File Download window appears. Configure the information for the file transfer (described below), and click the download icon to the right of the Progress field to begin the transfer.

After you click the upload icon, the SCP File Upload window appears.

Figure 119: SCP File Upload



The following information describes the fields in the SCP File Upload window for all protocols.

Table 107: SCP File Upload Fields

Field	Description
File Type	<p>Specify the type of file to transfer from the device to a remote system.</p> <ul style="list-style-type: none"> <li>Active Code – Select this option to transfer an active image.</li> <li>Backup Code – Select this option to transfer a backup image.</li> <li>Startup Configuration – Select this option to transfer a copy of the stored startup configuration from the device to a remote system.</li> <li>Backup Configuration – Select this option to transfer a copy of the stored backup configuration from the device to a remote system.</li> <li>Script File – Select this option to transfer a custom text configuration script from the device to a remote system.</li> <li>CLI Banner – Select this option to transfer the file containing the text to be displayed on the CLI before the login prompt to a remote system.</li> <li>Crash Log – Select this option to transfer the system crash log to a remote system.</li> <li>Operational Log – Select this option to transfer the system operational log to a remote system.</li> <li>Startup Log – Select this option to transfer the system startup log to a remote system.</li> <li>Trap Log – Select this option to transfer the system trap records to a remote system.</li> <li>Factory Defaults – Select this option to transfer the factory default configuration file to a remote system.</li> <li>Error Log – Select this option to transfer the system error (persistent) log, which is also known as the event log, to a remote system.</li> <li>Buffered Log – Select this option to transfer the system buffered (in-memory) log to a remote system.</li> </ul>
Image	<p>If the selected File Type is Code, specify whether to transfer the Active or Backup image to a remote system.</p>

**Table 107: SCP File Upload Fields (Continued)**

Field	Description
Server Address	Specify the IPv4 address, IPv6 address, or DNS-resolvable hostname of the remote server that will receive the file.
File Path	Specify the path on the server where you want to put the file.
File Name	Specify the name that the file will have on the remote server.
User Name	For FTP, SCP, and SFTP transfers, if the server requires authentication, specify the user name for remote login to the server that will receive the file.
Password	For FTP, SCP and SFTP transfers, if the server requires authentication, specify the password for remote login to the server that will receive the file.
Progress	Represents the completion percentage of the file transfer. The file transfer begins after you complete the required fields and click the upload icon to the right of this field.
Status	Provides information about the status of the file transfer.

After you click the download icon, the SFTP File Download window appears.

**Figure 120: SFTP File Download**

The screenshot shows a window titled "SFTP File Download" with a close button in the top right corner. The window contains a table of configuration fields:

File Type	Active Code	
Certificate Index	0	(1 to 8; 0 for None)
Server Address		(Host Name or IP Address)
File Path		(0 to 160 characters)
File Name		(1 to 32 characters)
User Name		(1 to 32 characters)
Password		(0 to 64 characters)
Digital Signature Verification	<input type="checkbox"/>	
Status		

Below the table is a button labeled "Begin Transfer" with a download icon. In the bottom right corner of the window is a "Close" button.

The following information describes the fields in the SFTP File Download window for all protocols.

Table 108: SFTP File Download Fields

Field	Description
File Type	<p>Specify the type of file to transfer to the device:</p> <ul style="list-style-type: none"> <li>Active Code – Select this option to transfer a new image to the device. The code file is stored as the active image.</li> <li>Backup Code – Select this option to transfer a new image to the device. The code file is stored as the backup image.</li> <li>Startup Configuration – Select this option to update the stored startup configuration file. If the file has errors, the update will be stopped.</li> <li>Backup Configuration – Select this option to update the stored backup configuration file. If the file has errors, the update will be stopped.</li> <li>Script File – Select this option to transfer a text-based configuration script to the device. You must use the command-line interface (CLI) to validate and activate the script.</li> <li>CLI Banner – Select this option to transfer the CLI banner file to the device. This file contains the text to be displayed on the CLI before the login prompt.</li> <li>IAS Users – Select this option to transfer an Internal Authentication Server (IAS) users database file to the device. The IAS user database stores a list of user name and (optional) password values for local port-based user authentication.</li> <li>SSH-1 RSA Key File – Select this option to transfer an SSH-1 Rivest-Shamir-Adleman (RSA) key file to the device. SSH key files contain information to authenticate SSH sessions for remote CLI-based access to the device.</li> </ul> <p><b>Note:</b> The SSH1-RSA can be downloaded, but they cannot be used.</p> <ul style="list-style-type: none"> <li>SSH-2 RSA Key PEM File – Select this option to transfer an SSH-2 Rivest-Shamir-Adleman (RSA) key file (PEM Encoded) to the device.</li> <li>SSH-2 DSA Key PEM File – Select this option to transfer an SSH-2 Digital Signature Algorithm (DSA) key file (PEM Encoded) to the device.</li> <li>Factory Defaults – Select this option to transfer the factory default configuration file to a remote system.</li> <li>CA Root Certificate – Select this option to transfer an CA certificate file to the device. This will be used as the root certificate for one of the syslog servers. Based on the index number the file will be named accordingly.</li> <li>Client Key – Select this option to transfer an client certificate file to the device. This will be used as the client certificate for one of the syslog servers. Based on the index number the file will be named accordingly.</li> <li>Client SSL Certificate – Select this option to transfer an client key file to the device. Based on the index number the file will be named accordingly.</li> <li>SSL Trusted Root Certificate PEM File – Select this option to transfer an SSL Trusted Root Certificate file (PEM Encoded) to the device. SSL files contain information to encrypt, authenticate, and validate HTTPS sessions.</li> <li>SSL Server Certificate PEM File – Select this option to transfer an SSL Server Certificate file (PEM Encoded) to the device.</li> <li>SSL DH Weak Encryption Parameter PEM File – Select this option to transfer an SSL Diffie-Hellman Weak Encryption Parameter file (PEM Encoded) to the device.</li> <li>SSL DH Strong Encryption Parameter PEM File – Select this option to transfer an SSL Diffie-Hellman Strong Encryption Parameter file (PEM Encoded) to the device.</li> <li>Public Key Image – Select this option to transfer the public key file used for code image validation to the device.</li> <li>Public Key Config – Select this option to transfer the public key file used for configuration script validation to the device.</li> </ul> <p><b>Note:</b></p> <ul style="list-style-type: none"> <li>To download SSH key files, SSH must be administratively disabled, and there can be no active SSH sessions.</li> <li>To download SSL related files, HTTPS must be administratively disabled.</li> </ul>
Select File	If HTTP is the Transfer Protocol, browse to the directory where the file is located and select the file to transfer to the device. This field is not present if the Transfer Protocol is TFTP or FTP.
Certificate Index	Index used to name a related group of certificate (PEM) or key files.

Table 108: SFTP File Download Fields (Continued)

Field	Description
Server Address	For TFTP, FTP, SCP, or SFTP transfers, specify the IPv4 address, IPv6 address, or DNS-resolvable hostname of the remote server.
File Path	For TFTP, FTP, SCP, or SFTP transfers, specify the path on the server where the file is located.
File Name	For TFTP, FTP, SCP, or SFTP transfers, specify the name of the file you want to transfer to the device.
User Name	For FTP, SCP, or SFTP transfers, if the server requires authentication, specify the user name for remote login to the server where the file resides.
Password	For FTP, SCP, or SFTP transfers, if the server requires authentication, specify the password for remote login to the server where the file resides.
Progress	Represents the completion percentage of the file transfer. The file transfer begins after you complete the required fields and click the download icon to the right of this field.
Digital Signature Verification	For Code and Startup Configuration file types, this option, when checked, will verify the file download with the digital signature.
Status	Provides information about the status of the file transfer.

#### 4.17.7 Digital Signature Verification

Use the Digital Signature Verification page to configure digital signature verification on downloading files from a remote system to the device.

To access the Digital Signature Verification page, click System > Utilities > Digital Signature Verification in the navigation menu.

Figure 121: Digital Signature Verification

Digital Signature Verification	
Digital Signature Verification	<input checked="" type="checkbox"/>
Code	<input type="checkbox"/>
Configuration	<input type="checkbox"/>

Table 109: Digital Signature Verification Fields

Field	Description
Digital Signature Verification	Provides option to verify the digital signature of a downloaded file.
Code	Verify the digital signature of downloaded code image files.
Configuration	Verify the digital signature of downloaded configuration script files.

#### 4.17.8 Core Dump

Use the Core Dump page to configure the Core Dump feature.

To access the Core Dump page, click System > Utilities > Core Dump in the navigation menu.

Figure 122: Core Dump

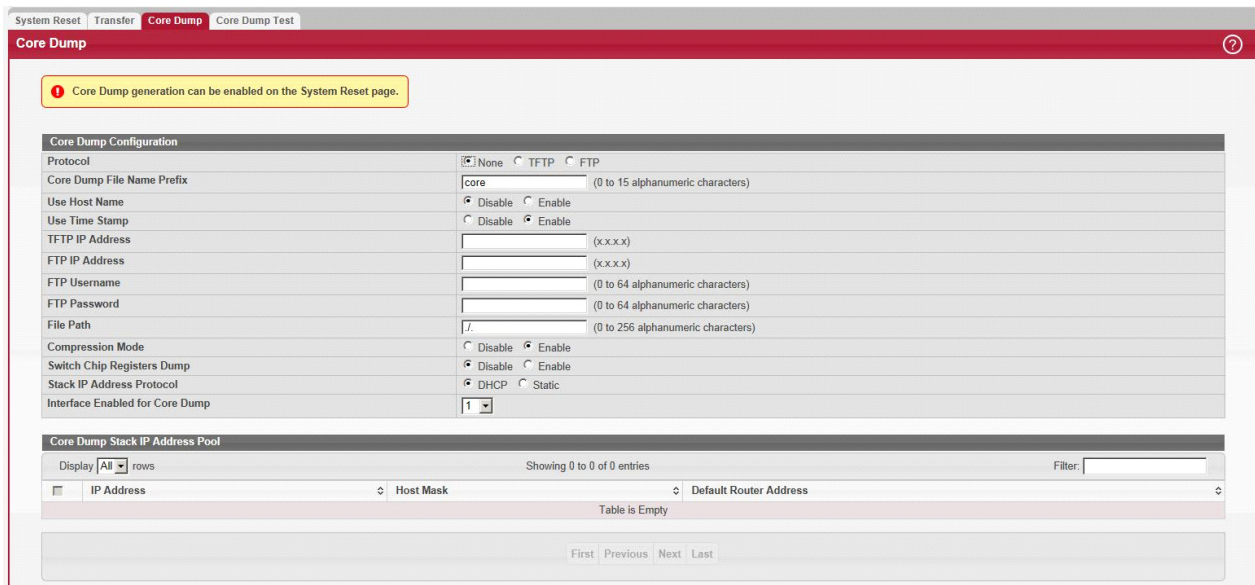


Table 110: Core Dump Fields

Field	Description
Protocol	The protocol used to store the core dump file. User can select: <ul style="list-style-type: none"> <li>• None–Disable Core Dump.</li> <li>• TFTP–Configure protocol to upload Core Dump to the TFTP server.</li> <li>• FTP–Configure protocol to upload Core Dump to the FTP server.</li> </ul>
Core Dump File Name Prefix	Prefix for the Core Dump file name. If hostname is configured, it takes else while generating Core Dump file. The prefix length is 15 characters.
Use Host Name	To use hostname (or MAC if hostname is not configured) to name Core Dump file.
Use Time Stamp	To use timestamp to name Core Dump file.
TFTP IP Address	IP address of remote TFTP server to dump core file to external server.
FTP IP Address	IP address of remote FTP server to dump core file to external server.
FTP Username	Username of remote FTP server.
FTP Password	Password of remote FTP server.
File Path	File path to dump core file to TFTP server, NFS mount or USB device sub-directory.
Compression Mode	To enable or disable compression mode.
Switch Chip Registers Dump	To enable or disable switch-chip-register dump in case of an exception. The switch-chip-register dump is taken only for master unit and not for member units.
Stack IP Address Protocol	Protocol (DHCP or Static) to be used to configure service port when a unit has crashed. If configured as DHCP, the unit gets the IP address from DHCP server available in the network. If configured as Static, an IP address from the Core Dump Stack IP Address Pool is used.



Table 111: Core Dump Stack IP Address Pool Fields

Field	Description
IP Address	Static IP address to be assigned to individual unit's service port in the stack when the switch has crashed. This IP address is used to perform the core dump.
Host Mask	The subnet mask.
Default Router Address	The IP address of the router.

To add a stack IP address, click Add and configure an IP address, netmask, and gateway address.

To delete a configured stack IP, select each entry to delete, click Remove, and confirm the action.

#### 4.17.9 Core Dump Test

Use the Core Dump Test page to test the core dump setup. For example, if protocol is configured as TFTP, it communicates with the TFTP server and informs the user if the TFTP server can be contacted.

To access the Core Dump Test page, click System > Utilities > Core Dump Test in the navigation menu.

Figure 123: Core Dump Test

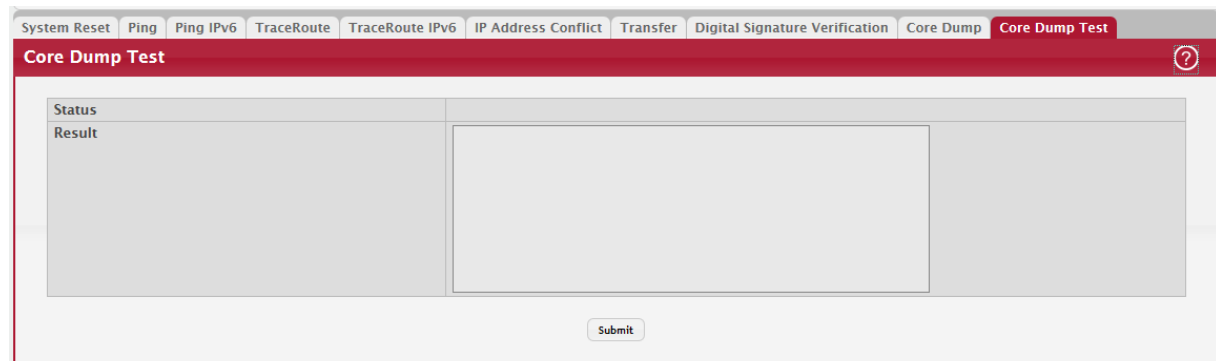


Table 112: Core Dump Test Fields

Field	Description
Status	Displays test status as OK if test passes and Error if test fails.
Result	Displays detailed error information with logs.

## 4.18 Managing SNMP Traps

The pages in the Trap Manager folder allow you to view and configure information about SNMP traps the system generates.

### 4.18.1 System Trap Log

Use the System Trap Log page to view the entries in the trap log.

To access the System Trap Log page, click System > Advanced Configuration > Trap Manager > Trap Log in the navigation menu.

Figure 124: System Trap Log

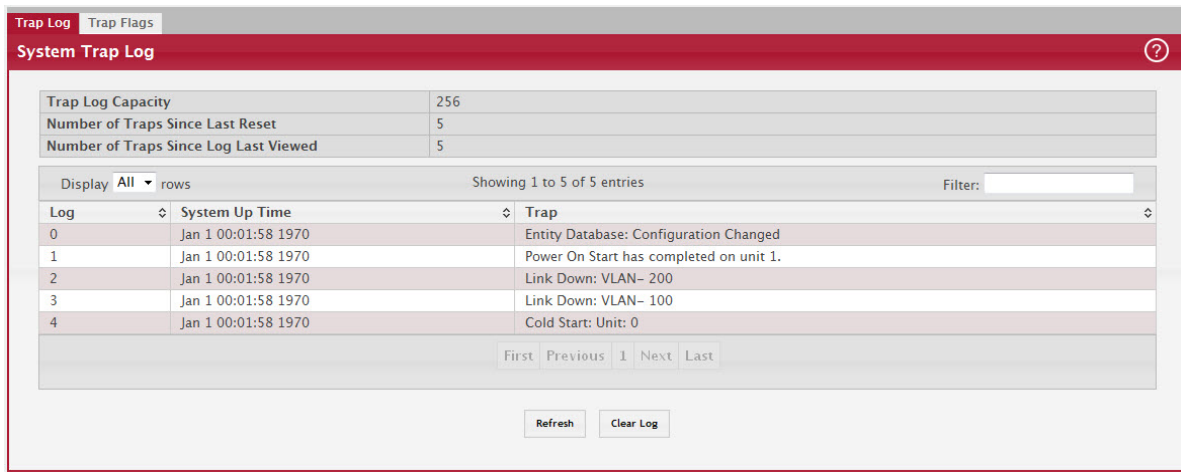


Table 113: System Trap Log Fields

Field	Description
Trap Log Capacity	The maximum number of traps stored in the log. If the number of traps exceeds the capacity, the entries will overwrite the oldest entries.
Number of Traps Since Last Reset	The number of traps generated since the trap log entries were last cleared.
Number of Traps Since Log Last Viewed	The number of traps that have occurred since the traps were last displayed. Displaying the traps by any method (terminal interface display, Web display, upload file from switch, etc.) will cause this counter to be cleared to 0.
Log	The sequence number of this trap.
System Up Time	The time at which this trap occurred, expressed in days, hours, minutes and seconds since the last reboot of the switch.
Trap	Displays the information identifying the trap.

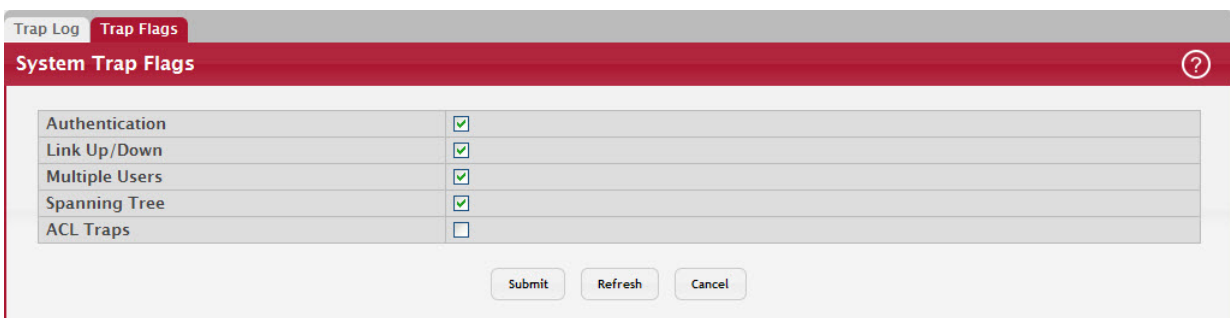
Click Clear Log to clear all entries in the log. Subsequent displays of the log will only show new log entries.

### 4.18.2 System Trap Flags

Use the System Trap Flags page to enable or disable traps the switch can send to an SNMP manager. When the condition identified by an active trap is encountered by the switch, a trap message is sent to any enabled SNMP Trap Receivers, and a message is written to the trap log.

To access the System Trap Flags page, click System > Advanced Configuration > Trap Manager > Trap Flags.

Figure 125: System Trap Flags Configuration



The fields available on the System Trap Flags page depends on the packages installed on your system. For example, if your system does not have the BGP4 package installed, the BGP Traps field is not available. [Figure 125: "System Trap Flags Configuration," on page 156](#) and Table 114, "System Trap Flags Fields," on page 157 shows the fields that are available on a system with all packages installed.

**Table 114: System Trap Flags Fields**

Field	Description
Authentication	Enable or disable activation of authentication failure traps by selecting the corresponding line on the pull-down entry field. The factory default is enabled.
Link Up/Down	Enable or disable activation of link status traps by selecting the corresponding line on the pull-down entry field. The factory default is enabled.
Multiple Users	Enable or disable activation of multiple user traps by selecting the corresponding line on the pull-down entry field. The factory default is enabled. This trap is triggered when the same user ID is logged into the switch more than once at the same time (either via telnet or the serial port).
Spanning Tree	Enable or disable activation of spanning tree traps by selecting the corresponding line on the pull-down entry field. The factory default is enabled.
ACL Traps	Enable or disable activation of ACL traps by selecting the corresponding line on the pull-down entry field. The factory default is disabled.

If you make any changes to this page, click Submit to apply the changes to the system.

## 4.19 Managing the DHCP Server

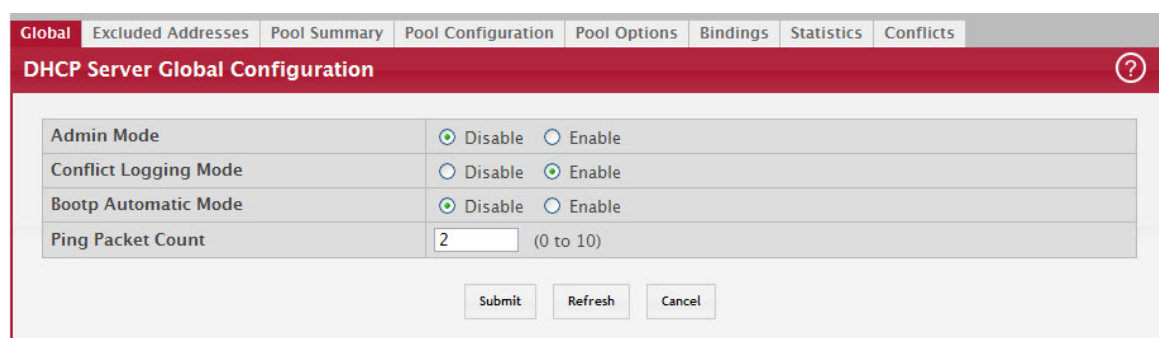
DHCP is generally used between clients (e.g., hosts) and servers (e.g., routers) for the purpose of assigning IP addresses, gateways, and other networking definitions such as DNS, NTP, and/or SIP parameters. The DHCP Server folder contains links to web pages that define and display DHCP parameters and data.

### 4.19.1 DHCP Server Global Configuration

Use the DHCP Server Global Configuration page to configure DHCP global parameters.

To display the page, click System > Advanced Configuration > DHCP Server > Global in the navigation menu.

**Figure 126: DHCP Server Global Configuration**



**Table 115: DHCP Server Global Configuration Fields**

Field	Description
Admin Mode	Enables or disables the DHCP server administrative mode. When enabled, the device can be configured to automatically allocate TCP/IP configurations for clients.
Conflict Logging Mode	Enables or disables the logging mode for IP address conflicts. When enabled, the system stores information IP address conflicts that are detected by the DHCP server.

Table 115: DHCP Server Global Configuration Fields (Continued)

Field	Description
Bootp Automatic Mode	Enables or disables the BOOTP automatic mode. When enabled, the DHCP server supports the allocation of automatic addresses for BOOTP clients. When disabled the DHCP server supports only static addresses for BOOTP clients.
Ping Packet Count	The number of packets the server sends to a pool address to check for duplication as part of a ping operation. If the server receives a response to the ping, the address is considered to be in conflict and is removed from the pool.

- If you change any settings or add an excluded address range, click Submit to apply the changes to the system. Each time you enter a value in the From or To fields, click Submit to add the address or address range to the excluded address list.
- To Delete an address or address range from the excluded address list, select one or more check box beneath the Delete Excluded Addresses field and click Submit.

## 4.19.2 DHCP Server Excluded Addresses

Use the DHCP Server Excluded Addresses page to view and configure the IP addresses that the DHCP server should not assign to clients.

To display the page, click System > Advanced Configuration > DHCP Server > Excluded Addresses in the navigation menu.

Figure 127: DHCP Server Excluded Addresses




Table 116: DHCP Server Excluded Addresses Fields

Field	Description
From	The DHCP Server excludes IP addresses beginning with this IP address.
To	The DHCP Server excludes IP addresses up to this IP address.
VRF Name	The VRF instance for whose DHCP pools this excluded address range is applicable.

### 4.19.2.1 Add Exclusion

Click the Add button on the DHCP Server Excluded Addresses page to open the Add Exclusion page to add one or more IP addresses to exclude. Specify the IPv4 address or range of addresses and the VRF name in the available fields.

Figure 128: Add Exclusion

From	<input type="text" value="10.0.0.1"/>	(x.x.x.x)
To	<input type="text" value="10.0.0.10"/>	(x.x.x.x)(0.0.0.0 to exclude single address)
VRF Name	<input type="text" value="vrt1-test"/>	(0 to 15 characters)

Table 117:

Parameter	Description
From	The IP address to exclude. In a range of addresses, this value is the lowest address to exclude.
To	The highest address to exclude in a range of addresses. If the excluded address is not part of a range, this field shows the same value as the From field. When adding a single IP address to exclude, you can enter the same address specified in the From field, or leave the field with the default value.
VRF Name	The name that identifies the VRF instance associated with the excluded IP address.

**NOTICE**

Applied changes update the device running configuration. If you want the device to retain these changes across a reset, you must first save the configuration.

### 4.19.3 DHCP Server Pool Summary

Use the DHCP Server Pool Summary page to view and configure the DHCP server pools. A DHCP server pool is a set of network configuration information available to DHCP clients that request the information. To display the page, click System > Advanced Configuration > DHCP Server > Pool Summary in the navigation menu.

Figure 129: DHCP Server Pool Summary

Name	Type	Network	Lease Time	VRF Name
p1	Dynamic	1.1.1.0	0 days, 1 hours, 0 mins	vrt2-test

Table 118: DHCP Server Pool Summary Fields

Field	Description
Name	Shows the names of all the existing DHCP server pools.
Type	Displays the type of binding for the pool. <ul style="list-style-type: none"> <li>Manual: The DHCP server assigns a specific IP address to the client based on the client's MAC address. This type is also known as Static.</li> <li>Dynamic: The DHCP server can assign the client any available IP address within the pool. This type is also known as Automatic.</li> </ul>
Network	<ul style="list-style-type: none"> <li>For a Manual pool, indicates the host IP address to assign the client.</li> <li>For a Dynamic pool, indicates the network portion of the IP address. A DHCP client can be offered any available IP address within the defined network as long as it has not been configured as an excluded address.</li> </ul>
Lease Time	The amount of time the information the DHCP server allocates is valid.
VRF Name	The name that identifies the VPN Routing and Forwarding (VRF) instance associated with the DHCP server pool. By default, the address pools are associated with the default VRF. See <a href="#">Section 4.19.3.1: "Add DHCP Server Pool"</a> .

### 4.19.3.1 Add DHCP Server Pool

Click the Add button on the DHCP Server Pool Summary page to open the Add DHCP Server Pool dialog box and configure the DHCP pool settings.

Figure 130: Add DHCP Server Pool

Field	Description
Pool Name	(1 to 31 characters)
Type of Binding	<input checked="" type="radio"/> Dynamic <input type="radio"/> Manual
Network Base Address	(x.x.x.x)
Network Mask	(x.x.x.x)
Client Name	(0 to 255 characters)
Hardware Address Type	Ethernet
Hardware Address	(xx:xx:xx:xx:xx:xx)
Client ID Type	<input type="radio"/> Char <input type="radio"/> HEX
Client ID	(0 to 255 characters)
Client ID HEX	(xx:xx:xx:xx:xx:xx)
Host IP Address	(x.x.x.x)
Host Mask	(x.x.x.x)
Lease Expiration	<input type="radio"/> Enable <input checked="" type="radio"/> Disable
Lease Duration	Days 0 Hours 0 Minutes 0
VRF Name	(0 to 15 characters)
Default Router Address	(x.x.x.x)
DNS Server Address 1	(x.x.x.x)
DNS Server Address 2	(x.x.x.x)

Table 119: Add DHCP Server Pool Fields

Field	Description
Pool Name	For a user with read/write permission, this field shows the names of all the existing pools along with an additional option <b>Create</b> . When the user selects <b>Create</b> , another text box, Pool Name, appears where the user may enter name for the pool to be created. The Pool Name is 1 to 31 characters. For a user with read-only permission, this field shows names of the existing pools only.
Type of Binding	Specifies the type of binding for the pool. <ul style="list-style-type: none"> <li>Dynamic: The DHCP server can assign the client any available IP address within the pool. This type is also known as Automatic.</li> <li>Manual: The DHCP server assigns a specific IP address to the client based on the client's MAC address. This type is also known as Static.</li> </ul> <b>NOTE:</b> The binding type you select determines the fields that are available to configure.
Network Base Address	(Dynamic pools only). The network portion of the IP address. A DHCP client can be offered any available IP address within the defined network as long as it has not been configured as an excluded address.
Network Mask	(Dynamic pools only). The subnet mask associated with the Network Base Address that separates the network bits from the host bits.

Table 119: Add DHCP Server Pool Fields (Continued)

Field	Description
Client Name	This field is optional. (Manual pools only). The system name of the client. The Client Name should not include the domain name.
Hardware Address Type	(Manual pools only). The protocol type used by the client's hardware platform of the DHCP client. Valid types are <b>Ethernet</b> and <b>IEEE802</b> . The default value is <b>Ethernet</b> . This value is used in response to requests from BOOTP clients.
Hardware Address	(Manual pools only). Specifies the MAC address of the hardware platform of the DHCP client.
Client ID Type	Select the option to designate Char or HEX Client ID Type.
Client ID	(Manual pools only) The value some DHCP clients send in the Client Identifier field of DHCP messages. This value is typically identical to the Hardware Address value. In some systems, such as Microsoft DHCP clients, the client identifier is required instead of the hardware address. If the client's DHCP request includes the client identifier, the Client ID field on the DHCP server must contain the same value, and the Hardware Address Type field must be set to the appropriate value. Otherwise, the DHCP server will not respond to the client's request.
Client ID HEX	Enter the hexadecimal number for the Client ID.
Host IP Address	(Manual pools only). The IP address to offer the client.
Host Mask	(Manual pools only). The subnet mask to be statically assigned to a DHCP client.
Lease Expiration	Indicates whether the information the server provides to the client should expire. <ul style="list-style-type: none"> <li>• Enable – Allows the lease to expire. If you select this option, you can specify the amount of time the lease is valid in the Lease Duration field.</li> <li>• Disable – Sets an infinite lease time. For Dynamic bindings, an infinite lease time implies a lease period of 60 days. For a Manual binding, an infinite lease period never expires.</li> </ul>
Lease Duration	The number of Days, Hours, and Minutes the lease is valid. This field cannot be configured if the Lease Expiration mode is disabled.
VRF Name	The name that identifies the VRF instance associated with the DHCP server pool. Enter the VRF name from 1 to 64 characters.
Default Router Address	This field is optional. The IP address of the router to which the client in the pool should send traffic. The default router should be in the same subnet as the client. To add additional default routers, use the <a href="#">DHCP Server Pool Configuration</a> page.
DNS Server Address	This field is optional. The IP addresses of up to two DNS servers the client in the pool should use to resolve host names into IP addresses. To add additional DNS servers, use the <a href="#">DHCP Server Pool Configuration</a> page.

**NOTICE**

Applied changes update the device running configuration. If you want the device to retain these changes across a reset, you must first save the configuration.

### 4.19.3.2 Remove DHCP Server Pool

To remove a pool, select each entry on the DHCP Server Pool Summary page to delete and click the Remove button.

### 4.19.4 DHCP Server Pool Configuration

Use the DHCP Server Pool Configuration page to edit pool settings or to configure additional settings for existing manual and dynamic pools. The additional settings on this page are considered advanced parameters because they are not typically used or configured. The fields that can be configured depend on the type of binding that is selected. The fields that do not apply to the selected binding type are disabled.



Figure 131: DHCP Server Pool Configuration

If you select `Automatic` or `Manual` from the `Type of Binding` drop-down menu, the screen refreshes and a slightly different set of fields appears.

Table 120: DHCP Server Pool Configuration Fields

Field	Description
Pool Name	For a user with read/write permission, this field would show names of all the existing pools along with an additional option <code>Create</code> . When the user selects <code>Create</code> , another text box, <code>Pool Name</code> , appears where the user may enter name for the Pool to be created. For a user with read-only permission, this field would show names of the existing pools only.
Type of Binding	Specifies the type of binding for the pool. <ul style="list-style-type: none"> <li>• <code>Unallocated</code>: The addresses are not assigned to a client.</li> <li>• <code>Automatic</code>: The IP address is automatically assigned to a client by the DHCP server.</li> <li>• <code>Manual</code>: You statically assign an IP address to a client based on the client's MAC address.</li> </ul>
Network Base Address	(Dynamic pools only) The network portion of the IP address. A DHCP client can be offered any available IP address within the defined network as long as it has not been configured as an excluded address.
Network Mask	For dynamic bindings, this field specifies the subnet mask for a DHCP address of a dynamic pool. You can enter a value in <code>Network Mask</code> or <code>Prefix Length</code> to specify the subnet mask, but do not enter a value in both fields.

Table 120: DHCP Server Pool Configuration Fields (Continued)

Field	Description
Client Name	For manual bindings, this field specifies a name for the client to which the DHCP server will statically assign an IP address. This field is optional.
Hardware Address Type	For manual bindings, this field specifies the protocol of the hardware platform of the DHCP client. Valid types are ethernet and ieee802. Default value is ethernet.
Hardware Address	For manual bindings, this field specifies the MAC address of the hardware platform of the DHCP client.
Client ID	(Manual pools only) The value some DHCP clients send in the Client Identifier field of DHCP messages. This value is typically identical to the Hardware Address value. In some systems, such as Microsoft DHCP clients, the client identifier is required instead of the hardware address. If the client's DHCP request includes the client identifier, the Client ID field on the DHCP server must contain the same value, and the Hardware Address Type field must be set to the appropriate value. Otherwise, the DHCP server will not respond to the client's request.
Host IP Address	(Manual pools only) The IP address to offer the client.
Host Mask	For manual bindings, this field specifies the subnet mask to be statically assigned to a DHCP client. You can enter a value in Host Mask or Prefix Length to specify the subnet mask, but do not enter a value in both fields.
Lease Expiration	Indicates whether the information the server provides to the client should expire. <ul style="list-style-type: none"> <li>• Enable – Allows the lease to expire. If you select this option, you can specify the amount of time the lease is valid in the Lease Duration field.</li> <li>• Disable – Sets an infinite lease time. For Dynamic bindings, an infinite lease time implies a lease period of 60 days. For a Manual binding, an infinite lease period never expires.</li> </ul>
Lease Duration	<ul style="list-style-type: none"> <li>• The number of Days, Hours, and Minutes the lease is valid. This field cannot be configured if the Lease Expiration is disabled.</li> </ul>
VRF Name	The name that identifies the VRF instance associated with the DHCP server pool. Enter the VRF name from 1 to 64 characters.
Next Server Address	<p>The IP address of the next server the client should contact in the boot process. For example, the client might be required to contact a TFTP server to download a new image file. To configure this field, click the Edit icon in the row. To reset the field to the default value, click the Reset icon in the row.</p> <p>To configure settings for one or more default routers, DNS servers, or NetBIOS servers that can be used by the client(s) in the pool, use the buttons available in the appropriate table to perform the following tasks:</p> <ul style="list-style-type: none"> <li>• To add an entry to the server list, click the + (plus) button and enter the IP address of the server to add.</li> <li>• To edit the address of a configured server, click the Edit icon associated with the entry to edit and update the address.</li> <li>• To delete an entry from the list, click the – (minus) button associated with the entry to remove.</li> <li>• To delete all entries from the list, click the – (minus) button in the heading row.</li> </ul>
Default Router	Lists the IP address of each router to which the client(s) in the pool should send traffic. The default router should be in the same subnet as the client.
DNS Server	Lists the IP address of each DNS server the client(s) in the pool can contact to perform address resolution.
NetBIOS Server	Lists the IP address of each NetBIOS Windows Internet Naming Service (WINS) name server that is available for the selected pool.

- After you configure values for the DHCP address pool, click Submit to create the pool and apply the changes to the system.
- To delete a pool, select the pool from the Pool Name drop-down menu and click Delete.

#### 4.19.5 DHCP Server Pool Options

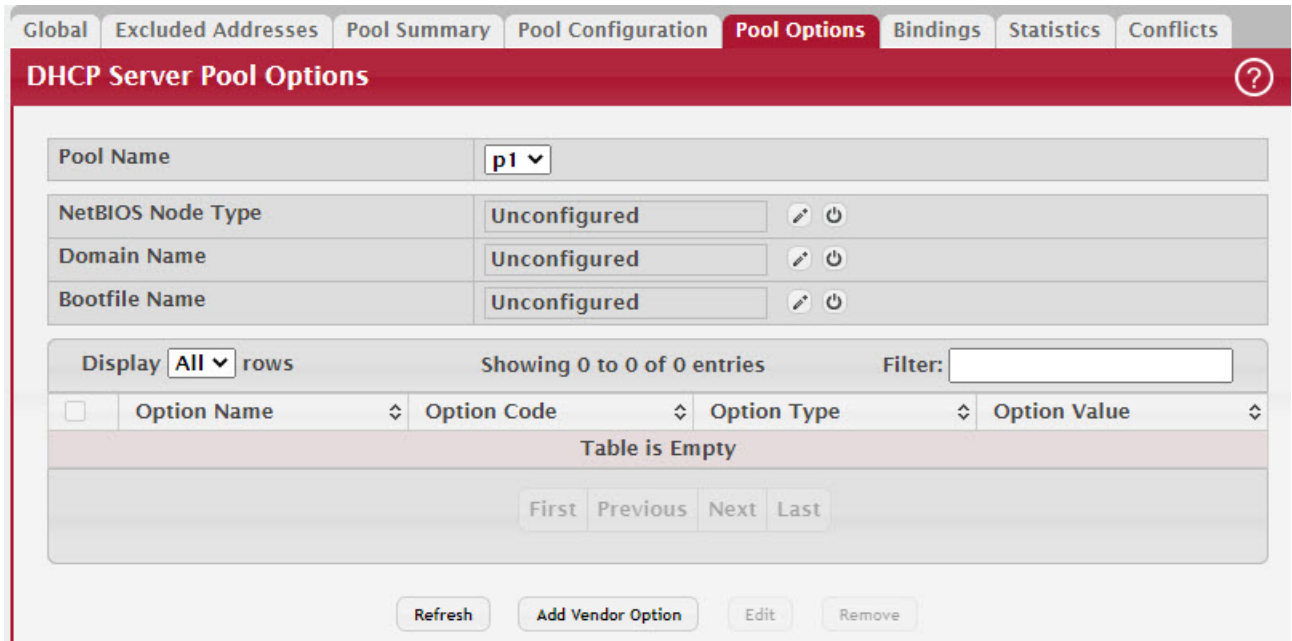
Use the DHCP Server Pool Options page to configure additional DHCP pool options, including vendor-defined options. DHCP options are collections of data with type codes that indicate how the options should be used. When a client broad-

casts a request for information, the request includes the option codes that correspond to the information the client wants the DHCP server to supply.

To access the DHCP Server Pool Options page, click System > Advanced Configuration > DHCP Server > Pool Options in the navigation menu.

If no DHCP pools exist, the DHCP Server Pool Options page does not display the fields shown in [Figure 132: "DHCP Server Pool Options,"](#) on page 165.

Figure 132: DHCP Server Pool Options



If any DHCP pools are configured on the system, the DHCP Server Pool Options page contains the following fields.

Table 121: DHCP Server Pool Options Fields

Field	Description
Pool Name	Select the DHCP pool to configure. The menu includes all pools that are configured on the device.
NetBIOS Node Type	The method the client should use to resolve NetBIOS names to IP addresses. To configure this field, click the <b>Edit</b> icon in the row. To reset the field to the default value, click the Reset icon in the row. The options are: <ul style="list-style-type: none"> <li>• B-Node Broadcast: Broadcast only.</li> <li>• P-Node Peer-to-Peer: NetBIOS name server only.</li> <li>• M-Node Mixed: Broadcast, then NetBIOS name server.</li> <li>• H-Node Hybrid: NetBIOS name server, then broadcast.</li> </ul>
Domain Name	The default domain name to configure for all clients in the selected pool.
Bootfile Name	The name of the default boot image that the client should attempt to download from a specified boot server.

The option table shows the Vendor Options that have been added to the selected pool. Use the buttons to perform the following tasks:

- To add a vendor option, click Add Vendor Option and configure the desired information in the available fields.
- To edit a vendor option, select the entry to change and click Edit.
- To remove a vendor option, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 122: Add Vendor Options Fields

Fields	Description
Option Name	Identifies whether the entry is a fixed option or a vendor-defined option (Vendor).
Option Code	The number that uniquely identifies the option.
Option Type	Specifies the type of data to associate with the Option Code configured for the selected pool. The possible values are as follows: <ul style="list-style-type: none"> <li>• ASCII: The option type is a text string.</li> <li>• HEX: The option type is a hexadecimal number.</li> <li>• IP Address: The option type is an IP address.</li> </ul>
Option Value	The data associated with the Option Code. When adding or editing a vendor option, the fields available for configuring the value depend on the selected Option Type. If the value you configure contains characters that are not allowed by the selected Option Type, the configuration cannot be applied.

**NOTICE**

Applied changes update the device running configuration. If you want the device to retain these changes across a reset, you must first save the configuration.

## 4.19.6 DHCP Server Bindings

Use the DHCP Server Bindings page to view information about the IP address bindings in the DHCP server database.

To access the DHCP Server Bindings page, click System > Advanced Configuration > DHCP Server > Bindings in the navigation menu.

Figure 133: DHCP Server Bindings

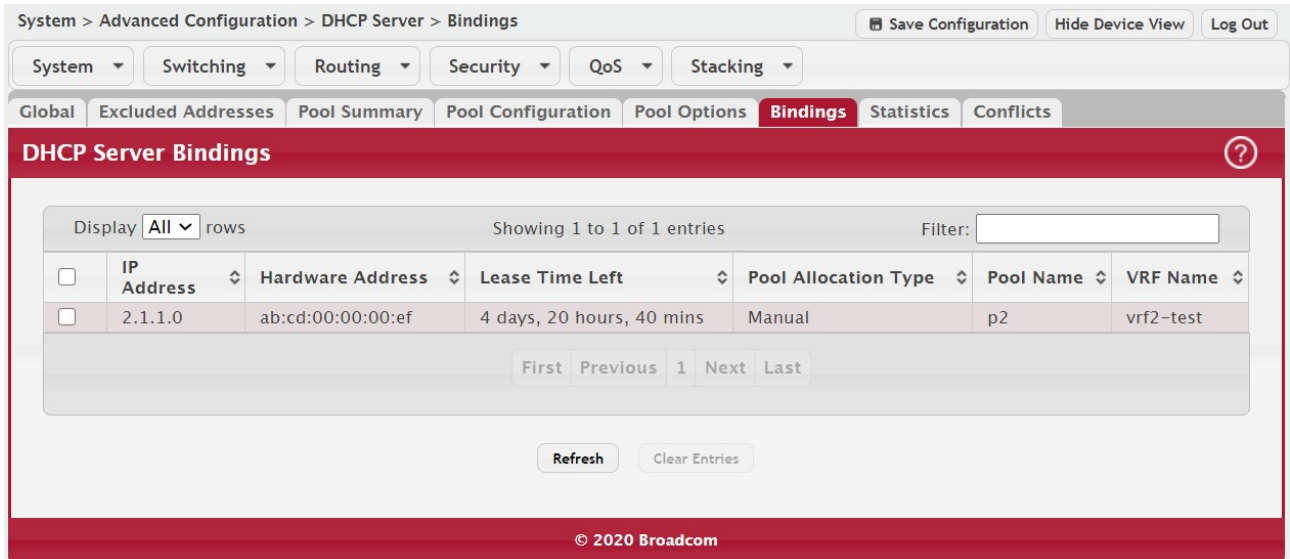


Table 123: DHCP Server Bindings Fields

Field	Description
IP Address	The IP Address of the DHCP client.
Hardware Address	The MAC address of the DHCP client.
Lease Time Left	The amount of time left until the lease expires in days, hours, and minutes.
Pool Allocation Type	The type of binding used: <ul style="list-style-type: none"> <li>Dynamic – The address was allocated dynamically from a pool that includes a range of IP addresses.</li> <li>Manual – A static IP address was assigned based on the MAC address of the client.</li> <li>Inactive – The pool is not in use.</li> </ul>
Clear Entries (Button)	To remove an entry from the table, select each entry to delete and click Clear Entries. You must confirm the action before the binding is deleted.
Pool Name	The name that identifies the DHCP server pool associated with the binding.
VRF Name	The name that identifies the VRF instance associated with the DHCP server pool.

If you change any settings, click Submit to apply the changes to the system.

## 4.19.7 DHCP Server Statistics

Use the DHCP Server Statistics page to view information about the DHCP server bindings and messages. To access the DHCP Server Statistics page, click System > Advanced Configuration > DHCP Server > Server Statistics in the navigation menu.

Figure 134: DHCP Server Statistics

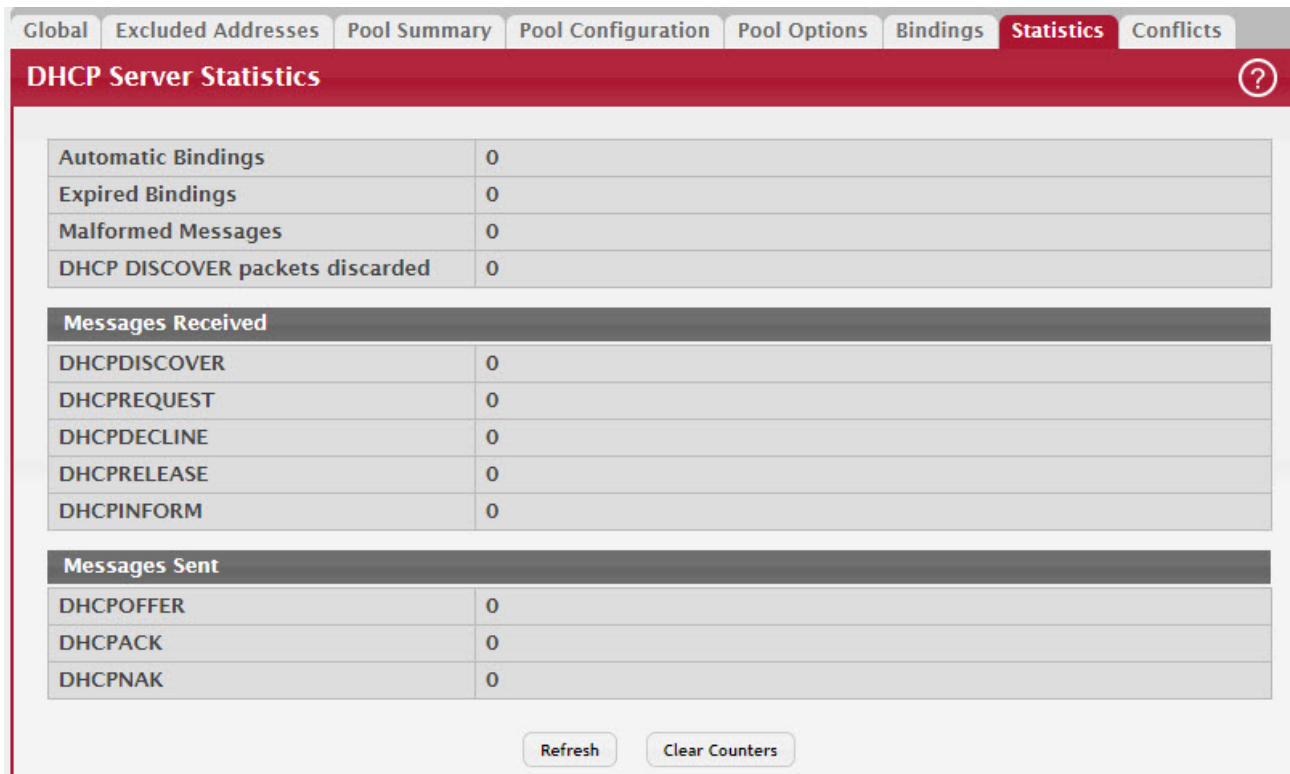


Table 124: DHCP Server Statistics Fields

Field	Description
Automatic Bindings	Shows the total number of IP addresses from all address pools with automatic bindings that the DHCP server has assigned to DHCP clients.
Expired Bindings	Shows the number of IP addresses that the DHCP server has assigned to DHCP clients that have exceeded the configured lease time.
Malformed Messages	Shows the number of messages received from one or more DHCP clients that were improperly formatted.
DHCP DISCOVER packets discarded	The number of messages discarded from one or more DHCP Discovers.
Message Received	
DHCPDISCOVER	Shows the number of DHCP discovery messages the DHCP server has received. A DHCP client broadcasts this type of message to discover available DHCP servers.
DHCPREQUEST	Shows the number of DHCP request messages the DHCP server has received. A DHCP client broadcasts this type of message in response to a DHCP offer message it received from a DHCP server.
DHCPDECLINE	Shows the number of DHCP decline messages the DHCP server has received from clients. A client sends a decline message if the DHCP client detects that the IP address offered by the DHCP server is already in use on the network. The server then marks the address as unavailable.
DHCPRELEASE	Shows the number of DHCP release messages the DHCP server has received from clients. This type of message indicates that a client no longer needs the assigned address.
DHCPINFORM	Shows the number of DHCP inform messages the DHCP server has received from clients. A client uses this type of message to obtain DHCP options.
Message Sent	

Table 124: DHCP Server Statistics Fields (Continued)

Field	Description
DHCPOFFER	The number of DHCP offer messages the DHCP server has sent to DHCP clients in response to DHCP discovery messages it has received.
DHCPACK	The number of DHCP acknowledgment messages the DHCP server has sent to DHCP clients in response to DHCP request messages it has received. The server sends this message after the client has accepted the offer from this particular server. The DHCP acknowledgment message includes information about the lease time and any other configuration information that the DHCP client has requested.
DHCPNAK	The number of negative DHCP acknowledgment messages the DHCP server has sent to DHCP clients. A server might send this type of message if the client requests an IP address that is already in use or if the server refuses to renew the lease.

- Click Refresh to update the information on the screen.
- Click Clear Server Statistics to reset all counters to zero.

### 4.19.8 DHCP Server Conflicts Information

Use the DHCP Server Conflicts Information page to view information on hosts that have address conflicts; i.e., when the same IP address is assigned to two or more devices on the network.

To access the DHCP Server Conflicts Information page, click System > Advanced Configuration > DHCP Server > Conflicts Information in the navigation menu.

Figure 135: DHCP Server Conflicts Information

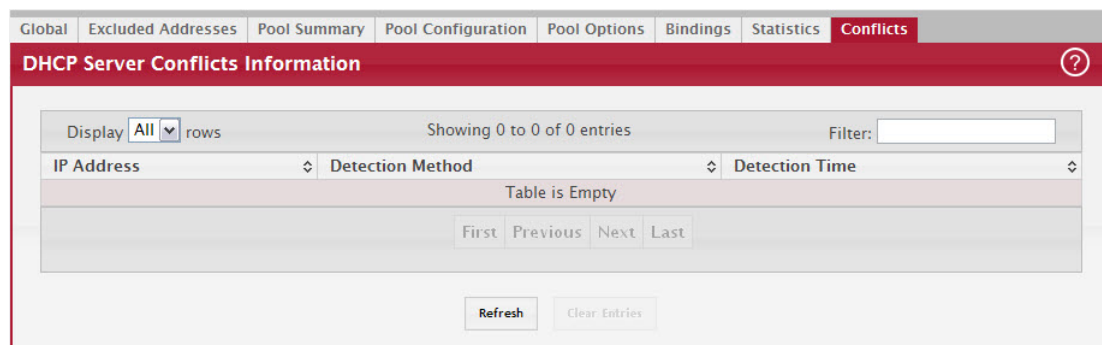


Table 125: DHCP Server Conflicts Information Fields

Field	Description
IP Address	The IP address that has been detected as a duplicate.
Detection Method	The method used to detect the conflict, which is one of the following: <ul style="list-style-type: none"> <li>• Gratuitous ARP – The DHCP client detected the conflict by broadcasting an ARP request to the address specified in the DHCP offer message sent by the server. If the client receives a reply to the ARP request, it declines the offer and reports the conflict.</li> <li>• Ping – The server detected the conflict by sending an ICMP echo message (ping) to the IP address before offering it to the DHCP client. If the server receives a response to the ping, the address is considered to be in conflict and is removed from the pool.</li> <li>• Host Declined – The server received a DHCPDECLINE message from the host. A DHCPDECLINE message indicates that the host has discovered that the IP address is already in use on the network.</li> </ul>
Detection Time	The time when the conflict was detected in days, hours, minutes and seconds since the system was last reset (i.e., system up time).
Clear Entries (Button)	Clears all of the address conflict entries.

## 4.20 Configuring Time Ranges

You can use these pages to configure time ranges to use in time-based access control list (ACL) rules. Time-based ACLs allow one or more rules within an ACL to be based on a periodic or absolute time. Each ACL rule within an ACL except for the implicit deny all rule can be configured to be active and operational only during a specific time period. The time range pages allow you to define specific times of the day and week to implement time-based ACLs. The time range is identified by a name and can then be referenced by an ACL rule defined within an ACL.

### 4.20.1 Time Range Summary

Use this page to create a named time range. Each time range can consist of one absolute time entry and/or one or more periodic time entries.

To access this page, click System > Advanced Configuration > Time Ranges > Configuration.

Figure 136: Time Range Summary

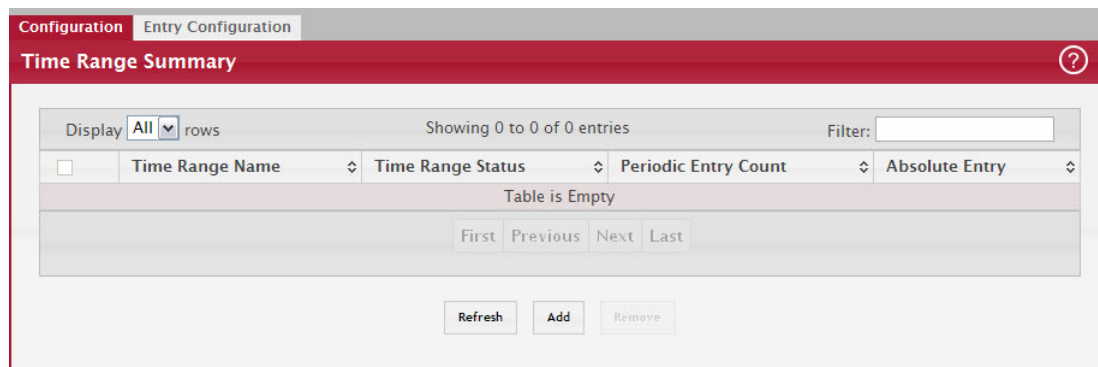


Table 126: Time Range Summary

Field	Description
Time Range Name	The unique ID or name that identifies this time range. A time-based ACL rule can reference the name configured in this field.
Time Range Status	Shows whether the time range is Active or Inactive. A time range is Inactive if the current day and time do not fall within any time range entries configured for the time range.
Periodic Entry Count	The number of periodic time range entries currently configured for the time range.
Absolute Entry	Shows whether an absolute time entry is currently configured for the time range.

Use the buttons to perform the following tasks:

- To add a time range, click Add and configure a name for the time range configuration.
- To delete a configured time range, select each entry to delete, click Remove, and confirm the action.
- Use Submit to add a new time range.

### 4.20.2 Time Range Entry Summary

Use this page to configure periodic and absolute time range entries and add them to named time ranges.

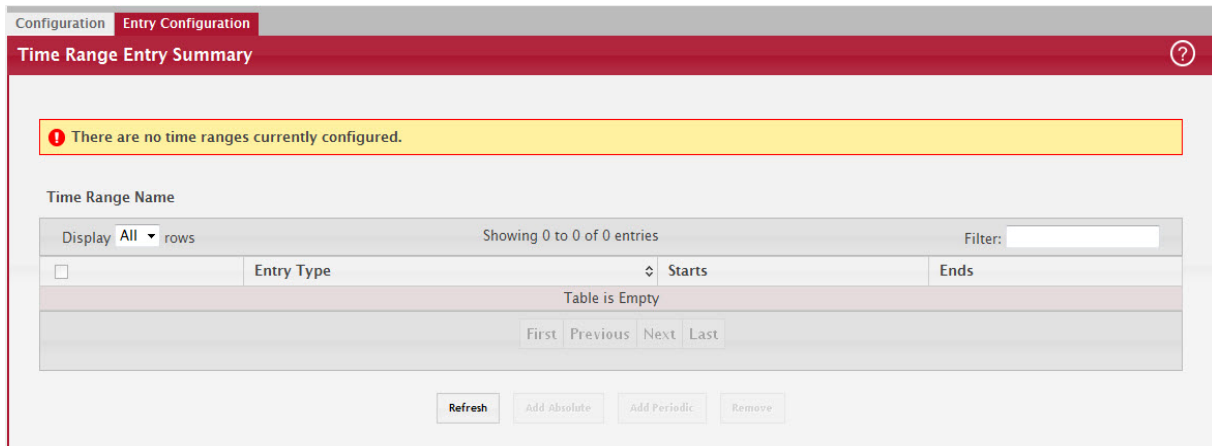
#### NOTICE

The time range entries use the system time for the time periods in which they take effect. Make sure you configure the SNTP server settings so that the SNTP client on the switch can obtain the correct date and time from the server.

To access this page, click System > Advanced Configuration > Time Ranges > Entry Configuration.



Figure 137: Time Range Entry Summary



To configure the time range entries for a time range configuration, select the time range configuration from the Time Range Name menu and use the buttons to perform the following tasks:

- To add an Absolute time range entry, click Add Absolute and configure information about when the Absolute entry occurs. If the Add Absolute button is not available, an Absolute entry already exists for the selected time range configuration.
- To add a Periodic time range entry, click Add Periodic and specify the days and times that the entry is in effect.
- To delete a time range entry, select each entry to delete, click Remove, and confirm the action.

Table 127: Time Range Entry Summary

Field	Description
Time Range Name	Select the name of the time range to which you want to add a time range entry.
Time Range Entry	Select Create New Time Range Entry to add a new entry to a time range. To view or delete an existing time range entry, select its ID from the menu.
Time Range Entry ID	When creating a new time range entry, assign a unique ID number from 1–10. This field does not appear if the entry has already been configured.
Time Range Entry Type.	Specifies whether the entry is periodic or absolute. A periodic entry occurs at the same time every day or on one or more days of the week. An absolute entry does not repeat.
<b>Periodic Time Range Entry</b>	
Applicable Days	Specify the day(s) when the time entry occurs: <ul style="list-style-type: none"> <li>• Daily–Has the same start and end time every day</li> <li>• Weekdays–Has the same start and end time Monday through Friday</li> <li>• Weekdays–Has the same start and end time on Saturday and Sunday</li> <li>• Days of the Week–Select the day of the week when the entry starts and stops. You do not need to use the same day of the week for the start and end time.</li> </ul>
Start Day	(Periodic Days of Week only) Select the day the time range entry starts. To select multiple days, hold the CTRL key and click the days.
Start Time	Specify the time when the entry begins. The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.
End Day	(Periodic Days of Week only) Select the day the time range entry ends.
End Time	Specify the time when the entry ends. The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.
<b>Absolute Time Range Entry</b>	
Absolute Start Date and Time	Select the check box to configure the date and time when the time range entry begins.
Start Month	Select the month when the time entry begins.

**Table 127: Time Range Entry Summary (Continued)**

Field	Description
Start Date	Select the day of the month when the time entry begins.
Start Year	Select the year when the time entry begins.
Start Time	Specify the time when the entry begins. The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.
Absolute End Date and Time	Select the check box to configure the date and time when the time range entry ends.
End Month	Select the month when the time entry ends.
End Date	Select the day of the month when the time entry ends.
End Year	Select the year when the time entry ends.
End Time	Specify the time when the entry ends. The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.

Click Submit to create the time range entry. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a save is performed.

## 4.21 Configuring DNS

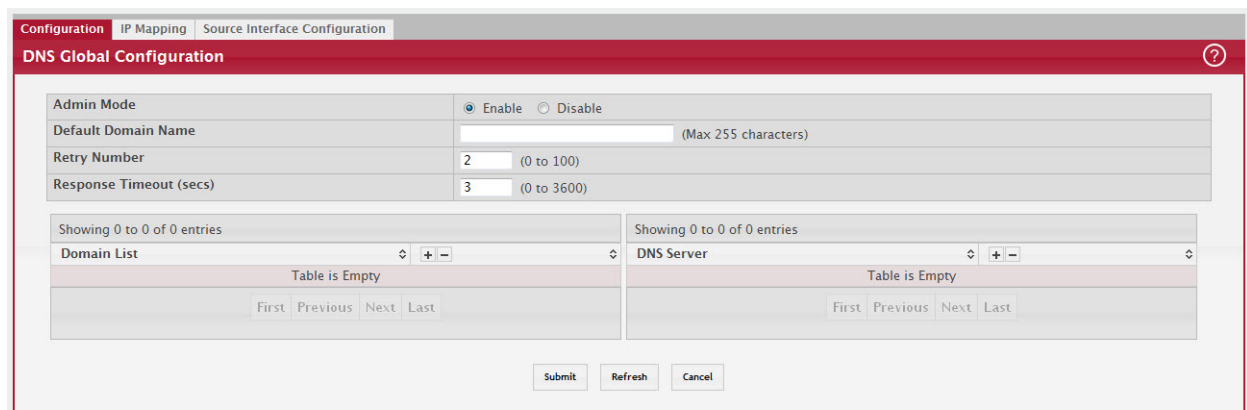
You can use these pages to configure information about DNS servers the network uses and how the switch/router operates as a DNS client.

### 4.21.1 Global Configuration

Use this page to configure global DNS settings and to view DNS client status information.

To access this page, click System > Advanced Configuration > DNS > Configuration.

**Figure 138: DNS Global Configuration**



**Table 128: DNS Global Configuration Fields**

Field	Description
Admin Mode	Select <i>Enable</i> or <i>Disable</i> from the pull-down menu to set the administrative status of DNS Client. The default is <i>Disable</i> .
Default Domain Name	Enter the default domain name for DNS client messages. The name should be no longer than 255 characters. When the system is performing a lookup on an unqualified hostname, this field is provided as the domain name (e.g., if default domain name is.com and the user enters hotmail, then hotmail is changed to hotmail.com to resolve the name). By default, no default domain name is configured in the system.

Table 128: DNS Global Configuration Fields (Continued)

Field	Description
Retry Number	Enter the number of times to retry sending DNS queries. The valid values are from 0 to 100. The default value is 2.
Response Timeout	Enter the number of seconds to allow a DNS server to respond to a request before issuing a retry. Valid values are 0 to 3600. The default value is 3.
Domain List	Enter a domain list to define the domain to use when performing a lookup on an unqualified hostname. Each name must be no more than 256 characters. Multiple default domain names can be configured using the default domain-name list. If there is no domain list, the default domain name configured is used.

- If you change any settings, click Submit to send the information to the system.
- To create a new list of domain names, click Create. Then enter a name of the list and click submit. Repeat this step to add multiple domains to the default domain list.
- To remove a domain from the default list select the Remove option next to the item you want to remove and click Submit.

#### 4.21.2 DNS IP Mapping Configuration

Use this page to configure DNS host names for hosts on the network and to view dynamic DNS entries. The host names are associated with IPv4 or IPv6 addresses on the network, which are statically assigned to particular hosts.

To access this page, click System > Advanced Configuration > DNS > HostName IP Mapping Summary in the menu.

Figure 139: DNS IP Mapping Summary

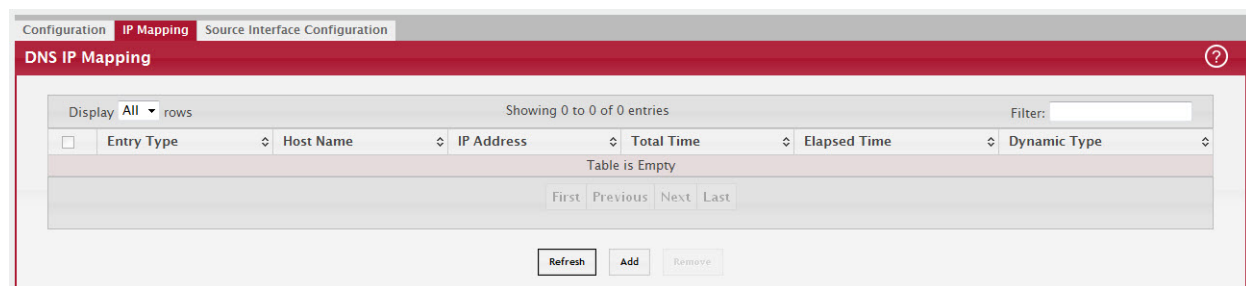


Table 129: DNS IP Mapping Summary Fields

Field	Description
DNS Static Entries	
Entry Type	Type of DNS entry: <ul style="list-style-type: none"> <li>• Static – An entry that has been manually configured on the device.</li> <li>• Dynamic – An entry that the device has learned by using a configured DNS server to resolve a hostname.</li> </ul>
Host Name	The name that identifies the system. For Static entries, specify the Host Name after you click Add.
IP Address	The IPv4 or IPv6 address associated with the configured Host Name. For Static entries, specify the IP Address after you click Add. You can specify either an IPv4 or an IPv6 address.
DNS Dynamic Entries	
Total Time	The number of seconds that the entry will remain in the table.

Table 129: DNS IP Mapping Summary Fields (Continued)

Field	Description
Elapsed Time	The number of seconds that have passed since the entry was added to the table. When the Elapsed Time reaches the Total Time, the entry times out and is removed from the table.
Dynamic Type	The type of address in the entry, for example IP or (less common) X.121.

### Command Buttons

The page includes the following command buttons:

- Click Add Static Entry to load the DNS IP Mapping Summary page to configure the Host Name IP Mapping entries.
- Click Submit to apply the new configuration and cause the change to take effect immediately. These changes will not be retained across a power cycle unless a Save is performed.
- Click Clear Dynamic Entries to remove all DNS IP Mapping Summary entries. A confirmation prompt will be displayed. Click the button to confirm removal and the DNS Host Name IP Mapping Summary dynamic entries are cleared.
- Click Refresh to refresh the page with the most current data from the switch.

If you click Add, the Add DNS Entry page appears.

Figure 140: Add DNS Entry

The screenshot shows a dialog box titled "Add DNS Entry". It has a close button (X) in the top right corner. Inside the dialog, there are two input fields. The first is labeled "Host Name" and has a placeholder text "(1 to 255 characters)". The second is labeled "IP Address" and has a placeholder text "(x.x.x.x or x::x::x::x::x)". At the bottom right of the dialog, there are two buttons: "Submit" and "Cancel".

Table 130: Add DNS Entry Fields

Field	Description
Host Name	Enter the host name to assign to the static entry.
IP Address	Enter the IP4 or IPv6 address associated with the host name.

### Command Buttons

The page includes the following command buttons:

- Click Submit to apply the new configuration and cause the change to take effect immediately. These changes will not be retained across a power cycle unless a Save is performed.
- Click Cancel to cancel and redisplay the hostname IP mapping page to see the configured hostname-IP mapping entries.

### 4.21.3 DNS Source Interface Configuration

Use this page to specify the physical or logical interface to use as the DNS client source interface. When an IP address is configured on the source interface, this address is used for all DNS communications between the local DNS client and the remote DNS server. The IP address of the designated source interface is used in the IP header of DNS management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access the DNS Source Interface Configuration page, click System > Advanced Configuration > DNS > Source Interface Configuration in the menu.

Figure 141: DNS Source Interface Configuration

Table 131: DNS Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> <li>None – The primary IP address of the originating (outbound) interface is used as the source address.</li> <li>Interface – The primary IP address of a physical port is used as the source address.</li> <li>Loopback – The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up.</li> <li>VLAN – The primary IP address of a VLAN routing interface is used as the source address.</li> </ul>
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.

If you change any of the settings on the page, click Submit to apply the changes to system.

## 4.22 Configuring SNTP Settings

FASTPATH software supports the Simple Network Time Protocol (SNTP). SNTP assures accurate network device clock time synchronization up to the millisecond. Time synchronization is performed by a network SNTP server. FASTPATH software operates only as an SNTP client and cannot provide time services to other systems.

Time sources are established by Stratum. Stratum defines the accuracy of the reference clock. The higher the stratum (where zero is the highest), the more accurate the clock. The device receives time from stratum 1 and above since it is itself a stratum 2 device.

The following is an example of stratum:

- Stratum 0: A real time clock is used as the time source, for example, a GPS system.
- Stratum 1: A server that is directly linked to a Stratum 0 time source is used. Stratum 1 time servers provide primary network time standards.
- Stratum 2: The time source is distanced from the Stratum 1 server over a network path. For example, a Stratum 2 server receives the time over a network link, via NTP, from a Stratum 1 server.

Information received from SNTP servers is evaluated based on the time level and server type.

SNTP time definitions are assessed and determined by the following time levels:

- T1: Time at which the original request was sent by the client.
- T2: Time at which the original request was received by the server.
- T3: Time at which the server sent a reply.
- T4: Time at which the client received the server's reply.

The device can poll Unicast and Broadcast server types for the server time.

Polling for Unicast information is used for polling a server for which the IP address is known. SNTP servers that have been configured on the device are the only ones that are polled for synchronization information. T1 through T4 are used to determine server time. This is the preferred method for synchronizing device time because it is the most secure method. If this method is selected, SNTP information is accepted only from SNTP servers defined on the device using the SNTP Server Configuration page.

Broadcast information is used when the server IP address is unknown. When a Broadcast message is sent from an SNTP server, the SNTP client listens to the message. If Broadcast polling is enabled, any synchronization information is accepted, even if it has not been requested by the device. This is the least secure method.

The device retrieves synchronization information, either by actively requesting information or at every poll interval. If Unicast and Broadcast polling are enabled, the information is retrieved in this order:

- Information from servers defined on the device is preferred. If Unicast polling is not enabled or if no servers are defined on the device, the device accepts time information from any SNTP server that responds.
- If more than one Unicast device responds, synchronization information is preferred from the device with the lowest stratum.
- If the servers have the same stratum, synchronization information is accepted from the SNTP server that responded first.

MD5 (Message Digest 5) Authentication safeguards device synchronization paths to SNTP servers. MD5 is an algorithm that produces a 128-bit hash. MD5 is a variation of MD4, and increases MD4 security. MD5 verifies the integrity of the communication, authenticates the origin of the communication.

### 4.22.1 SNTP Global Configuration

Use the SNTP Global Configuration page to view and adjust SNTP parameters.

To display the SNTP Global Configuration page, click System > Advanced Configuration > SNTP > Global Configuration in the navigation menu.

Figure 142: SNTP Global Configuration

SNTP Global Configuration	
Client Mode	Unicast
Port	None
Unicast Poll Timeout (Seconds)	6 (6 to 10)
Broadcast Poll Interval (Seconds)	6 (6 to 10)
Unicast Poll Timeout (Seconds)	5 (1 to 30)
Unicast Poll Retry	1 (0 to 10)
Number of Servers Configured	1

Table 132: SNTP Global Configuration Fields

Field	Description
Client Mode	Use drop-down list specify the SNTP client mode, which is one of the following modes: <ul style="list-style-type: none"> <li>• Disable: SNTP is not operational. No SNTP requests are sent from the client nor are any received SNTP messages processed.</li> <li>• Unicast: SNTP operates in a point to point fashion. A unicast client sends a request to a designated server at its unicast address and expects a reply from which it can determine the time and, optionally the round-trip delay and local clock offset relative to the server.</li> <li>• Broadcast: SNTP operates in the same manner as multicast mode but uses a local broadcast address instead of a multicast address. The broadcast address has a single subnet scope while a multicast address has Internet wide scope.</li> </ul>
Port	Specifies the local UDP port to listen for responses/broadcasts. Allowed range is (1 to 65535). Default value is 123.
Unicast Poll Interval	Specifies the number of seconds between unicast poll requests expressed as a power of two when configured in unicast mode. Allowed range is (6 to 10). Default value is 6.
Broadcast Poll Interval	Specifies the number of seconds between broadcast poll requests expressed as a power of two when configured in broadcast mode. Broadcasts received prior to the expiry of this interval are discarded. Allowed range is (6 to 10). Default value is 6.
Unicast Poll Timeout	Specifies the number of seconds to wait for an SNTP response when configured in unicast mode. Allowed range is (1 to 30). Default value is 5.
Unicast Poll Retry	Specifies the number of times to retry a request to an SNTP server after the first time-out before attempting to use the next configured server when configured in unicast mode. Allowed range is (0 to 10). Default value is 1.
Number of Servers Configured	Specifies the number of current valid unicast server entries configured for this client.

If you change any of the settings on the page, click Submit to apply the changes to system.

#### 4.22.2 SNTP Global Status

Use the SNTP Global Status page to view information about the system's SNTP client.

To access the SNTP Global Status page, click System > Advanced Configuration > SNTP > Global Status in the navigation menu.

Figure 143: SNTP Global Status

SNTP Global Status	
Version	4
Supported Mode	Unicast and Broadcast
Last Update Time	Dec 31 19:00:00 1969
Last Attempt Time	May 1 18:40:04 1919
Last Attempt Status	Request Timed Out
Server IP Address	
Address Type	Unknown
Server Stratum	0
Reference Clock ID	
Server Mode	Reserved
Unicast Server Max Entries	3
Unicast Server Current Entries	1
Broadcast Count	0

Table 133: SNTP Global Status Fields

Field	Description
Version	Specifies the SNTP Version the client supports.
Supported Mode	Specifies the SNTP modes the client supports. Multiple modes may be supported by a client.
Last Update Time	Specifies the local date and time (UTC) the SNTP client last updated the system clock.
Last Attempt Time	Specifies the local date and time (UTC) of the last SNTP request or receipt of an unsolicited message.
Last Attempt Status	Specifies the status of the last SNTP request or unsolicited message for both unicast and broadcast modes. If no message has been received from a server, a status of Other is displayed. These values are appropriate for all operational modes: <ul style="list-style-type: none"> <li>• Other: None of the following enumeration values.</li> <li>• Success: The SNTP operation was successful and the system time was updated.</li> <li>• Request Timed Out: A directed SNTP request timed out without receiving a response from the SNTP server.</li> <li>• Bad Date Encoded: The time provided by the SNTP server is not valid.</li> <li>• Version Not Supported: The SNTP version supported by the server is not compatible with the version supported by the client.</li> <li>• Server Unsynchronized: The SNTP server is not synchronized with its peers. This is indicated via the 'leap indicator' field on the SNTP message.</li> <li>• Server Kiss Of Death: The SNTP server indicated that no further queries were to be sent to this server. This is indicated by a stratum field equal to 0 in a message received from a server.</li> </ul>
Server IP Address	Specifies the IP address of the server for the last received valid packet. If no message has been received from any server, an empty string is shown.
Address Type	Specifies the address type of the SNTP Server address for the last received valid packet.
Server Stratum	Specifies the claimed stratum of the server for the last received valid packet.
Reference Clock Id	Specifies the reference clock identifier of the server for the last received valid packet.
Server Mode	Specifies the mode of the server for the last received valid packet.
Unicast Sever Max Entries	Specifies the maximum number of unicast server entries that can be configured on this client.
Unicast Server Current Entries	Specifies the number of current valid unicast server entries configured for this client.
Broadcast Count	Specifies the number of unsolicited broadcast SNTP messages that have been received and processed by the SNTP client since last reboot.

Click Refresh to display the latest information from the router.

### 4.22.3 SNTP Server Configuration

Use the SNTP Server Configuration page to view and modify information for adding and modifying Simple Network Time Protocol SNTP servers.

To display the SNTP Server Configuration page, click System > Advanced Configuration > SNTP > Server Configuration in the navigation menu.



Figure 144: SNMP Server Configuration

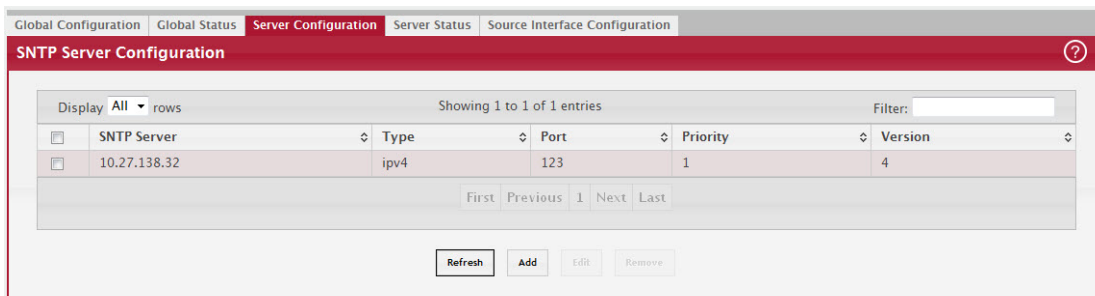


Table 134: SNMP Server Configuration Fields

Field	Description
SNMP Server	Select the IP address of a user-defined SNMP server to view or modify information about an SNMP server, or select Add to configure a new SNMP server. You can define up to three SNMP servers.
Type	Select IPv4 if you entered an IPv4 address, IPv6 if you entered an IPv6 address or DNS if you entered a hostname.
Port	Enter a port number from 1 to 65535. The default is 123.
Priority	Enter a priority from 1 to 3, with 1 being the highest priority. The switch will attempt to use the highest priority server and, if it is not available, will use the next highest server.
Version	Enter the protocol version number.

- To add an SNMP server, select Add from the Server list, complete the remaining fields as desired, and click Submit. The SNMP server is added, and is now reflected in the Server list. You must perform a save to retain your changes over a power cycle.
- To removing an SNMP server, select the IP address of the server to remove from the Server list, and then click Remove. The entry is removed, and the device is updated.

#### 4.22.4 SNMP Server Status

The SNMP Server Status page displays status information about the SNMP servers configured on your switch. To access the SNMP Server Status page, click System > Advanced Configuration > SNMP > Server Status in the navigation menu.

Figure 145: SNMP Server Status

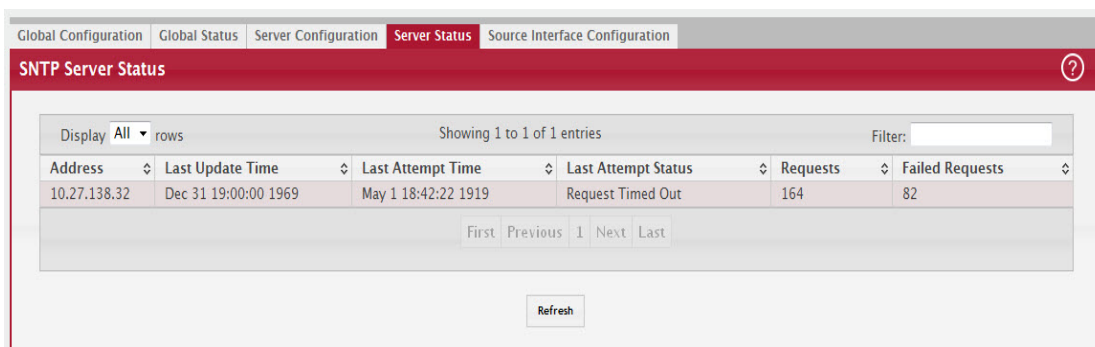


Table 135: SNTP Server Status Fields

Field	Description
Address	Specifies all the existing Server Addresses. If no Server configuration exists, a message saying No SNTP server exists flashes on the screen.
Last Update Time	Specifies the local date and time (UTC) that the response from this server was used to update the system clock.
Last Attempt Time	Specifies the local date and time (UTC) that this SNTP server was last queried.
Last Attempt Status	Specifies the status of the last SNTP request to this server. If no packet has been received from this server, a status of Other is displayed: <ul style="list-style-type: none"> <li>• Other: None of the following enumeration values.</li> <li>• Success: The SNTP operation was successful and the system time was updated.</li> <li>• Request Timed Out: A directed SNTP request timed out without receiving a response from the SNTP server.</li> <li>• Bad Date Encoded: The time provided by the SNTP server is not valid.</li> <li>• Version Not Supported: The SNTP version supported by the server is not compatible with the version supported by the client.</li> <li>• Server Unsynchronized: The SNTP server is not synchronized with its peers. This is indicated via the 'leap indicator' field on the SNTP message.</li> <li>• Server Kiss Of Death: The SNTP server indicated that no further queries were to be sent to this server. This is indicated by a stratum field equal to 0 in a message received from a server.</li> </ul>
Requests	Specifies the number of SNTP requests made to this server since last agent reboot.
Failed Requests	Specifies the number of failed SNTP requests made to this server since last reboot.

Click Refresh to display the latest information from the switch.

#### 4.22.5 SNTP Source Interface Configuration

Use this page to specify the physical or logical interface to use as the SNTP client source interface. When an IP address is configured on the source interface, this address is used for all SNTP communications between the local SNTP client and the remote SNTP server. The IP address of the designated source interface is used in the IP header of SNTP management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access the SNTP Source Interface Configuration page, click System > Advanced Configuration > SNTP > Source Interface Configuration in the navigation menu.

Figure 146: SNTP Source Interface Configuration

Table 136: SNTP Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> <li>None – The primary IP address of the originating (outbound) interface is used as the source address.</li> <li>Interface – The primary IP address of a physical port is used as the source address.</li> <li>Loopback – The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up.</li> <li>VLAN – The primary IP address of a VLAN routing interface is used as the source address.</li> </ul>
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.

Click Refresh to display the latest information from the switch.

Click Submit to apply the settings to the running configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.

## 4.23 Configuring the Time Zone

This page displays information about the current system time, the time zone, and the daylight saving time (also known as summer time) settings configured on the device.

To access the Time Zone Summary page, click System > Advanced Configuration > Time Zone > Summary in the navigation menu.

Figure 147: Time Zone Summary

Current Time	
Time	07:23:47
Zone	
Date	January 02, 1970
Time Source	No time source

Time Zone	
Zone	
Offset	UTC+0:00

Summer Time	
Summer Time	No Summer Time

Table 137: Time Zone Summary Fields

Field	Description
Current Time	<p>This section contains information about the system time and date on the device. If the current time has not been acquired by the SNTP client on the device or configured manually, this section shows the default time and date plus the amount of time since the system was reset.</p> <ul style="list-style-type: none"> <li>• Time — The current time on the system clock. This time is used to provide time stamps on log messages. Additionally, some CLI show commands include the time in the command output.</li> <li>• Zone — The acronym that represents the time zone.</li> <li>• Date — The current date on the system.</li> <li>• Time Source — The time source from which the time update is taken:                             <ul style="list-style-type: none"> <li>- SNTP — The time has been acquired from an SNTP server.</li> <li>- No Time Source — The time has either been manually configured or not configured at all.</li> </ul> </li> </ul>
Time Zone	<p>This section contains information about the time zone and offset.</p> <p>Zone — The acronym that represents the time zone.</p> <p>Offset — The number of hours offset from Coordinated Universal Time (UTC), which is also known as Greenwich Mean Time (GMT).</p>

Click Refresh to display the latest information from the router.

### 4.23.1 Time Zone Configuration

Use this page to manually configure the system clock settings. The SNTP client must be disabled to allow manual configuration of the system time and date.

To access the Time Zone Configuration page, click System > Advanced Configuration > Time Zone > Time Zone in the navigation menu.

Figure 148: Time Zone Configuration

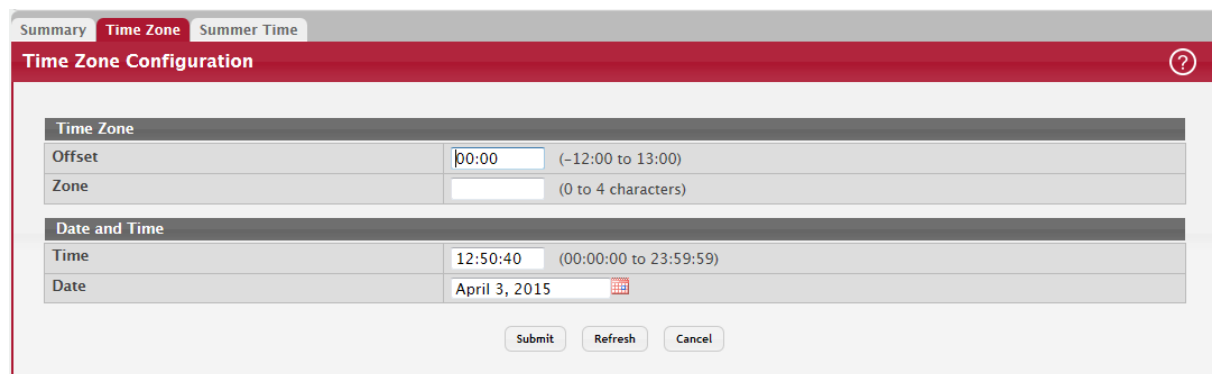


Table 138: Time Zone Configuration Fields

Field	Description
Time Zone	<p>The time zone settings include the amount of time the system clock is offset from Coordinated Universal Time (UTC) and the time zone acronym.</p> <ul style="list-style-type: none"> <li>Offset — The number of hours the system clock is offset from UTC, which is also known as Greenwich Mean Time (GMT).</li> <li>Zone — The acronym that represents the time zone. This field is not validated against an official list of time zone acronyms.</li> </ul>
Date and Time	<p>Use the fields in this section to manually configure the system time and date. If the SNTP client is enabled (Unicast mode or Broadcast mode), these fields cannot be configured.</p> <ul style="list-style-type: none"> <li>Time — The current time in hours, minutes, and seconds on the system clock.</li> <li>Date — The current date in month, day, and year on the system clock. To change the date, click the calendar icon to the right of the field, select the year from the menu, browse to the desired month, and click the date.</li> </ul>

Click Refresh to display the latest information from the router.

Click Submit to apply the settings to the running configuration and cause the change to take effect.

### 4.23.2 Summer Time Configuration

Use this page to configure settings for summer time, which is also known as daylight saving time. Used in some countries around the world, summer time is the practice of temporarily advancing clocks during the summer months. Typically clocks are adjusted forward one or more hours near the start of spring and are adjusted backward in autumn.

To access the Summer Time Configuration page, click System > Advanced Configuration > Time Zone > Summer Time in the navigation menu.

Figure 149: Summer Time Configuration

The screenshot shows the 'Summer Time Configuration' page. At the top, there are navigation tabs: 'Summary', 'Time Zone', and 'Summer Time'. The main title is 'Summer Time Configuration'. Below the title, there is a dropdown menu for 'Summer Time' currently set to 'Disable'. The configuration is divided into three main sections:

- Date Range:** Contains fields for 'Start Date', 'Starting Time of Day' (00:00 to 23:59), 'End Date', and 'Ending Time of Day' (00:00 to 23:59).
- Recurring Date:** Contains fields for 'Start Week' (First), 'Start Day' (Sunday), 'Start Month' (January), 'Starting Time of Day' (00:00 to 23:59), 'End Week' (First), 'End Day' (Sunday), 'End Month' (January), and 'Ending Time of Day' (00:00 to 23:59).
- Zone:** Contains fields for 'Offset' (1 to 1440) and 'Zone' (0 to 4 characters).

At the bottom of the form are three buttons: 'Submit', 'Refresh', and 'Cancel'.

Table 139: Summer Time Configuration Fields

Field	Description
Summer Time	<p>The summer time mode on the system:</p> <ul style="list-style-type: none"> <li>• Disable – Summer time is not active, and the time does not shift based on the time of year.</li> <li>• Recurring – Summer time occurs at the same time every year. The start and end times and dates for the time shift must be manually configured.</li> <li>• EU – The system clock uses the standard recurring summer time settings used in countries in the European Union. When this field is selected, the rest of the applicable fields on the page are automatically populated and cannot be edited.</li> <li>• USA – The system clock uses the standard recurring daylight saving time settings used in the United States. When this field is selected, the rest of the applicable fields on the page are automatically populated and cannot be edited.</li> <li>• Non-Recurring – Summer time settings are in effect only between the start date and end date of the specified year. When this mode is selected, the summer time settings do not repeat on an annual basis.</li> </ul>
Date Range	<p>The fields in this section are available only if the Non-Recurring mode is selected from the Summer Time menu.</p> <ul style="list-style-type: none"> <li>• Start Date – The day, month, and year that summer time begins. To change the date, click the calendar icon to the right of the field, select the year from the menu, browse to the desired month, and click the date.</li> <li>• Starting Time of Day – The time, in hours and minutes, to start summer time on the specified day.</li> <li>• End Date – The day, month, and year that summer time ends. To change the date, click the calendar icon to the right of the field, select the year from the menu, browse to the desired month, and click the date.</li> <li>• Ending Time of Day – The time, in hours and minutes to end summer time on the specified day.</li> </ul>
Recurring Date	<p>The fields in this section are available only if the Recurring mode is selected from the Summer Time menu.</p> <ul style="list-style-type: none"> <li>• Start Week – The week of the month within which summer time begins.</li> <li>• Start Day – The day of the week on which summer time begins.</li> <li>• Start Month – The month of the year within which summer time begins.</li> <li>• Starting Time of Day – The time, in hours and minutes, to start summer time.</li> <li>• End Week – The week of the month within which summer time ends.</li> <li>• End Day – The day of the week on which summer time ends.</li> <li>• End Month – The month of the year within which summer time ends.</li> <li>• Ending Time of Day – The time, in hours and minutes, to end summer time.</li> </ul>
Zone	<p>The fields in this section are available only if the Recurring or Non-Recurring modes are selected from the Summer Time menu.</p> <ul style="list-style-type: none"> <li>• Offset – The number of minutes to shift the summer time from the standard time.</li> <li>• Zone – The acronym associated with the time zone when summer time is in effect.</li> </ul>

Click Refresh to display the latest information from the router.

Click Submit to apply the settings to the running configuration and cause the change to take effect.

## 4.24 Configuring and Viewing ISDP Information

The Industry Standard Discovery Protocol (ISDP) is a proprietary Layer 2 network protocol which inter-operates with Cisco devices running the Cisco Discovery Protocol (CDP). ISDP is used to share information between neighboring devices. FASTPATH software participates in the CDP protocol and is able to both discover and be discovered by other CDP supporting devices.

## 4.24.1 ISDP Global Configuration

To access the ISDP Global Configuration page, click System > Advanced Configuration > ISDP > Global in the navigation menu.

Figure 150: ISDP Global Configuration

ISDP Global Configuration	
ISDP Mode	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
ISDP V2 Mode	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Message Interval (Seconds)	30 (5 to 254)
Hold Time Interval (Seconds)	180 (10 to 255)
Device ID	none
Device ID Format Capability	Serial Number, Host Name
Device ID Format	Serial Number

The following table describes the fields available on the ISDP Global Configuration page.

Table 140: ISDP Global Configuration

Field	Description
ISDP Mode	Use this field to enable or disable the Industry Standard Discovery Protocol on the switch.
ISDP V2 Mode	Use this field to enable or disable the Industry Standard Discovery Protocol v2 on the switch.
Message Interval	Specifies the ISDP transmit interval. The range is (5–254). Default value is 30 seconds.
Hold Time Interval	The receiving device holds ISDP message during this time period. The range is (10–255). Default value is 180 seconds.
Device ID	The Device ID advertised by this device. The format of this Device ID is characterized by the value of Device ID Format object.
Device ID Format Capability	Indicates the Device ID format capability of the device. <ul style="list-style-type: none"> <li>serialNumber—Indicates that the device uses serial number as the format for its Device ID.</li> <li>macAddress—Indicates that the device uses Layer 2 MAC address as the format for its Device ID.</li> <li>other—Indicates that the device uses its platform specific format as the format for its Device ID.</li> </ul>
Device ID Format	Indicates the Device ID format of the device. <ul style="list-style-type: none"> <li>serialNumber—Indicates that the value is in the form of an ASCII string containing the device serial number.</li> <li>macAddress—Indicates that the value is in the form of Layer 2 MAC address.</li> <li>other—Indicates that the value is in the form of a platform specific ASCII string containing info that identifies the device. For example: ASCII string contains serialNumber appended/prepended with system name.</li> </ul>

## 4.24.2 ISDP Cache Table

From the ISDP Cache Table page, you can view information about other devices the switch has discovered through the ISDP.

To access the ISDP Cache Table page, click System > Advanced Configuration > ISDP > Cache Table in the navigation menu.

Figure 151: ISDP Cache Table

Device ID	Interface	IP Address	Version	Hold Time (Seconds)	Capability	Platform	Port ID	Protocol Version	Last Time Changed
	2/0/8	10.27.15.8	3.2.0.7	176	R	PCT6248	2/0/24	2	0d:06:20:12
1225098	2/0/8	10.27.226.154	F.4.30.2	165	S	BCM-56224	1/0/2	2	0d:06:20:01
none	2/0/8	10.27.21.27	5.10.8.20	162	R	BCM956143R	0/1	2	0d:06:19:58

The following table describes the fields available on the ISDP Cache Table page.

Table 141: ISDP Cache Table

Field	Description
Device ID	Displays the string with Device ID which is reported in the most recent ISDP message.
Interface	Displays the interface that this neighbor is attached to.
IP Address	The (first) network-layer address that is reported in the Address TLV of the most recently received ISDP message.
Version	Displays the Version string for the neighbor.
Hold Time	Displays the ISDP hold time for the neighbor.
Capability	Displays the ISDP Functional Capabilities for the neighbor.
Platform	Displays the ISDP Hardware Platform for the neighbor.
Port ID	Displays the ISDP port ID string for the neighbor.
Protocol Version	Displays the ISDP Protocol Version for the neighbor.
Last Time Changed	Displays when entry was last modified.
Clear (Button)	Clears all entries from the table. The table is repopulated as ISDP messages are received from neighbors.

## 4.24.3 ISDP Interface Configuration

From the ISDP Interface Configuration page, you can configure the ISDP settings for each interface.

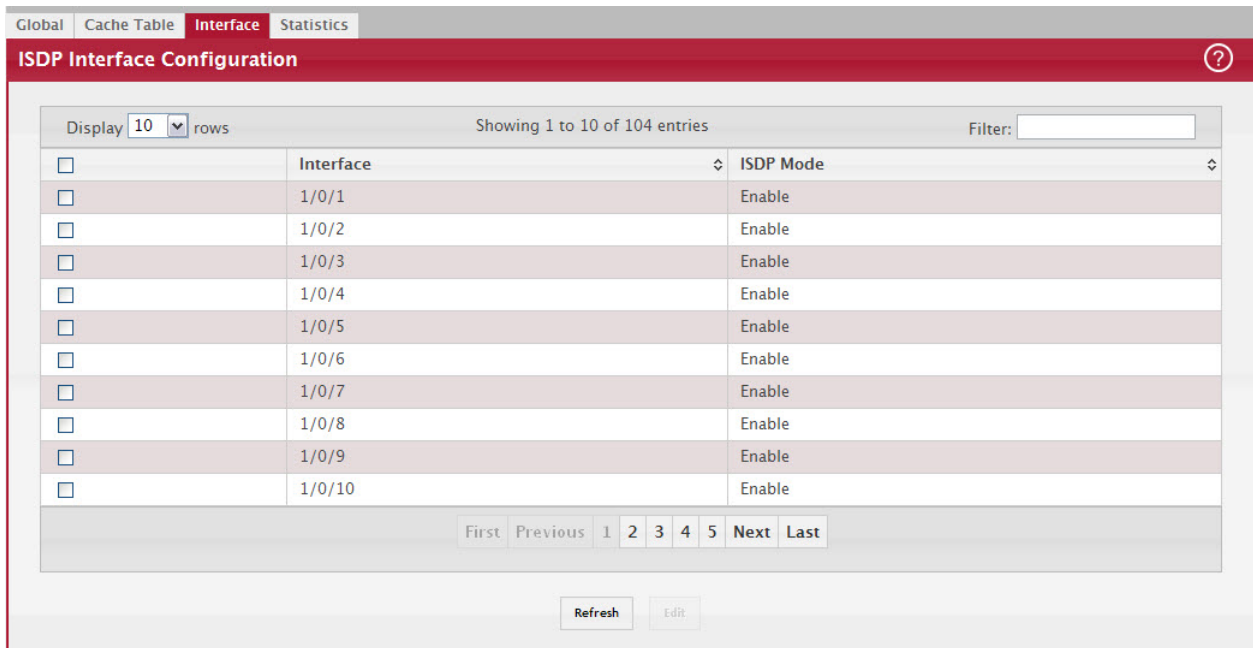
### NOTICE

If ISDP is enabled on an interface, it must also be enabled globally for the interface to transmit ISDP packets. If the ISDP mode on the ISDP Global Configuration page is disabled, the interface will not transmit ISDP packets, regardless of the mode configured on the interface.

To access the ISDP Interface Configuration page, click System > Advanced Configuration > ISDP > Interface in the navigation menu.



Figure 152: ISDP Interface Configuration



The following table describes the fields available on the ISDP Interface Configuration page.

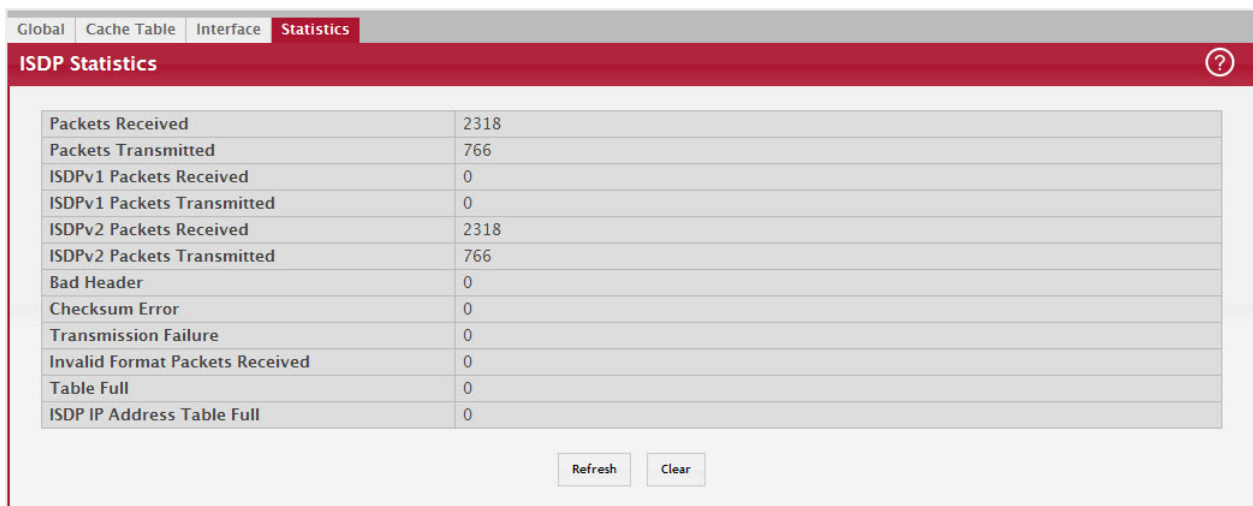
Table 142: ISDP Interface Configuration Fields

Field	Description
Interface	Select the interface with the ISDP mode status to configure or view.
ISDP Mode	Use this field to enable or disable the Industry Standard Discovery Protocol on the selected interface.

#### 4.24.4 Statistics

From the ISDP Statistics page, you can view information about the ISDP packets sent and received by the switch. To access the ISDP Statistics page, click System > Advanced Configuration > ISDP > Statistics in the navigation menu.

Figure 153: ISDP Statistics



The following table describes the fields available on the ISDP Statistics page.

**Table 143: ISDP Statistics Fields**

Field	Description
ISDP Packets Received	Displays the number of all ISDP protocol data units (PDUs) received.
ISDP Packets Transmitted	Displays the number of all ISDP PDUs transmitted.
ISDPv1 Packets Received	Displays the number of v1 ISDP PDUs received.
ISDPv1 Packets Transmitted	Displays the number of v1 ISDP PDUs transmitted.
ISDPv2 Packets Received	Displays the number of v2 ISDP PDUs received.
ISDPv2 Packets Transmitted	Displays the number of v2 ISDP PDUs transmitted.
ISDP Bad Header	Displays the number of ISDP PDUs that were received with bad headers.
ISDP Checksum Error	Displays the number of ISDP PDUs that were received with checksum errors.
ISDP Transmission Failure	Displays the number of ISDP PDUs transmission failures.
Invalid Format ISDP Packets Received	Displays the number of ISDP PDUs that were received with an invalid format.
Table Full	Displays the number of times the system tried to add an entry to the ISDP table but was unsuccessful because the table was full.
ISDP IP Address Table Full	Displays the number of times the system tried to add an entry to the ISDP IP Address table but was unsuccessful because the table was full.
Clear (Button)	Resets all statistics to zero.

## 4.25 Link Dependency

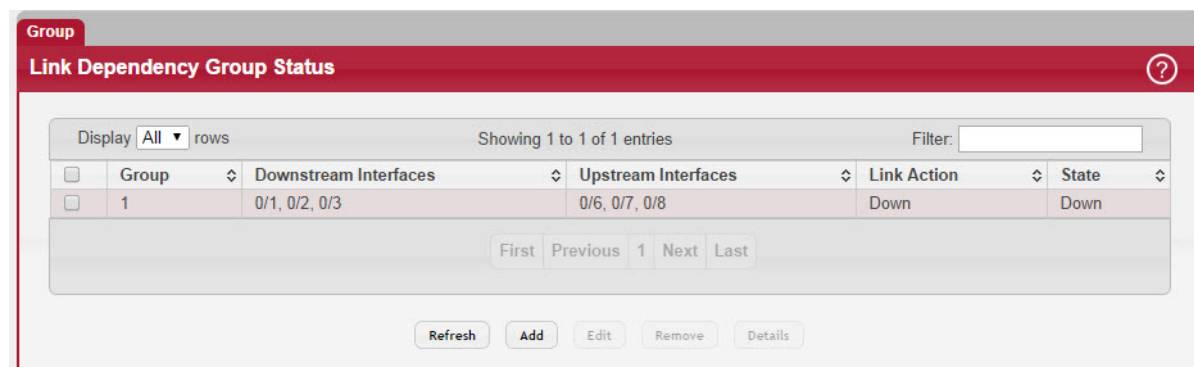
The link dependency feature provides the ability to enable or disable one or more ports based on the link state of one or more different ports. With link dependency enabled on a port, the link state of that port is dependent on the link state of another port. For example, if port A is dependent on port B and the switch detects a link loss on port B, the switch automatically brings down the link on port A. When the link is restored to port B, the switch automatically restores the link to port A.

### 4.25.1 Link Dependency Group Status

Use this page to configure link dependency groups. Link dependency allows the link status of one interface to be dependent on the link status of another interface. Link state groups define the interface link dependency.

To access the Link Dependency Group Status page, click System > Advanced Configuration > Link Dependency > Group in the navigation menu.

**Figure 154: Link Dependency Group Status**



Use the buttons to perform the following tasks:

- To add a link dependency group, click Add. Then, specify a group number, link action, and the interfaces that share a dependency.
- To change the settings for a group, select the check box associated with the group and click Edit.
- To delete a link dependency group, select the check box associated with each entry to delete and click Remove.
- To view additional information about a group, select the check box associated with the group and click Details.

**Table 144: Link Dependency Group Status**

Field	Description
Group	The unique link dependency group identifier.
Downstream Interfaces	The set of interfaces that depend on other interfaces. In other words, the link state of the downstream interfaces depends on the link state of the upstream interfaces.
Upstream Interfaces	The set of interfaces that determine the link state of the downstream interfaces.
Link Action	The action performed on downstream interfaces when the upstream interfaces are down, which can be one of the following: <ul style="list-style-type: none"> <li>• Up: Downstream interfaces are up when upstream interfaces are down.</li> <li>• Down: Downstream interfaces go down when upstream interfaces are down.</li> </ul> Creating a link dependency group with the up link action essentially creates a backup link for the dependent link and alleviates the need to implement STP to handle the fail-over.
State	The group state, which can be one of the following: <ul style="list-style-type: none"> <li>• Up: Link action is up, and no upstream interfaces have their link up, or link action is down and there are upstream interfaces that have their link up.</li> <li>• Down: Link is down when the above conditions are not true.</li> </ul>
Available Interfaces	Available in the Add Group dialog, this field lists the interfaces that can be added to the group. An interface defined as an upstream interface cannot be defined as a downstream interface in the same link state group or in a different group. Similarly, an interface defined as a downstream interface cannot be defined as an upstream interface.  To move an interface between the Available Interfaces and Downstream Interfaces or Upstream Interfaces fields, click the interface (or CTRL + click to select multiple interfaces), and then click the appropriate arrow to move the selected interfaces to the desired field.
Link Up	Available in the Group Entry Details dialog, this field lists the upstream and downstream interfaces that currently have their link up.
Link Down	Available in the Group Entry Details dialog, this field lists the upstream and downstream interfaces that currently have their link down.

## 4.26 Link Local Protocol Filtering

Link Local Protocol Filtering (LLPF) allows the switch to filter out multiple proprietary protocol PDUs, such as Port Aggregation Protocol (PAgP), if the problems occur with proprietary protocols running on standards-based switches. If certain protocol PDUs cause unexpected results, LLPF can be enabled to prevent those protocol PDUs from being processed by the switch.

---

The LLPF feature is not supported on all platforms.

### **NOTICE**

---

## 4.26.1 LLPF Interface Configuration

Use the Link Local Protocol Filtering Configuration page to enable or disable the filtering of various proprietary protocols. To access the Link Local Protocol Filtering Configuration page, click System > Advanced Configuration > LLPF > Configuration in the navigation menu.

Figure 155: Link Local Protocol Filtering Configuration

Interface	Blocked Protocols
<input type="checkbox"/> 0/1	UDLD
<input type="checkbox"/> 0/2	UDLD
<input type="checkbox"/> 0/3	UDLD
<input type="checkbox"/> 0/4	UDLD
<input type="checkbox"/> 0/5	UDLD
<input type="checkbox"/> 0/6	UDLD
<input type="checkbox"/> 0/7	UDLD
<input type="checkbox"/> 0/8	UDLD
<input type="checkbox"/> 0/9	UDLD
<input type="checkbox"/> 0/10	UDLD

Use the buttons to perform the following tasks:

- To select the protocols for LLPF to block on an interface, click Add. Then, select an interface to configure and select each protocol to block on that interface.
- To change which protocols are blocked on an interface, select the check box associated with the interface and click Edit.
- To delete an entry from the list, select the check box associated with each entry to delete and click Remove.

The following table describes the fields available on the Link Local Protocol Filtering Configuration page.

Table 145: Link Local Protocol Filtering Configuration

Field	Description
Interface	Identifies the physical or LAG interface.
ISDP	When enabled, the select port blocks ISDP PDUs.
VTP	When enabled, the select port blocks VTP PDUs.
DTP	When enabled, the select port blocks DTP PDUs.
UDLD	When enabled, the select port blocks UDLD PDUs.
PAGP	When enabled, the select port blocks PAgP PDUs.
SSTP	When enabled, the select port blocks SSTP PDUs.
All Protocols	All the above mentioned protocols will be dropped in addition to protocols with a Destination MAC of 01:00:0C:CC:CC:CX.

When you configure the blocked protocols on the Add LLPF Interface or Edit LLPF Interface page, select the check box for each protocol to block, or clear the box to allow the protocol on the selected interface. If you select the All Protocols option, all protocols are blocked whether their associated box is checked or unchecked.

Click Submit to apply the settings to the running configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.

## 5/ Configuring Switching Information

### 5.1 Managing VLANs

Adding Virtual LAN (VLAN) support to a Layer 2 switch offers some of the benefits of both bridging and routing. Like a bridge, a VLAN switch forwards traffic based on the Layer 2 header, which is fast, and like a router, it partitions the network into logical segments, which provides better administration, security and management of multicast traffic.

A VLAN is a set of end stations and the switch ports that connect them. You may have many reasons for the logical division, such as department or project membership. The only physical requirement is that the end station and the port to which it is connected both belong to the same VLAN.

Each VLAN in a network has an associated VLAN ID, which appears in the IEEE 802.1Q tag in the Layer 2 header of packets transmitted on a VLAN. An end station may omit the tag, or the VLAN portion of the tag, in which case the first switch port to receive the packet may either reject it or insert a tag using its default VLAN ID. A given port may handle traffic for more than one VLAN, but it can only support one default VLAN ID.

#### 5.1.1 VLAN Status

Use the VLAN Status page to view information about the VLANs configured on your system, and to configure the statistics collection mode on VLANs.

To access the VLAN Status page, click Switching > VLAN > Status in the navigation menu.

Figure 156: VLAN Status

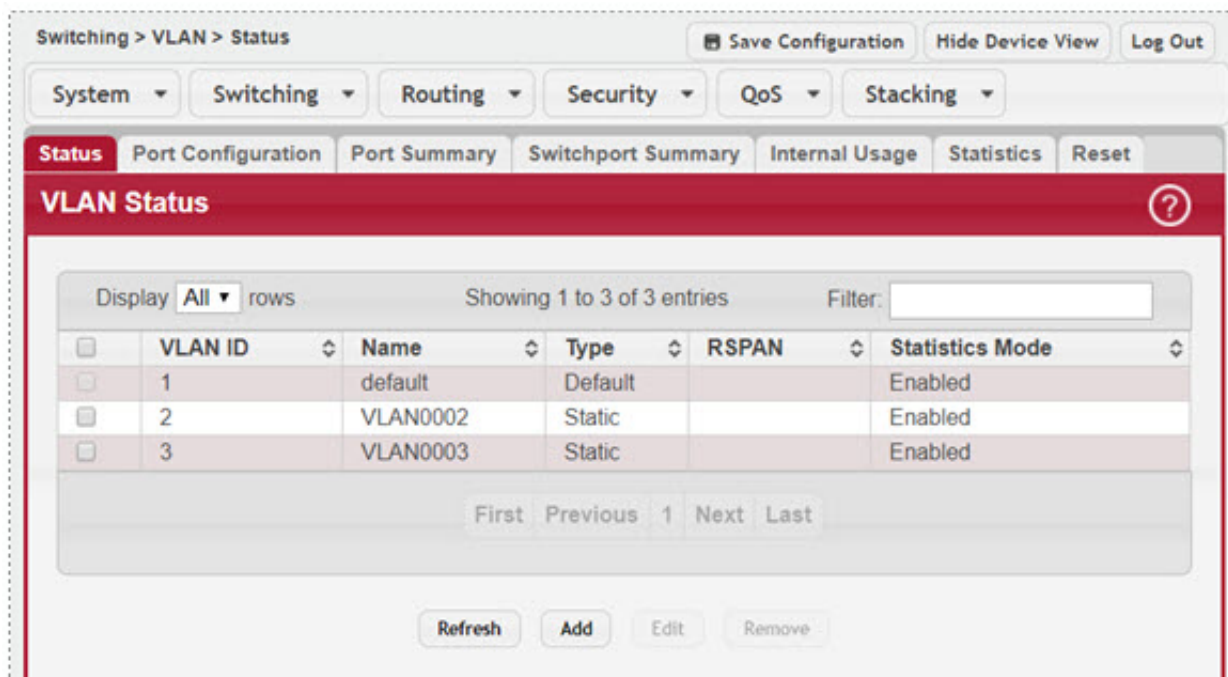


Table 146: VLAN Status Fields

Field	Description
VLAN ID	The VLAN Identifier (VID) of the VLAN. The range of the VLAN ID is 1 to 3965.
VLAN Name	The name of the VLAN. VLAN ID 1 is always named Default.
VLAN Type	The VLAN type, which can be one of the following: <ul style="list-style-type: none"> <li>• Default: (VLAN ID = 1) -- always present</li> <li>• Static: A VLAN you have configured</li> <li>• Dynamic: A VLAN created by GVRP registration that you have not converted to static, and that GVRP may therefore remove</li> </ul>

Table 146: VLAN Status Fields (Continued)

Field	Description
RSPAN	Lists the status of RSPAN, enabled or disabled.
Statistics Mode	Lists the status of Statistics Mode, enabled or disabled.

Click Refresh to display the latest information from the router.

### 5.1.1.1 Add a VLAN

To add a VLAN, click the Add button and specify a VLAN ID in the available field. For static VLANs, specify a name for the VLAN. This field is optional and is used to help identify the VLAN. This field is not available for other VLAN types.

Figure 157: Add VLAN

Click Submit to add the VLAN to the system.

### 5.1.1.2 Edit VLAN Configuration

To edit the VLAN Configuration, select the entry to modify and click the Edit button.

Figure 158: Edit VLAN Configuration

Edit the configured VLAN settings, as follows.

Table 147: Edit VLAN Configuration

Field	Description
Name	For static VLANs, specify a name for the VLAN. The name can be 1 to 32 alpha-numeric characters. This field is optional and is used to help identify the VLAN. This field is not available for other VLAN types.
Convert VLAN Type to Static	For dynamic VLANs, select this option to convert the dynamic VLAN to a static VLAN. This option is not available for other VLAN types. A dynamic VLAN is learned by using GVRP, which is an industry-standard protocol that propagates VLAN information from one network device to another. GVRP can also remove dynamic VLANs. If you convert a dynamic VLAN to a static VLAN, it cannot be removed by GVRP.
Statistics Mode	Select the option to enable or disable Statistics Mode. Use the VLAN Statistics page to view and clear statistical information on which the statistics mode is enabled. See <a href="#">Section 5.1.6: "Configure VLAN Statistics"</a> .

Click Submit to submit the VLAN configuration changes. Click Cancel to cancel the changes.

### 5.1.1.3 Remove VLAN Configuration

To remove one or more configured VLANs, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

## 5.1.2 VLAN Port Configuration

Use the VLAN Port Configuration page to configure a virtual LAN on a port.

To access the VLAN Port Configuration page, click Switching > VLAN > Port Configuration in the navigation menu.

Figure 159: VLAN Port Configuration

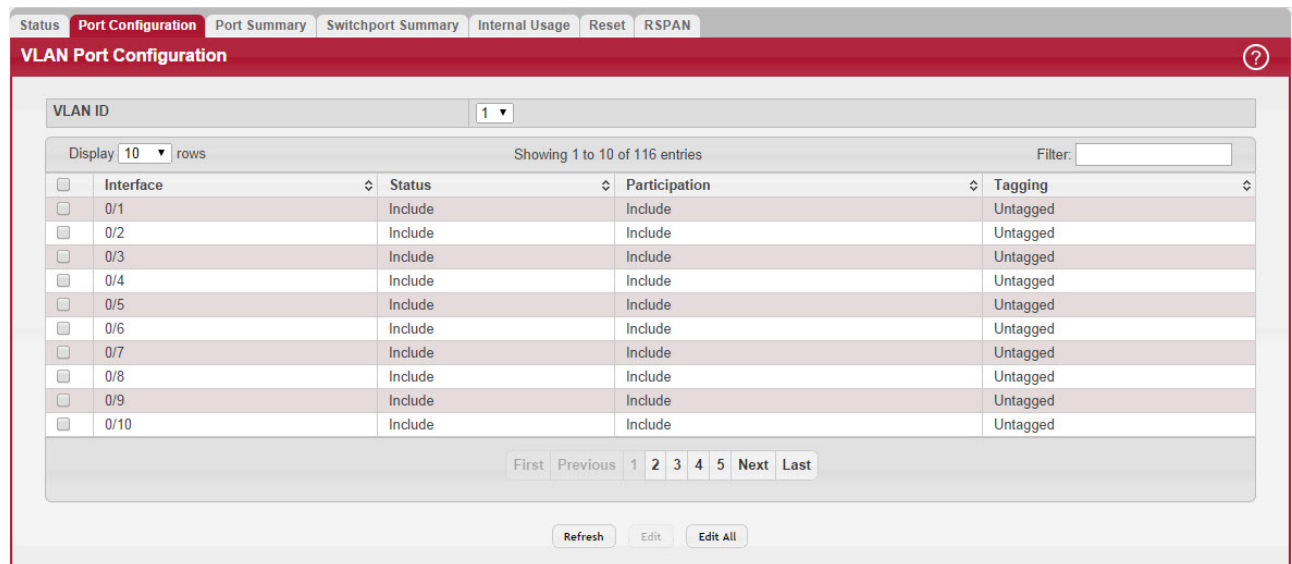


Table 148: VLAN Port Configuration Fields

Field	Description
VLAN ID	The menu includes the VLAN ID for all VLANs configured on the device. To view or configure settings for a VLAN, be sure to select the correct VLAN from the menu.
Interface	Select the interface for which you want to display or configure data. Select All to set the parameters for all ports to same values.
Status	The current participation mode of the interface in the selected VLAN. The value of the Status field differs from the value of the Participation field only when the Participation mode is set to Auto Detect. The Status is one of the following: <ul style="list-style-type: none"> <li>• Include – The port is a member of the selected VLAN.</li> <li>• Exclude – The port is not a member of the selected VLAN.</li> </ul>
Participation	The participation mode of the interface in the selected VLAN, which is one of the following: <ul style="list-style-type: none"> <li>• Include – The port is always a member of the selected VLAN. This mode is equivalent to registration fixed in the IEEE 802.1Q standard.</li> <li>• Exclude – The port is never a member of the selected VLAN. This mode is equivalent to registration forbidden in the IEEE 802.1Q standard.</li> <li>• Auto Detect – The port can be dynamically registered in the selected VLAN through GVRP. The port will not participate in this VLAN unless it receives a GVRP request. This mode is equivalent to registration normal in the IEEE 802.1Q standard.</li> </ul>
Tagging	The tagging behavior for all the ports in this VLAN, which is one of the following: <ul style="list-style-type: none"> <li>• Tagged – The frames transmitted in this VLAN will include a VLAN ID tag in the Ethernet header.</li> <li>• Untagged – The frames transmitted in this VLAN will be untagged.</li> </ul>

Use the buttons to perform the following tasks:

- To configure settings for one or more interfaces, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.
- To apply the same settings to all interfaces, click Edit All and configure the desired settings.
- To reload the page and view the most current information, click Refresh.

### 5.1.3 VLAN Port Summary

Use the VLAN Port Summary page to view VLAN configuration information for all the ports on the system. To access the VLAN Port Summary page, click Switching > VLAN > Port Summary in the navigation menu.

Figure 160: VLAN Port Summary

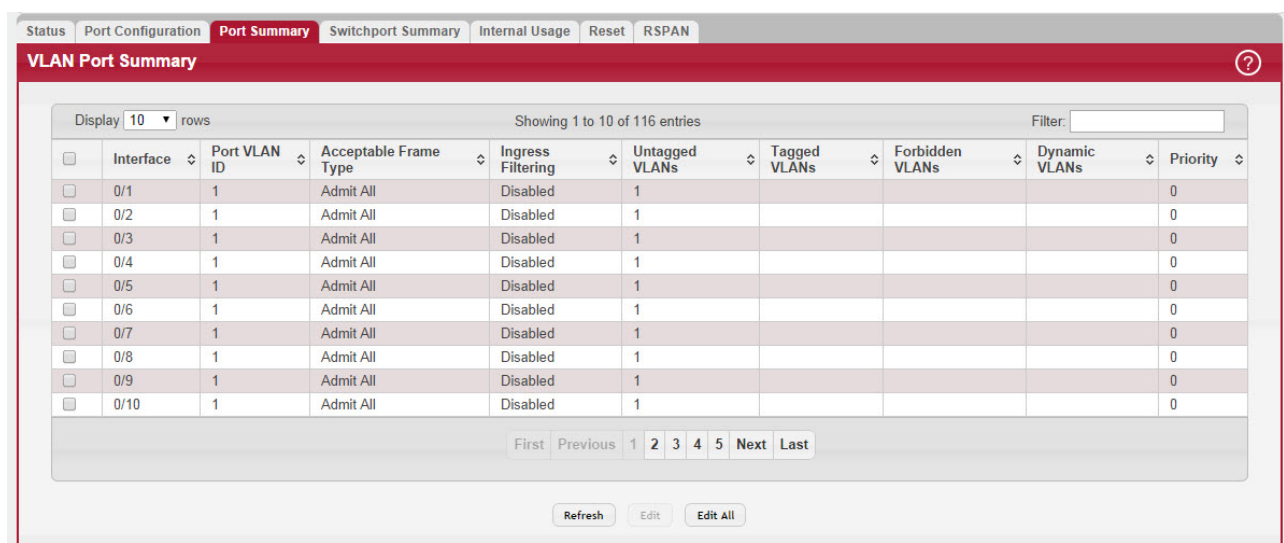




Table 149: VLAN Port Summary Fields

Field	Description
Interface	Identifies the physical interface associated with the rest of the data in the row.
Port VLAN ID	The VLAN ID assigned to untagged or priority tagged frames received on this port. This value is also known as the Port VLAN ID (PVID). In a tagged frame, the VLAN is identified by the VLAN ID in the tag.
Acceptable Frame Types	Indicates how the interface handles untagged and priority tagged frames. The options include the following: <ul style="list-style-type: none"> <li>• Admit All – Untagged and priority tagged frames received on the interface are accepted and assigned the value of the Port VLAN ID for this interface.</li> <li>• Only Tagged – The interface discards any untagged or priority tagged frames it receives.</li> <li>• Only Untagged – The interface discards any tagged frames it receives.</li> </ul> For all options, VLAN tagged frames are forwarded in accordance with the IEEE 802.1Q VLAN standard.
Ingress Filtering	Shows how the port handles tagged frames. <ul style="list-style-type: none"> <li>• Enable: A tagged frame is discarded if this port is not a member of the VLAN identified by the VLAN ID in the tag.</li> <li>• Disable: All tagged frames are accepted, which is the factory default.</li> </ul>
Untagged VLANs	VLANs that are configured on the port to transmit egress packets as untagged.
Tagged VLANs	VLANs that are configured on the port to transmit egress packets as tagged.
Forbidden VLANs	When configuring port memberships in VLANs, you can specify one or more VLANs to be excluded from the available VLANs for the port. The forbidden VLANs list shows the VLANs to which the port cannot be assigned membership.
Dynamic VLANs	The list of VLANs of which the port became a member as result of the operations of dynamic VLAN protocols. When a VLAN is created as a dynamic VLAN, any port that is configured as switchport type Trunk or General automatically becomes a member of the VLAN, unless the VLAN port is excluded from the VLAN.
Priority	Identifies the default 802.1p priority assigned to untagged packets arriving at the port.

Use the buttons to perform the following tasks:

- To configure settings for one or more interfaces, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.
- To apply the same settings to all interfaces, click Edit All and configure the desired settings.
- To reload the page and view the most current information, click Refresh.

### 5.1.4 Switchport Summary

Use the Switchport Summary page to configure switchport mode settings on interfaces. The switchport mode defines the purpose of the port based on the type of device it connects to and constraints the VLAN configuration of the port accordingly. Assigning the appropriate switchport mode helps simplify VLAN configuration and minimize errors.

To access the Switchport Summary page, click Switching > VLAN > Switchport Summary in the navigation menu.

Figure 161: VLAN Switchport Summary

Interface	Switchport Mode	Operational Switchport Mode	Access VLAN ID	Native VLAN ID	Native VLAN Tagging	Trunk Allowed VLANs
0/1	General	General	1	1	Disabled	1-4093
0/2	General	General	1	1	Disabled	1-4093
0/3	General	General	1	1	Disabled	1-4093
0/4	General	General	1	1	Disabled	1-4093
0/5	General	General	1	1	Disabled	1-4093
0/6	General	General	1	1	Disabled	1-4093
0/7	General	General	1	1	Disabled	1-4093
0/8	General	General	1	1	Disabled	1-4093
0/9	General	General	1	1	Disabled	1-4093
0/10	General	General	1	1	Disabled	1-4093

Table 150: VLAN Switchport Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When editing information for one or more interfaces, this field identifies the interfaces that are being configured.
Switchport Mode	<p>The switchport mode of the interface, which is one of the following:</p> <ul style="list-style-type: none"> <li>• Access–Access mode is suitable for ports connected to end stations or end users. Access ports participate only in one VLAN. They accept both tagged and untagged packets, but always transmit untagged packets.</li> <li>• Trunk–Trunk mode is intended for ports that are connected to other switches. Trunk ports can participate in multiple VLANs and accept both tagged and untagged packets.</li> <li>• General–General mode enables a custom configuration of a port. The user configures the General port VLAN attributes such as membership, PVID, tagging, ingress filter, etc., using the settings on the Port Configuration page. By default, all ports are initially configured in General mode.</li> <li>• Private VLAN Host–The interface belongs to a secondary VLAN and, depending upon the type of secondary VLAN, can either communicate with other ports in the same community (if the secondary VLAN is a community VLAN) and with the promiscuous ports, or is able to communicate only with the promiscuous ports (if the secondary VLAN is an isolation VLAN).</li> <li>• Private VLAN Promiscuous–The interface belongs to a primary VLAN and can communicate with all interfaces in the private VLAN, including other promiscuous ports, community ports, and isolated ports.</li> <li>• Private VLAN Promiscuous Trunk–The interface belongs to a primary VLAN and can communicate with all interfaces in the private VLAN, including other promiscuous trunk ports, community ports, and isolated ports.</li> </ul>

Table 150: VLAN Switchport Summary Fields (Continued)

Field	Description
Operational Switchport Mode	The operational switchport mode of the interface, which is one of the following: <ul style="list-style-type: none"> <li>• Access</li> <li>• Trunk</li> <li>• General</li> <li>• Private VLAN Host</li> <li>• Private VLAN Promiscuous</li> <li>• Private VLAN Promiscuous Trunk</li> </ul>
Access VLAN ID	The access VLAN for the port, which is valid only when the port switchport mode is Access.
Native VLAN ID	The native VLAN for the port, which is valid only when the port switchport mode is Trunk.
Native VLAN Tagging	When enabled, if the trunk port receives untagged frames, it forwards them on the native VLAN with no VLAN tag. When disabled, if the port receives untagged frames, it includes the native VLAN ID in the VLAN tag when forwarding.
Trunk Allowed VLANs	The set of VLANs of which the port can be a member, when configured in Trunk mode. By default, this list contains all possible VLANs even if they have not yet been created.

Use the buttons to perform the following tasks:

- To configure settings for one or more interfaces, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.
- To apply the same settings to all interfaces, click Edit All and configure the desired settings.
- To reload the page and view the most current information, click Refresh.

### 5.1.5 VLAN Internal Usage

Use the VLAN Internal Usage Configuration page to assign a Base VLAN ID for internal allocation of VLANs to the routing interface.

To access the VLAN Internal Usage page, click Switching > VLAN > Internal Usage in the navigation menu.

Figure 162: VLAN Internal Usage Configuration

Table 151: VLAN Internal Usage Configuration Fields

Field	Description
Base VLAN ID	The first VLAN ID to be assigned to a port-based routing interface.
Allocation Policy	Determines whether VLAN IDs assigned to port-based routing interfaces start at the base and decrease in value (Descending) or start at the base and increase in value (Ascending).
VLAN ID	The VLAN ID assigned to a port-based routing interface. The device automatically assigns an unused VLAN ID when the routing interface is created.
Routing Interface	The port-based routing interface associated with the VLAN.

If you change any information on the page, click Submit to apply the changes to the system.

## 5.1.6 Configure VLAN Statistics

Use the [Section 5.1.1.2: "Edit VLAN Configuration"](#) page to Enable or Disable the statistics collection mode on VLANs.VLAN Edit Configuration.

Use the VLAN Statistics page to view and clear the statistical information for VLANs on which the statistics mode is enabled.

### NOTICE

The VLAN Statistics page is only available if the VLAN Statistics feature is supported on the platform.

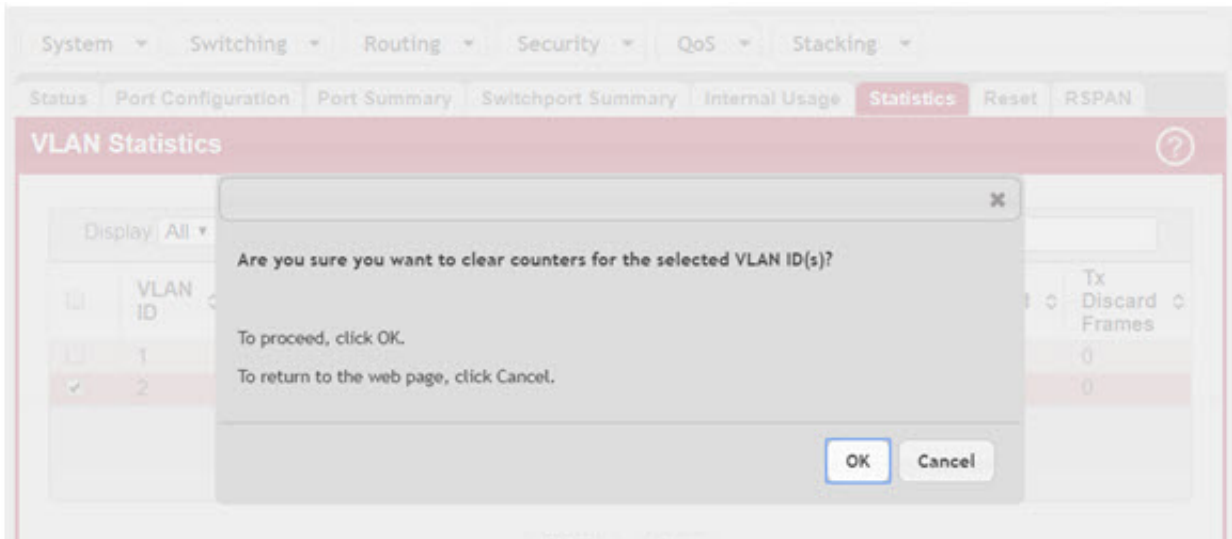
To access the VLAN Statistics page, click Switching > VLAN > Statistics in the navigation menu.

Figure 163: VLAN Statistics

VLAN ID	Rx Bytes	Rx Frames	Rx Discard Bytes	Rx Discard Frames	Tx Bytes	Tx Frames	Tx Discard Bytes	Tx Discard Frames
1	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0

To reset the counter values on one or more VLANs to the default values, select the VLAN(s) from the list and click the Clear button. This opens a modal page as shown in [Figure 164: "Clear VLAN Statistics," on page 199](#). Confirm the action to reset the counter values for the VLAN(s).

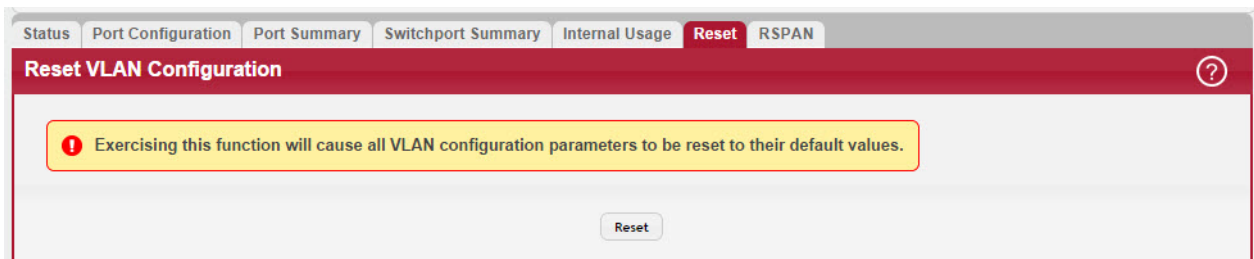
Figure 164: Clear VLAN Statistics



### 5.1.7 Reset VLAN Configuration

Use the Reset Configuration page to return all VLAN parameters for all interfaces to the factory default values. To access the Reset Configuration page, click Switching > VLAN > Reset in the navigation menu.

Figure 165: Reset VLAN Configuration



When you click Reset, the screen refreshes, and you are asked to confirm the reset. Click Reset again to restore all default VLAN settings for the ports on the system.

### 5.1.8 RSPAN Configuration

Use this page to configure the VLAN to use as the Remote Switched Port Analyzer (RSPAN) VLAN. RSPAN allows you to mirror traffic from multiple source ports (or from all ports that are members of a VLAN) from different network devices and send the mirrored traffic to a destination port (a probe port connected to a network analyzer) on a remote device. The mirrored traffic is tagged with the RSPAN VLAN ID and transmitted over trunk ports in the RSPAN VLAN.

To access the RSPAN page, click Switching > VLAN > RSPAN in the navigation menu.

Figure 166: RSPAN VLAN Configuration

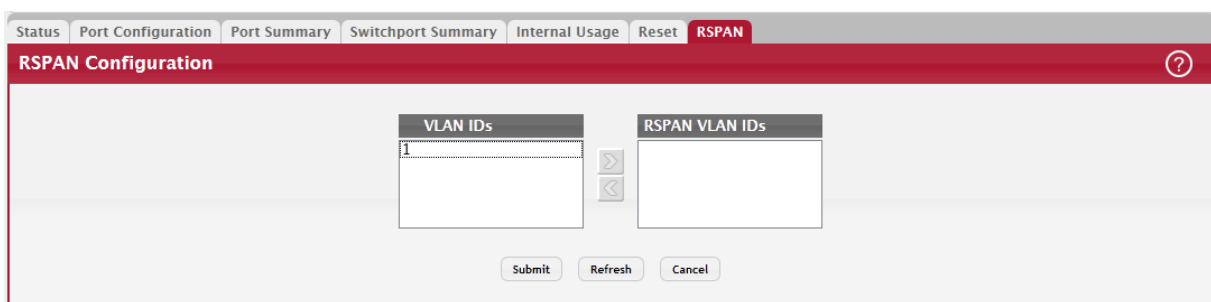


Table 152: RSPAN VLAN Configuration Fields

Field	Description
VLAN IDs	The VLANs configured on the system that are not currently enabled as Private VLANs. To enable a VLAN as a RSPAN VLAN, click the VLAN ID to select it (or CTRL + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the RSPAN VLAN IDs window.
RSPAN VLAN IDs	The VLANs that are enabled as RSPAN VLAN. To disable a VLAN as a RSPAN VLAN, click the VLAN ID to select it (or CTRL + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the VLAN IDs window.

Click Refresh to display the latest information from the router.

If you change any information on the page, click Submit to apply the changes to the system.

## 5.2 Configuring UDLD

The UDLD feature detects unidirectional links on physical ports by exchanging packets containing information about neighboring devices. The purpose of the UDLD feature is to detect and avoid unidirectional links. A unidirectional link is a forwarding anomaly in a Layer 2 communication channel in which a bidirectional link stops passing traffic in one direction.

To access the UDLD Configuration page, click Switching > UDLD > Configuration in the navigation menu.

Figure 167: UDLD Configuration

Table 153: UDLD Configuration Fields

Field	Description
Admin Mode	The administrative mode of UDLD on the device. UDLD must be administratively enabled on the device and on an interface for that interface to send UDLD messages. Additionally, UDLD must be enabled on the both sides of the link for the device to detect a unidirectional link.
Message Interval (Seconds)	The amount of time to wait between sending UDLD probe messages on ports that are in the advertisement phase.
Timeout Interval (Seconds)	The amount of time to wait to receive a UDLD message before considering the UDLD link to be unidirectional.

Click Refresh to display the latest information from the router.

If you change any information on the page, click Submit to apply the changes to the system.

## 5.2.1 UDLD Interface Configuration

Use this page to configure the per-port UDLD settings.

To access the UDLD Interface Configuration page, click Switching > UDLD > Interface Configuration in the navigation menu.

Figure 168: UDLD Interface Configuration

Interface	Admin Mode	UDLD Mode	UDLD Status
0/1	Disabled	Normal	Not Applicable
0/2	Disabled	Normal	Not Applicable
0/3	Disabled	Normal	Not Applicable
0/4	Disabled	Normal	Not Applicable
0/5	Disabled	Normal	Not Applicable
0/6	Disabled	Normal	Not Applicable
0/7	Disabled	Normal	Not Applicable
0/8	Disabled	Normal	Not Applicable

Use the buttons to perform the following tasks:

- To configure UDLD settings for one or more interfaces, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.
- To reset all UDLD ports that have a UDLD Status of Shutdown, click UDLD Port Reset. If the global and interface UDLD administrative mode is enabled and the port link is up, the port restarts the exchange of UDLD messages with its link partner. The UDLD port status is Shutdown if UDLD has detected an unidirectional link and has put the port in a disabled state.

Table 154: UDLD Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. In the Edit UDLD Interface Configuration window, this field identifies each interface that is being configured.
Admin Mode	The administrative mode of UDLD on the port.

Table 154: UDLD Interface Configuration Fields (Continued)

Field	Description
UDLD Mode	<p>The UDLD mode for the port, which is one of the following:</p> <ul style="list-style-type: none"> <li>• Normal – The state of the port is classified as Undetermined if an anomaly exists. An anomaly might be the absence of its own information in received UDLD messages or the failure to receive UDLD messages. An Undetermined state has no effect on the operation of the port. The port is not disabled and continues operating. When operating in UDLD normal mode, a port will be put into a disabled (Shutdown) state only in the following situations: <ul style="list-style-type: none"> <li>- The UDLD PDU received from a partner does not have its own details (echo).</li> <li>- When there is a loopback, and information sent out on a port is received back exactly as it was sent.</li> </ul> </li> <li>• Aggressive – The port is put into a disabled state for the same reasons that it occurs in normal mode. Additionally, a port in UDLD aggressive mode can be disabled if the port does not receive any UDLD echo packets even after bidirectional connection was established. If a bidirectional link is established, and packets suddenly stop coming from partner device, the UDLD aggressive-mode port assumes that link has become unidirectional.</li> </ul>
UDLD Status	<p>The UDLD status on the port, which is one of the following:</p> <ul style="list-style-type: none"> <li>• Not Applicable – The administrative status of UDLD is globally disabled or disabled on the interface.</li> <li>• Bidirectional – UDLD has detected a bidirectional link.</li> <li>• Shutdown – UDLD has detected a unidirectional link, and the port is in a disabled state. To clear the disabled state, click UDLD Port Reset.</li> <li>• Undetermined – UDLD has not collected enough information to determine the state of the port.</li> <li>• Unknown – The port link has physically gone down, but it is not because it was put in a disabled state by the UDLD feature.</li> </ul>

Click Refresh to display the latest information from the router.

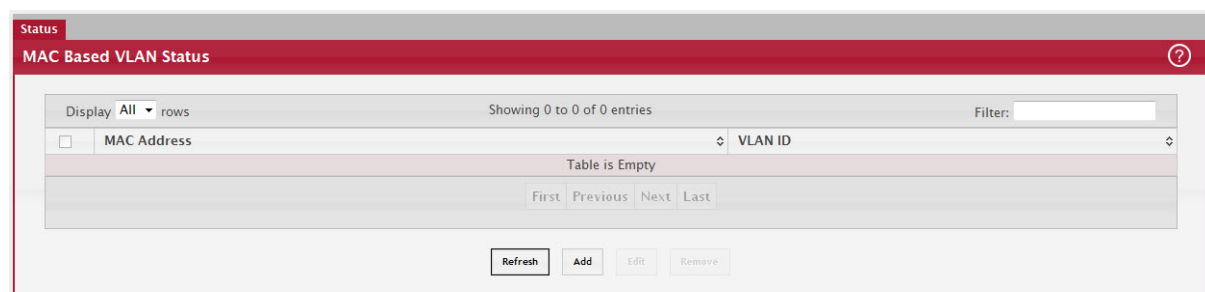
If you change any information on the page, click Submit to apply the changes to the system.

### 5.3 MAC Based VLAN Status

Use this page to add, edit, or remove MAC-based VLANs. MAC-based VLANs allow incoming untagged packets to be assigned to a VLAN based on the source MAC address of the packet. This type of VLAN is useful when a host might not always connect to the network through the same port but needs to be on the same VLAN.

To access the MAC Based VLAN Status page, click Switching > MAC Based VLAN > Status in the navigation menu.

Figure 169: MAC Based VLAN Status



Use the buttons to perform the following tasks:

- To add a MAC-based VLAN, click Add and specify a MAC address and a VLAN ID in the available fields.
- To change the VLAN ID of a configured MAC-based VLAN, select the entry to modify and click Edit. Then, configure the desired VLAN ID.
- To remove one or more configured MAC-based VLANs, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.



Table 155: MAC Based VLAN Status Fields

Field	Description
MAC Address	The source MAC address of the host. All untagged traffic that includes this address in the source MAC address field of the Ethernet frame is placed in the associated VLAN.
VLAN ID	The VLAN ID of the MAC-based VLAN. If an untagged frame received on any port or LAG matches the associated MAC address, it is tagged with this VLAN ID.

Click Refresh to display the latest information from the router.

## 5.4 Double VLAN (DVLAN) Tunneling

DVLAN Tunneling allows the use of a second tag on network traffic. The additional tag helps differentiate between customers in the Metropolitan Area Networks (MAN) while preserving individual customer's VLAN identification when they enter their own 802.1Q domain.

With the introduction of this second tag, you do not need to divide the 4k VLAN ID space to send traffic on an Ethernet-based MAN.

With DVLAN Tunneling enabled, every frame that is transmitted from an interface has a new VLAN tag (S-tag) attached while every packet that is received from an interface has a VLAN tag (S-tag) removed (if one or more tags are present).

DVLAN also supports up to 4 Tag Protocol Identifier (TPID) values per switch and the ability to map these values to ports. This allows you to configure the same or different TPIDs for different ports. Use the DVLAN Tunneling page to configure DVLAN frame tagging on one or more ports.

### 5.4.1 DVLAN Configuration

The DVLAN Config page allows you to configure the TPID with an associated Global EtherType for all ports on the system. To access the DVLAN Configuration page, click Switching > DVLAN > Configuration in the navigation menu.

Figure 170: DVLAN Configuration

The screenshot shows the DVLAN Configuration page. At the top, there are tabs for 'Configuration', 'Summary', and 'Interface Summary'. The main title is 'DVLAN Configuration'. Below this, there is a 'Primary TPID' field containing the value '0x8100'. Underneath is a table for 'Secondary TPIDs' which is currently empty, with a '+ -' button to the right. At the bottom of the form is a 'Refresh' button.

Table 156: DVLAN Configuration Fields

Field	Description
Primary TPID	<p>The two-byte hex EtherType value to be used as the first 16 bits of the DVLAN tag. The value configured in this field is used as the primary TPID for all interfaces that are enabled for DVLAN tagging. The Primary TPID can be one of the following:</p> <ul style="list-style-type: none"> <li>0x8100 – IEEE 802.1Q customer VLAN tag type</li> <li>0x88a8 – Virtual Metropolitan Area Network (VLAN) tag type</li> <li>Custom Tag – User-defined EtherType value</li> </ul> <p>To change the Primary TPID, click the Edit icon and select an option from the menu.</p>
Secondary TPIDs	<p>The two-byte hex EtherType values available to be configured as secondary TPIDs. Only the options you configure as Secondary TPIDs can be selected as the Primary TPID. To add Secondary TPIDs to the list, click the + (plus) symbol and select one or more of the following options:</p> <ul style="list-style-type: none"> <li>802.1Q Tag – IEEE 802.1Q customer VLAN tag type, represented by the EtherType value 0x8100. This value indicates that the frame includes a VLAN tag. If this value is already configured as a primary or secondary TPID, it cannot be selected.</li> <li>vMAN Tag – Virtual Metropolitan Area Network (VLAN) tag type, represented by the EtherType value 0x88a8. This value indicates that the frame is DVLAN tagged. If this value is already configured as a primary or secondary TPID, it cannot be selected.</li> <li>Custom Tag – User-defined EtherType value. If you select this option, specify the EtherType value in the available field.</li> </ul> <p>To remove a TPID from the list, click the – (minus) symbol associated with the entry. To remove all TPID entries from the list, select the – (minus) symbol in the header row and confirm the action.</p>

If you make any changes to the page, click Submit to apply the changes to the system.

## 5.4.2 DVLAN Summary

The DVLAN Summary page allows you to view the Global and Default TPIDs configured for all ports on the system.

To access the DVLAN Summary page, click Switching > DVLAN > Summary in the navigation menu.

Figure 171: DVLAN Summary




Table 157: DVLAN Summary Fields

Field	Description
Primary TPIDs	The two-byte hex EtherType value used as the first 16 bits of the DVLAN tag. This value identifies the frame as one of the following types: <ul style="list-style-type: none"> <li>0x8100 – IEEE 802.1Q VLAN tag type. This value indicates that the frame includes a VLAN tag.</li> <li>0x88a8 – Virtual Metropolitan Area Network (VMAN) tag type. This value indicates that the frame is double VLAN tagged.</li> <li>Custom Tag – Any TPID value other than 0x8100 or 0x88a8 is a user-defined EtherType value.</li> </ul>
Secondary TPID	The two-byte hex EtherType values configured as secondary TPIDs.

Click Refresh to display the latest information from the router.

### 5.4.3 DVLAN Interface Summary

Use this page to view and configure the double VLAN (DVLAN) tag settings for each interface. Double VLAN tagging allows service providers to create Virtual Metropolitan Area Networks (VMANs). With DVLAN tagging, service providers can pass VLAN traffic from one customer domain to another through a metro core. By using an additional tag on the traffic, the interface can differentiate between customers in the MAN while preserving an individual customer's VLAN identification that is used when the traffic enters the customer's 802.1Q domain.

To access the DVLAN Interface Summary page, click Switching > DVLAN > Interface Summary in the navigation menu.

To configure the DVLAN settings for an interface, select the interface to configure and click Edit.

Figure 172: DVLAN Interface Summary

The screenshot displays the 'DVLAN Interface Summary' page. At the top, there are navigation tabs for 'Configuration', 'Summary', and 'Interface Summary'. Below the tabs, the page title 'DVLAN Interface Summary' is shown with a help icon. The main content area features a table with the following data:

Display	10 rows	Showing 1 to 10 of 168 entries	Filter:
<input type="checkbox"/>	Interface	Interface Mode	Interface EtherType
<input type="checkbox"/>	1/0/1	Disable	0x8100
<input type="checkbox"/>	1/0/2	Disable	0x8100
<input type="checkbox"/>	1/0/3	Disable	0x8100
<input type="checkbox"/>	1/0/4	Disable	0x8100
<input type="checkbox"/>	1/0/5	Disable	0x8100
<input type="checkbox"/>	1/0/6	Disable	0x8100
<input type="checkbox"/>	1/0/7	Disable	0x8100
<input type="checkbox"/>	1/0/8	Disable	0x8100
<input type="checkbox"/>	1/0/9	Disable	0x8100
<input type="checkbox"/>	1/0/10	Disable	0x8100

Below the table, there are navigation buttons: 'First', 'Previous', '1', '2', '3', '4', '5', 'Next', and 'Last'. At the bottom of the page, there are 'Edit' and 'Refresh' buttons.

Table 158: DVLAN Interface Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Interface Mode	The administrative mode of double VLAN tagging on the interface. When DVLAN tagging is enabled, every frame that is transmitted from the interface has a DVLAN tag attached, and every packet that is received from the interface has a tag removed (if one or more tags are present).
Interface EtherType	The EtherType value to be used as the first 16 bits of the DVLAN tag. If one or more secondary TPIDs have been configured for the interface, these EtherType values are also displayed.
EtherType (Primary TPID)	The EtherType value to be used as the first 16 bits of the DVLAN tag. This is a global value that is configured on the DVLAN Configuration page.
Secondary TPIDs	The EtherType value(s) available to be configured as secondary TPIDs. To add a secondary TPID, the DVLAN Interface Mode must first be enabled. Then, select the entry in the Secondary TPIDs field and click the right arrow button. The entry moves into the Configured TPIDs field.
Configured TPIDs	The EtherType value(s) configured as secondary TPIDs. To remove a configured secondary TPID, enable the DVLAN Interface Mode, select the entry to remove from the Configured TPIDs field and click the left arrow button. The entry returns to the Secondary TPIDs field.

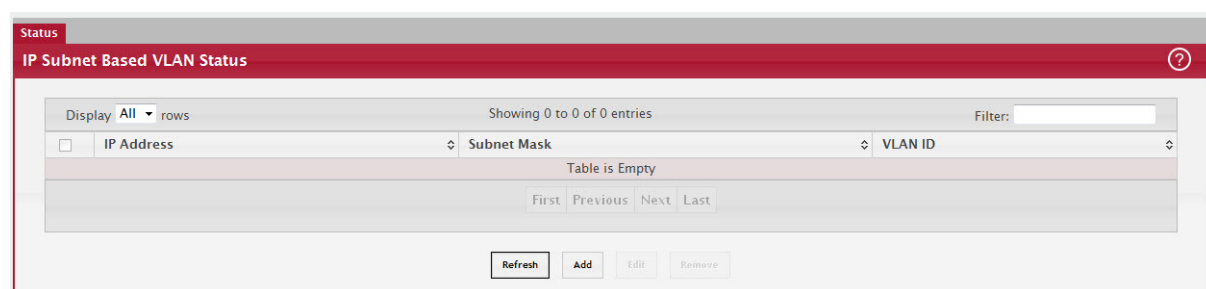
Click Refresh to redisplay the most current information from the router.

## 5.5 IP Subnet Based VLAN

Use this page to add, edit, and remove IP subnet-based VLANs. IP subnet-based VLANs allow incoming untagged packets to be assigned to a VLAN based on the source IP address of the packet. All hosts in the same subnet are members of the same VLAN.

To display the IP Subnet Based VLAN Status page, click Switching > IP Subnet Based VLAN > Status.

Figure 173: IP Subnet Based VLAN Status



Use the buttons to perform the following tasks:

- To add an IP subnet-based VLAN, click Add and specify an IP address, subnet mask, and VLAN ID in the available fields.
- To change the VLAN ID of a configured IP subnet-based VLAN, select the entry to modify and click Edit. Then, configure the desired VLAN ID.
- To remove one or more configured IP subnet-based VLANs, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 159: IP Subnet Based VLAN Status Fields

Field	Description
IP Address	The network address for the IP subnet. All incoming untagged packets that have a source IP address within the defined subnetwork are placed in the same VLAN.
Subnet Mask	The subnet mask that defines the network portion of the IP address.
VLAN ID	The VLAN ID of the IP subnet-based VLAN. If the source IP address of untagged traffic received on any port or LAG is within the associated IP subnet, the traffic is tagged with this VLAN ID.

Click Refresh to redisplay the most current information from the router.

## 5.6 Protocol Based VLAN Configuration

This page is divided into two sections:

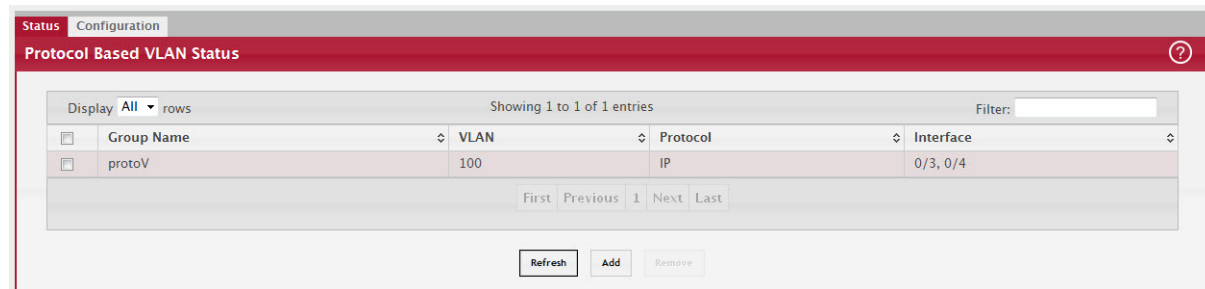
- Status
- Configuration

### 5.6.1 Status

Use this page to add and remove Protocol-based Virtual Local Area Networks (PBVLANS). In a PBVLAN, traffic is bridged through specified ports based on the protocol. PBVLANS allow you to define a packet filter that the device uses as the matching criteria to determine whether a particular packet belongs to a particular VLAN. PBVLANS are most often used in environments where network segments contain hosts running multiple protocols. PBVLANS can help optimize network traffic patterns because protocol-specific broadcast messages are sent only to hosts that use the protocols specified in the PBVLAN.

To display the Protocol Based VLAN Status page, click Switching > Protocol Based VLAN > Status.

Figure 174: Protocol Based VLAN Status



#### 5.6.1.1 Adding a PBVLAN

1. To add a PBVLAN, click Add and specify a group name, VLAN ID, protocol, and interfaces in the available fields.
2. To remove one or more configured PBVLANS, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 160: Protocol Based VLAN Status Fields

Field	Description
Group Name	The user-configured name that identifies the PBVLAN group.
VLAN	The VLAN ID associated with the PBVLAN. VLAN tagging for the PBVLAN works as follows: <ul style="list-style-type: none"> <li>• If the frame received over a port is tagged, normal processing takes place.</li> <li>• If the frame received over a port is untagged, the frame type is matched according to the protocol(s) assigned to the group on that port. <ul style="list-style-type: none"> <li>- If a match is found, the frame is assigned the VLAN ID specified for the group.</li> <li>- If a match is not found, the frame is assigned the port VID (PVID) as its VLAN ID.</li> </ul> </li> </ul>
Protocol	The protocol or protocols to use as the match criteria for an Ethernet frame. The protocol is included in the two-byte EtherType field of the frame. When adding a PBVLAN, you can specify the EtherType hex value or (for IP, ARP, and IPX) the protocol keyword.
Interface	The interfaces that are members of the PBVLAN group. An interface can be a member of multiple groups as long as the same protocol is not specified in more than one group that includes the interface. When adding a PBVLAN group, use the Available Interfaces and Group Interfaces fields to configure the interfaces that are members of the PBVLAN group. To move an interface between the Available Interfaces and Group Interfaces fields, click the interface (or CTRL + click to select multiple interfaces), and then click the appropriate arrow to move the selected interfaces to the desired field.

- Click Refresh to display the latest information from the router.

## 5.6.2 Configuration

Use this page to configure existing Protocol-based VLAN (PBVLAN) groups. You can change the group name, VLAN ID, protocol information, and interfaces associated with the PBVLAN group.

To display the Protocol Based VLAN Status page, click Switching > Protocol Based VLAN > Configuration.

Figure 175: Protocol Based VLAN Configuration

The screenshot displays the 'Protocol Based VLAN Group Configuration' interface. At the top, there are tabs for 'Status' and 'Configuration'. The main title is 'Protocol Based VLAN Group Configuration'. Below the title, there are several input fields: 'Group Name' (set to 'protoV'), 'VLAN' (set to '100'), and 'Protocol' (set to 'IP'). There are also two lists: 'Available Interfaces' (0/1, 0/2, 0/5, 0/6, 0/7, 0/8, 1/1, 1/2) and 'Group Interfaces' (0/3, 0/4). At the bottom, there are three buttons: 'Submit', 'Refresh', and 'Cancel'.

Table 161: Protocol Based VLAN Configuration Fields

Field	Description
Group Name	To change the properties of a PBVLAN, select its name from the Group Name menu. The Group Name field allows you to update the name of the PBVLAN group.
VLAN	The VLAN ID associated with the PBVLAN. Untagged traffic that matches the protocol criteria is tagged with this VLAN ID.
Protocol	<p>The protocol or protocols to use as the match criteria to determine whether a particular packet belongs to the PBVLAN. The protocols in this list are checked against the two-byte EtherType field of ingress Ethernet frames on the PVBLAN Group Interfaces. When adding a protocol, you can specify the EtherType hex value or (for IP, ARP, and IPX) the protocol keyword.</p> <p>To configure the protocols associated with a PBVLAN group, use the buttons available in the protocol table:</p> <ul style="list-style-type: none"> <li>• To add a protocol to the group, click the + (plus) button and enter the protocol to add.</li> <li>• To delete an entry from the list, click the – (minus) button associated with the entry to remove.</li> <li>• To delete all entries from the list, click the – (minus) button in the heading row.</li> </ul>
Available Interfaces	The interfaces that can be added to the PBVLAN group. An interface can be a member of multiple groups as long as the same protocol is not specified in more than one group that includes the interface. To move an interface between the Available Interfaces and Group Interfaces fields, click the interface (or CTRL + click to select multiple interfaces), and then click the appropriate arrow to move the selected interfaces to the desired field.
Group Interfaces	The interfaces that are members of the PBVLAN group.

- If you make any changes, click Submit to apply the change to the system.
- Click Refresh to display the latest information from the router.

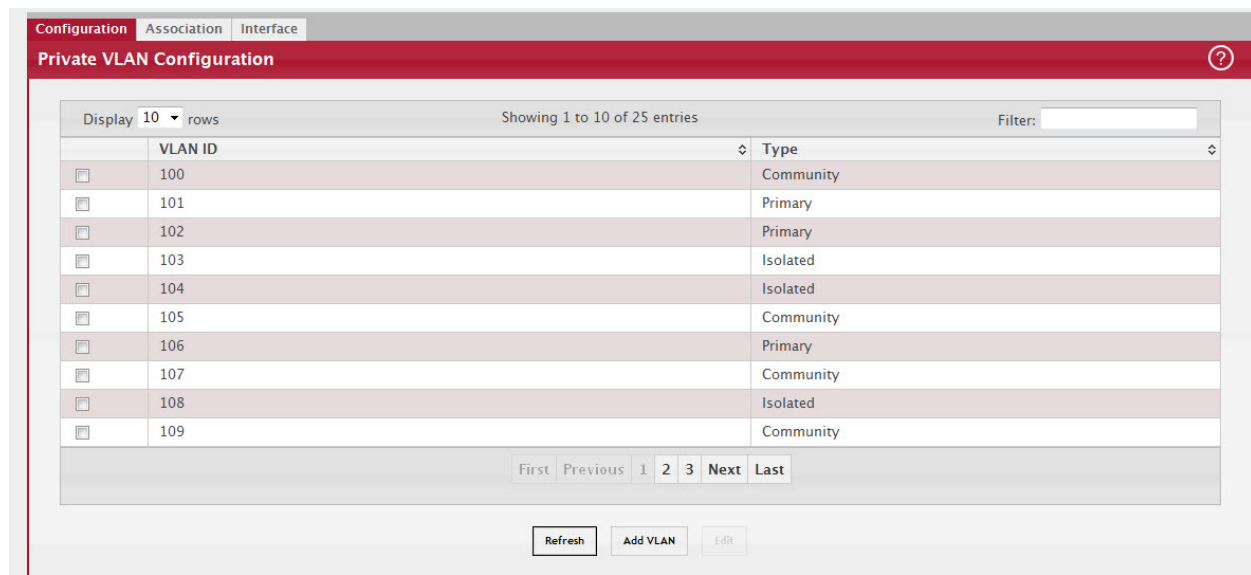
## 5.7 Private VLAN

Use this screen to add Virtual Local Area Networks (VLANs) to the device and to configure existing VLANs as private VLANs. Private VLANs provide Layer 2 isolation between ports that share the same broadcast domain. In other words, a private VLAN allows a VLAN broadcast domain to be partitioned into smaller point-to-multipoint subdomains. The ports participating in a private VLAN can be located anywhere in the Layer 2 network. Each subdomain is defined (represented) by a primary VLAN and a secondary VLAN. The primary VLAN ID is the same for all subdomains that belong to a private VLAN. The secondary VLAN ID differentiates subdomains from each another and provides Layer 2 isolation between ports that are members of the same private VLAN.

### 5.7.1 Private VLAN Configuration

To access the Private VLAN Configuration page, click Switching > Private VLAN > Configuration in the navigation menu.

Figure 176: Private VLAN Configuration



Use the buttons to perform the following tasks:

- To add a VLAN, click Add VLAN and specify the VLAN ID(s) in the available field.
- To configure a private VLAN, select the entry to modify and click Edit. Then, configure the desired private VLAN setting.

### NOTICE

Default VLAN and management VLAN cannot be configured as a private VLANs and hence are not displayed on this page.

Table 162: Private VLAN Configuration Fields

Field	Description
VLAN ID	Displays the VLAN ID for which Private VLAN type is being set.
Type	<p>Use the Private VLAN Type menu to select the type of private VLAN. The factory default is Unconfigured.</p> <ul style="list-style-type: none"> <li>• Primary – A private VLAN that forwards the traffic from the promiscuous ports to isolated ports, community ports, and other promiscuous ports in the same private VLAN. Only one primary VLAN can be configured per private VLAN. All ports within a private VLAN share the same primary VLAN.</li> <li>• Isolated – A secondary VLAN that carries traffic from isolated ports to promiscuous ports. Only one isolated VLAN can be configured per private VLAN.</li> <li>• Community – A secondary VLAN that forwards traffic between ports that belong to the same community and to the promiscuous ports. Multiple community VLANs can be configured per private VLAN.</li> <li>• Unconfigured – The VLAN is not configured as a private VLAN.</li> </ul>

Click Refresh to display the latest information from the router.

## 5.7.2 Private VLAN Association

Use this page to configure the association between the primary VLAN and secondary VLANs. Associating a secondary VLAN with a primary VLAN allows host ports in the secondary VLAN to communicate outside the private VLAN. To configure a primary VLAN association, select the entry to modify and click Edit.

To access the Private VLAN Association page, click Switching > Private VLAN > Association in the navigation menu.



Figure 177: Private VLAN Association

Primary VLAN	Isolated VLAN	Community VLAN
101	103	107, 109
102	108	100
106	104	105

Use the buttons to perform the following tasks:

- To configure a primary VLAN association, select each entry to modify and click Edit.

---

Isolated VLANs and Community VLANs are collectively called Secondary VLANs.

### NOTICE

Table 163: Private VLAN Association Fields

Field	Description
Primary VLAN	The VLAN ID of each VLAN configured as a primary VLAN.
Isolated VLAN	The VLAN ID of the isolated VLAN associated with the primary VLAN. If the field is blank, no isolated VLAN has been associated with the primary VLAN. An isolated VLAN is a secondary VLAN that carries traffic from isolated ports to promiscuous ports. Only one isolated VLAN can be configured per private VLAN.
Community VLAN	The VLAN ID of each community VLAN associated with the primary VLAN. If the field is blank, no community VLANs have been associated with the primary VLAN. A community VLAN is a secondary VLAN that forwards traffic between ports that belong to the same community and to the promiscuous ports. Multiple community VLANs can be configured per private VLAN.

After you click Edit, the Edit Private VLAN Association window opens and allows you to create associations with the selected primary VLAN. The following information describes the field in this window.

- Secondary VLAN – The isolated or community VLANs that can be associated with the primary VLAN. Secondary VLANs that are already associated with a primary VLAN do not appear in the list and cannot be associated with another primary VLAN. To select multiple secondary VLANs, Ctrl + click each VLAN to associate with the primary VLAN.

Click Refresh to display the latest information from the router.

### 5.7.3 Private VLAN Interface

The private VLAN interface association page allows you to configure the port mode for the ports and LAGs that belong to a private VLAN and to configure associations between interfaces and primary/secondary private VLANs.

To access the Private VLAN Interface page, click Switching > Private VLAN > Interface in the navigation menu.

Figure 178: Private VLAN Interface

Interface	Mode	Host Primary VLAN	Host Secondary VLAN	Promiscuous Primary VLAN	Promiscuous Secondary VLAN	Promiscuous Trunk Primary VLAN	Promiscuous Trunk Secondary VLAN	Promiscuous Trunk Native VLAN	Promiscuous Trunk Allowed VLAN	Operational Private VLAN
1/0/1	Host	4	8							4, 8
1/0/2	General									
1/0/3	Promiscuous Trunk							2		
1/0/4	General									
1/0/5	Promiscuous			4	8					4, 8
1/0/6	General									
1/0/7	General									
1/0/8	General									
1/0/9	General									
1/0/10	General									

Use the buttons to perform the following tasks:

- To configure the port mode and private VLAN-to-interface associations, select the entry to modify and click Edit.
- To remove the association between an interface and the primary/secondary private VLANs that the interface belongs to when it operates in host mode, select each interface with the association to clear and click Remove Host Association. You must confirm the action before the host association for the entry is cleared.
- To remove the association between an interface and the primary/secondary private VLANs that the interface belongs to when it operates in promiscuous mode, select each interface with the association to clear and click Remove Promiscuous Association. You must confirm the action before the promiscuous association for the entry is cleared.
- To remove the association between an interface and the primary/secondary promiscuous trunk private VLANs that the interface belongs to when it operates in promiscuous trunk mode, select each interface with the association to clear and click Remove Promiscuous Trunk Association. You must confirm the action before the promiscuous association for the entry is cleared.
- To remove the association between an interface and the primary/secondary private VLANs that the interface belongs to when it operates in isolated trunk mode, select each interface with the association to clear and click Remove Isolated Trunk Association. You must confirm the action before the isolated association for the entry is cleared.

Table 164: Private VLAN Interface Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When editing interface settings, this field identifies the interface being configured.
Mode	The private VLAN mode of the interface, which is one of the following: <ul style="list-style-type: none"> <li>• General – The interface is in general mode and is not a member of a private VLAN.</li> <li>• Promiscuous – The interface belongs to a primary VLAN and can communicate with all interfaces in the private VLAN, including other promiscuous ports, community ports, and isolated ports.</li> <li>• Isolated Trunk – The interface also belongs to a primary VLAN. It carries traffic from isolated ports to promiscuous ports. Only one isolated VLAN can be configured per private VLAN. An isolated trunk port carries tagged traffic of multiple isolated VLANs and normal VLANs.</li> <li>• Promiscuous Trunk – The interface belongs to a primary VLAN and can communicate with all interfaces in the private VLAN, including other promiscuous trunk ports, community ports, and isolated ports.</li> <li>• Host – The interface belongs to a secondary VLAN and, depending upon the type of secondary VLAN, can either communicate with other ports in the same community (if the secondary VLAN is a community VLAN) and with the promiscuous ports or is able to communicate only with the promiscuous ports (if the secondary VLAN is an isolated VLAN).</li> </ul>

Table 164: Private VLAN Interface Fields (Continued)

Field	Description
Host Primary VLAN	The primary private VLAN the port is a member of when it is configured to operate in Host mode.
Host Secondary VLAN	The secondary private VLAN the port is a member of when it is configured to operate in Host mode. The secondary private VLAN is either an isolated or community VLAN.
Isolated Trunk Primary VLAN	The primary private VLAN in which the port is a member when it is configured to operate in Isolated Trunk mode.
Isolated Trunk Secondary VLAN	The secondary private VLAN the port is a member of when it is configured to operate in Isolated Trunk mode. The secondary private VLAN must be an isolated VLAN.
Promiscuous Primary VLAN	The primary private VLAN in which the port is a member when it is configured to operate in Promiscuous mode.
Promiscuous Secondary VLAN	The secondary private VLAN the port is a member of when it is configured to operate in Promiscuous mode. The secondary private VLAN is either an isolated or community VLAN.
Promiscuous Trunk Primary VLAN	The primary private VLAN in which the port is a member when it is configured to operate in Promiscuous Trunk mode.
Promiscuous Trunk Secondary VLAN	The secondary private VLANs the port is a member of when it is configured to operate in Promiscuous Trunk mode. The secondary private VLANs are either isolated or community VLANs.
Trunk Native VLAN	When it is configured to operate in Isolated or Promiscuous Trunk mode, defines VLAN association for untagged packets. If not configured, untagged packets are dropped.
Trunk Allowed VLAN	The list of allowed normal VLANs on the trunk port when it is configured to operate in Promiscuous or Isolated Trunk mode.
Operational Private VLAN	The primary and secondary operational private VLANs for the interface. The VLANs that are operational depend on the configured mode for the interface and the private VLAN type.

Click Refresh to display the latest information from the router.

## 5.8 Voice VLAN Configuration

The voice VLAN feature enables switch ports to carry voice traffic with defined settings so that voice and data traffic are separated when coming onto the port. A voice VLAN ensures that the sound quality of an IP phone is safeguarded from deterioration when data traffic on the port is high.

The inherent isolation provided by VLANs ensures that inter-VLAN traffic is under management control and that network-attached clients cannot initiate a direct attack on voice components. A QoS protocol based on the IEEE 802.1P class-of-service (CoS) protocol uses classification and scheduling to send network traffic from the switch in a predictable manner. The system uses the source MAC of the traffic traveling through the port to identify the IP phone data flow.

Voice VLAN is enabled per-port basis. A port can participate only in one voice VLAN at a time. The Voice VLAN feature is disabled by default.

To display the Voice VLAN Configuration page, click Switching > Voice VLAN > Configuration.

Figure 179: Voice VLAN Configuration

The screenshot shows the 'Voice VLAN Configuration' page. At the top, there are two tabs: 'Configuration' (active) and 'Interface Summary'. Below the tabs is a red header bar with the title 'Voice VLAN Configuration' and a help icon. The main content area contains a form for 'Voice VLAN Admin Mode' with two radio buttons: 'Enable' and 'Disable'. The 'Disable' radio button is selected. Below the form are three buttons: 'Submit', 'Refresh', and 'Cancel'.

Table 165: Voice VLAN Configuration Fields

Field	Description
Voice VLAN Admin Mode	Click Enable or Disable to administratively turn the Voice VLAN feature on or off for all ports. The administrative mode of the Voice VLAN feature. When Voice VLAN is enabled globally and configured on interfaces that carry voice traffic, this feature can help ensure that the sound quality of an IP phone does not deteriorate when data traffic on the port is high.

- If you make any changes, click Submit to apply the change to the system.
- Click Refresh to display the latest information from the router.

## 5.9 Voice VLAN Interface

Use this page to configure the per-port settings for the Voice VLAN feature. When Voice VLAN is configured on a port that receives both voice and data traffic, it can help ensure that the voice traffic has priority.

Use the buttons to perform the following tasks:

- To configure Voice VLAN settings on a port, click Add. Select the interface to configure from the Interface menu, and then configure the desired settings.
- To change the Voice VLAN settings, select the interface to modify and click Edit.
- To remove the Voice VLAN configuration from one or more ports, select each entry to delete and click Remove.

To display the Voice VLAN Interface page, click Switching > Voice VLAN > Interface Summary.

Figure 180: Voice VLAN Interface

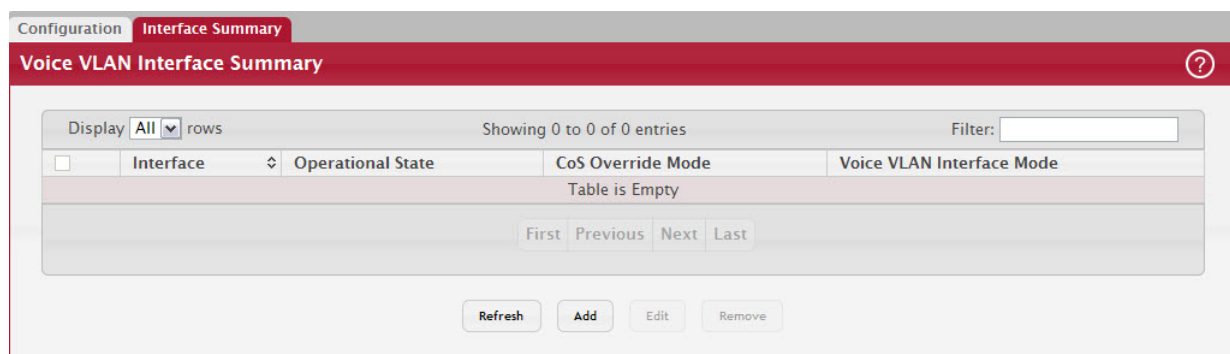


Table 166: Voice VLAN Interface Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When adding a Voice VLAN configuration to a port, the Interface menu allows you to select the port to configure. Only interfaces that have not been configured with Voice VLAN settings can be selected from the menu.
Operational State	The operational status of the Voice VLAN feature on the interface. To be enabled, Voice VLAN must be globally enabled and enabled on the interface. Additionally, the interface must be up and have a link.
CoS Override Mode	The Class of Service override mode: <ul style="list-style-type: none"> <li>• Enabled – The port ignores the 802.1p priority value in the Ethernet frames it receives from connected devices.</li> <li>• Disabled – The port trusts the priority value in the received frame.</li> </ul>

Table 166: Voice VLAN Interface Fields (Continued)

Field	Description
Voice VLAN Interface Mode	Indicates how an IP phone connected to the port should send voice traffic: <ul style="list-style-type: none"> <li>• LAN ID – Forward voice traffic in the specified voice VLAN.</li> <li>• Dot1p – Tag voice traffic with the specified 802.1p priority value.</li> <li>• None – Use the settings configured on the IP phone to send untagged voice traffic.</li> <li>• Untagged – Send untagged voice traffic.</li> <li>• Disable – Operationally disables the Voice VLAN feature on the interface.</li> </ul>
Voice VLAN Interface Value	When adding or editing Voice VLAN settings for an interface and either VLAN ID or Dot1p is selected as the Voice VLAN Interface Mode, specify the voice VLAN ID or the Dot1p priority value that the connected IP phone should use for voice traffic.

- If you make any changes, click Submit to apply the change to the system.
- Click Refresh to display the latest information from the router.

## 5.10 Virtual Port Channel Configuration

Use this page to view and manage global virtual port channel (VPC) settings on the device. VPCs are also known as multichassis or multiswitch link aggregation groups (MLAGs). Like port channels (also known as link aggregation groups or LAGs), VPCs allow one or more Ethernet links to be aggregated together to increase speed and provide redundancy. With port channels, the aggregated links must be on the same physical device, but VPCs do not share that requirement. The VPC feature allows links on two different switches to pair with links on a partner device. The partner device is unaware that it is pairing with two different devices to form a port channel.

To display the Virtual Port Channel Configuration page, click Switching > Virtual Port Channel > Global.

Figure 181: Virtual Port Channel Global Configuration

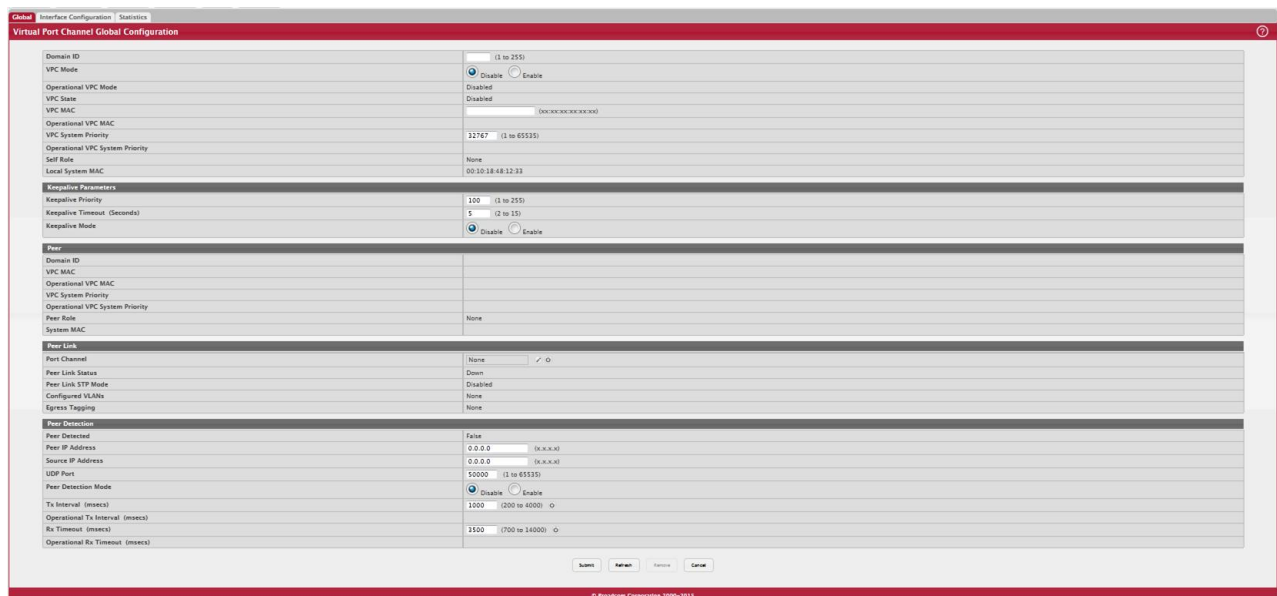


Table 167: Virtual Port Channel Configuration Fields

Field	Description
Domain ID	The ID of the VPC domain. Only one VPC domain can be created on a given device. The VPC domain ID should be equal to the domain ID of the peer to form a VPC pair. The domain IDs are exchanged during role election and if different, VPC does not become operational.
VPC Mode	The administrative mode of VPC on the system.
Operational VPC Mode	The operational mode of VPC on the system. For the VPC to be operational, several conditions must be met including the following: <ul style="list-style-type: none"> <li>• The VPC administrative mode is globally enabled.</li> <li>• Peer links are configured.</li> <li>• The Keepalive mode is enabled.</li> </ul>
VPC State	The VPC state, which is one of the following: <ul style="list-style-type: none"> <li>• Disable – The VPC mode is not operational.</li> <li>• Listen – The keepalive component does not advertise any packets. It listens for advertisements from a peer.</li> <li>• Ready – The keepalive component starts sending periodic keepalive messages.</li> <li>• Primary – Traffic over VPC interfaces is allowed to be forwarded in this state. The keepalive component continues to advertise keepalive messages with the state as Primary and monitors the health of the secondary device.</li> <li>• Secondary – Traffic over VPC interfaces is allowed to be forwarded in this state. The keepalive component continues to advertise keepalive messages with the state as Secondary and monitors the health of the primary device.</li> </ul>
VPC MAC	The MAC address of the VPC domain. VPC MAC must be same on both the peer devices. The MAC address should be unicast and not be equal to the system MAC of either the primary or secondary VPC device. MAC addresses are exchanged during role election and if different, VPC does not become operational.
Operational VPC MAC	The VPC MAC address agreed upon by both peers during role election. This field is present in the keepalive message only if the transmitting peer is either primary or secondary.
VPC System Priority	The system priority of the VPC domain. System priority should be same on both peer devices for VPC to become operational.
Operational VPC System Priority	The VPC system priority agreed upon by both peers during role election. This field is present in the keepalive message only if the transmitting peer is either primary or secondary.
Self Role	The role of the local device in the VPC domain, which is Primary, Secondary, or None. The role is determined by an election between the two devices after a keepalive link is established. The primary device owns the VPC member ports on the secondary device and handles the control plane functionality of supported protocols for the VPC member ports on the secondary device.
Local System MAC	The MAC address of the local system.
Keepalive Parameters	The VPC feature sends periodic keepalive messages over the peer link between the primary and secondary devices in the VPC domain to determine the device roles (primary and secondary) and to monitor the health of the link.
Keepalive Priority	The priority value of the keepalive component on the local device. The device with lower priority value becomes the Primary device in the VPC role election.
Keepalive Timeout (Seconds)	The number of seconds that must pass without receiving a keepalive message before the peer device is considered to be down.
Keepalive Mode	The administrative mode of the keepalive component on the device.
Peer	The peer fields provide information about the peer device.
Domain ID	The ID of the peer VPC domain.
VPC MAC	The MAC address of the peer VPC domain.

**Table 167: Virtual Port Channel Configuration Fields (Continued)**

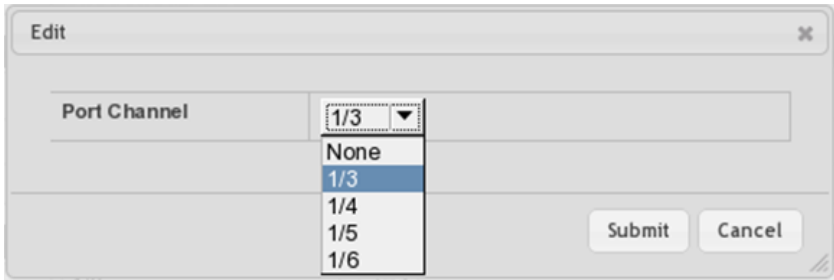
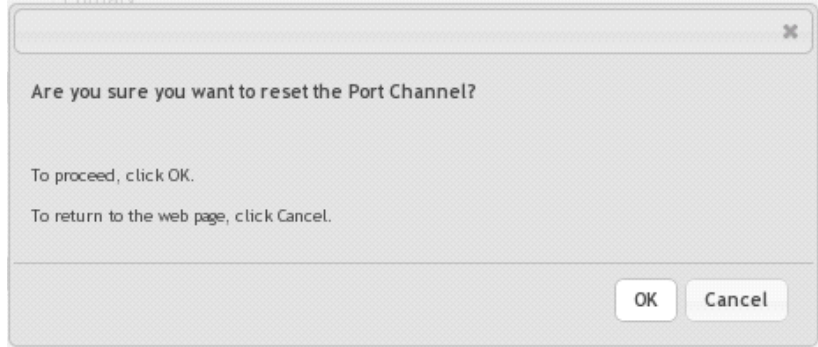
Field	Description
Operational VPC MAC	The VPC MAC address agreed upon by both peers during role election.
VPC System Priority	The system priority of the peer VPC domain.
Operational VPC System Priority	The VPC system priority agreed upon by both peers during role election.
Peer Role	The role of the peer device in the VPC domain, which is Primary, Secondary, or None.
System MAC	The MAC address of the peer system.
Peer Link	<p>The peer link is a port channel that serves as the link between the two devices in the VPC domain. Using a multimember port channel as the peer link helps protect it from link-level failures. The peer link is used:</p> <ul style="list-style-type: none"> <li>To carry the keepalive messages between the two peer devices.</li> <li>To carry the BPDUs and LACPDUs between the secondary and primary VPC devices.</li> <li>To carry control messages like VPC member port related events, FDB/MFDB entries, and configuration details.</li> <li>To carry data traffic over the peer's VPC member ports when the member ports of the VPC interface are all down on the local device.</li> </ul>
Port Channel	<p>The port channel on the local device used for the peer link. To configure the peer link, click the Edit icon next to the field.</p>  <p>The Edit window opens and allows you to select an available port channel from the Port Channel menu.</p> <p>To reset the port channel to the default value, click the Reset icon.</p> 
Peer Link Status	The port channel cannot be changed or reset when the Operational VPC Mode is Enabled.
Peer Link STP Mode	The operational status of the peer link, which is either Up or Down.
Configured VLANs	The spanning tree protocol (STP) mode of the port channel. When enabled, the port channel participates in the STP operation to help prevent network loops.
Egress Tagging	The VLAN ID of each VLAN in which the port channel participates.
	The VLAN ID tags included in the frames transmitted from the port channel.

Table 167: Virtual Port Channel Configuration Fields (Continued)

Field	Description
Peer Detection	The peer detection feature uses the dual control plane detection protocol (DCPDP), a UDP-based protocol, to detect peer links. You must configure peer detection on an IP interface with a VLAN that is not shared by any of the VPC interfaces.
Peer Detected	Indicates whether a peer link has been detected by DCPDP.
Peer IP Address	The IP address of the peer VPC device. This is the destination IP address in the DCPDP messages.
Source IP Address	The source IP address to be used by DCPDP.
UDP Port	The local UDP port to be used for listening to DCPDP packets.
Peer Detection Mode	The administrative mode of the peer detection feature (DCPDP).
Tx Interval	The interval in milliseconds between the DCPDP messages transmitted.
Operational Tx Interval	The operational transmit interval in milliseconds.
Rx Timeout	The DCPDP reception timeout in milliseconds.
Operational Rx Timeout	The operational timeout value in milliseconds.

- If you make any changes, click Submit to apply the change to the system.
- Click Refresh to display the latest information from the router.

### 5.10.1 Interface Configuration

Use this page to configure the VPC interfaces on the device. A VPC interface is created by combining a port channel on the local device with a port channel on the peer device. The VPC interface on the local and peer devices share a common VPC identifier. You can configure multiple instances of VPC interfaces on each peer device in the VPC domain.

To display the Virtual Port Channel Configuration page, click Switching > Virtual Port Channel > Interface Configuration.

Figure 182: Virtual Port Channel Interface Configuration

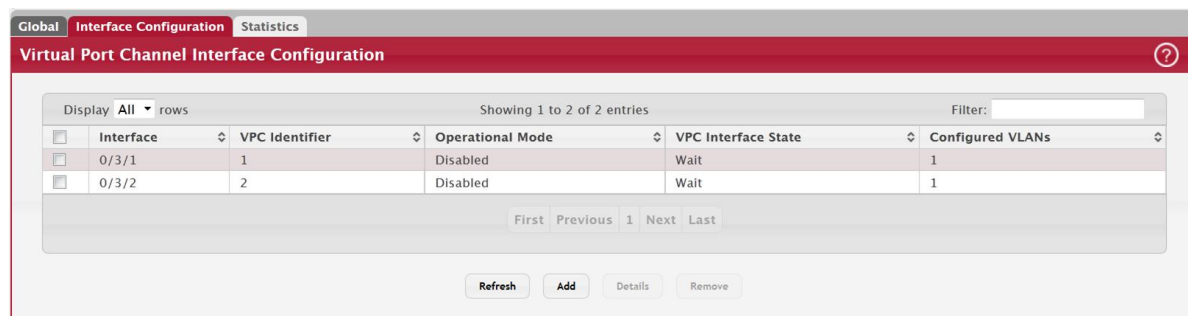


Table 168: Virtual Port Channel Interface Configuration Fields

Field	Description
Interface	The ID of the local port channel configured as a VPC interface.
VPC Identifier	The VPC identifier of the interface. To form a VPC interface with a port channel on the peer device, the port channel on the peer device must use the same VPC identifier.
Operational Mode	The operational mode of the VPC interface.

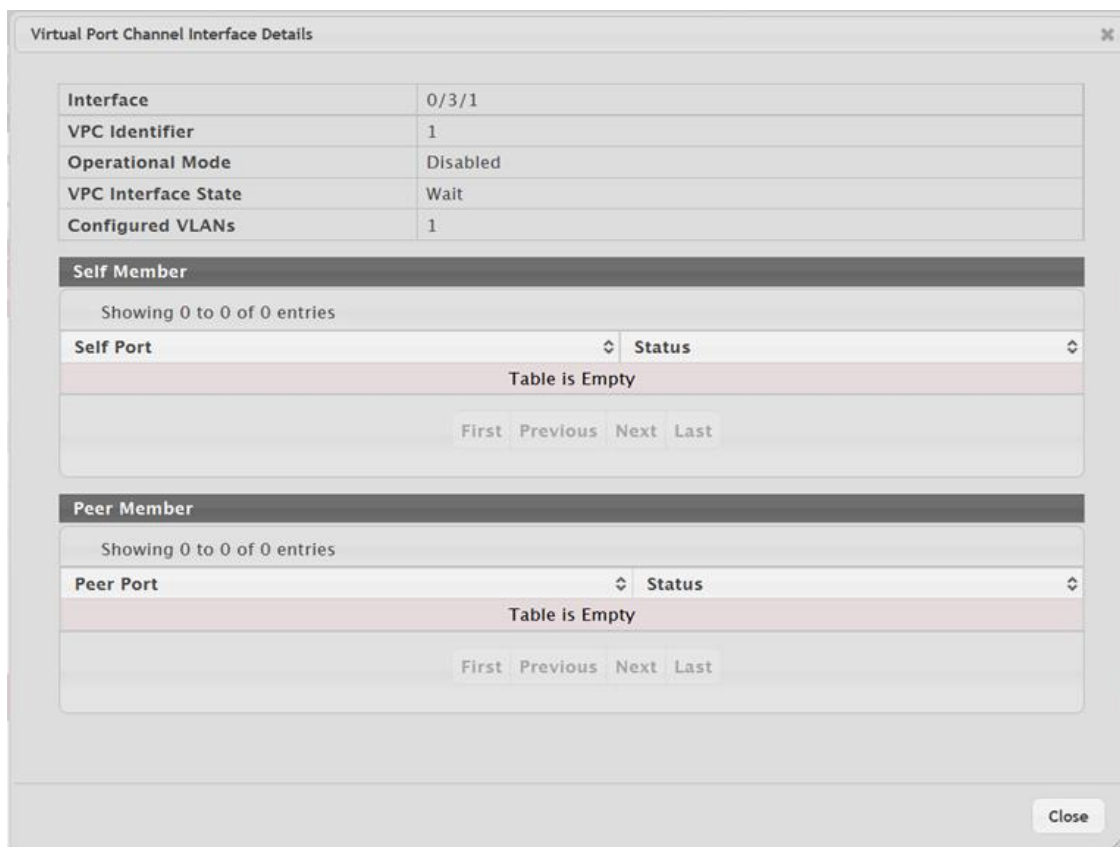


Table 168: Virtual Port Channel Interface Configuration Fields (Continued)

Field	Description
VPC Interface State	<p>The VPC interface state, which is one of the following:</p> <ul style="list-style-type: none"> <li>Disabled – VPC functionality is operationally disabled on the VPC interface.</li> <li>Wait – The port channel is waiting for VPC functionality to be enabled on a port channel on the peer device.</li> <li>Error – VPC functionality is enabled on a port channel on both peer devices, but not all entry criteria are met for the port channel to be operational. For example, if the combined number of member ports for the VPC interface is more than the maximum allowed, then the state is set to Error on both devices.</li> <li>Active – VPC functionality is enabled on a port channel on both peer devices, and all entry criteria are satisfied. The VPC interface is operationally enabled, and traffic is allowed to flow through the VPC member ports.</li> <li>Inactive – The links connected to the VPC member ports are down, but the VPC interface on the peer remains active.</li> </ul>
Configured VLANs	The VLAN ID of each VLAN in which the port channel participates.

After you select an interface and click Details, the Virtual Port Channel Interface Details window opens and displays additional information about the interface.

Figure 183: Virtual Port Channel Interface Details



The following information describes the additional fields in this window.

Table 169: Virtual Port Channel Interface Details

Field	Description
Self Member	The Self Member fields provide information about the VPC member ports on the local device.
Self Port	The ID of each port that is a member of the port channel configured as a VPC interface.
Status	The operational status of the port.
Peer Member	The Peer Member fields provide information about the VPC member ports on the peer device.
Peer Port	The ID of each port that is a member of the port channel configured as a VPC interface.
Status	The operational status of the port.

Use the buttons to perform the following tasks:

- To configure a port channel as a VPC interface, click Add and configure the desired settings.

Figure 184: Add VPC Interface Configuration

Table 170: Add VPC Interface Configuration Fields

Field	Description
Interface	The ID of the local port channel configured as a VPC interface.
VPC Identifier	The VPC identifier of the interface. To form a VPC interface with a port channel on the peer device, the port channel on the peer device must use the same VPC identifier.

- To view details about a VPC interface, select the interface with the information to view and click Details.
- To remove the VPC functionality from one or more port channels, select each entry to change and click Remove.

## 5.10.2 Statistics

This page shows information about the number of messages of various types sent between the two VPC peer devices over the peer link.

To display the Virtual Port Channel Statistics page, click Switching > Virtual Port Channel > Statistics.

Figure 185: Virtual Port Channel Statistics

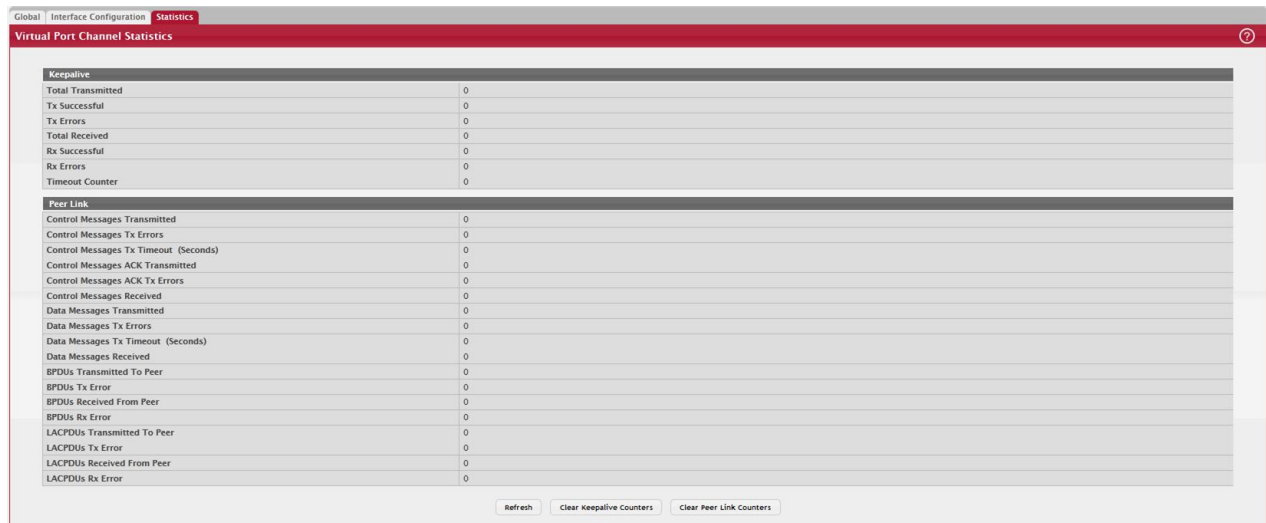


Table 171: Virtual Port Channel Statistics Fields

Field	Description
Keapalive	The VPC feature sends periodic keapalive messages over the peer link between the primary and secondary devices in the VPC domain to determine the device roles (primary and secondary) and to monitor the health of the link.
Total Transmitted	The total number of keapalive messages the local device has sent to the peer device.
Tx Successful	The number of keapalive messages that have been successfully transmitted from the local device.
Tx Errors	The number of keapalive messages that the local device attempted to send to the peer device that were not transmitted due to an error.
Total Received	The total number of keapalive messages the local device has received from the peer device.
Rx Successful	The number of keapalive messages the local device has successfully received from the peer device.
Rx Errors	The number of keapalive messages the local device has received from the peer device that contained errors.
Timeout Counter	The number of times the keapalive timeout timer has expired.
Peer Link	In addition to keapalive messages, the peer link is used to send and receive control messages, data messages, BPDUs, and LACPDU's between the peer devices.
Control Messages Transmitted	The number of control messages successfully sent from the local device to the peer device over the peer link.
Control Messages Tx Errors	The number of errors encountered when sending peer-link control messages from the local device to the peer device over the peer link.
Control Messages Tx Timeout (Seconds)	The number of peer-link control messages that did not receive an ACK from the peer device.
Control Messages ACK Transmitted	The number of ACKs sent to the peer device in response to peer-link control messages that were received.
Control Messages ACK Tx Errors	The number of errors encountered when sending ACKs in response to peer-link control messages.
Control Messages Received	The number of control messages successfully received by the local device from the peer device over the peer link.
Data Messages Transmitted	The number of data messages successfully sent from the local device to the peer device over the peer link.

**Table 171: Virtual Port Channel Statistics Fields (Continued)**

Field	Description
Data Messages Tx Errors	The number of errors encountered when sending peer-link data messages from the local device to the peer device over the peer link.
Data Messages Tx Timeout (Seconds)	The number of peer-link data messages that did not receive an ACK from the peer device.
Data Messages Received	The number of data messages successfully received by the local device from the peer device over the peer link.
BPDU's Transmitted To Peer	The number of BPDU's successfully sent to the peer device over the peer link.
BPDU's Tx Error	The number of errors encountered when sending BPDU's to the peer device.
BPDU's Received From Peer	The number of BPDU's successfully received from the peer device over the peer link.
BPDU's Rx Error	The number of errors encountered when receiving BPDU's from the peer device.
LACPDU's Transmitted To Peer	The number of LACPDU's successfully sent to the peer device over the peer link.
LACPDU's Tx Error	The number of errors encountered when sending LACPDU's to the peer device.
LACPDU's Received From Peer	The number of LACPDU's successfully received from the peer device over the peer link.
LACPDU's Rx Error	The number of errors encountered when receiving LACPDU's from the peer device.
Clear Keepalive Counters (Button)	Click this button to reset all keepalive message counters to 0.
Clear Peer Link Counters (Button)	Click this button to reset all peer link message counters to 0.

## 5.11 Port Auto Recovery

The Auto Recovery feature can automatically enable a disabled interface when the error conditions that caused the interface to be disabled are no longer detected. If Auto Recovery is not used (disabled), the interface remains disabled until an administrator manually enables it.

The switch supports an interface error disable feature that allows an interface to be automatically placed into a diagnostically disabled state when certain error conditions are detected on that interface. When an interface has been placed in a diagnostically disabled state, the interface is shut down, and no traffic is sent or received on that interface until it is either manually enabled by the administrator or re-enabled by the Auto Recovery feature after the recovery time interval has expired.

If the interface continues to encounter errors, it may be placed back into the diagnostically disabled state, and the interface will be disabled (link down). An interface in the diagnostically disabled state may also be manually recovered by enabling it from the Port Status page

### 5.11.1 Port Auto Recovery Configuration

Use the Port Auto Recovery Configuration page to allow a port to attempt to become re-enabled if it has been placed into a diagnostically disabled state due to the detection of certain error conditions.

To access the Port Auto Recovery Configuration page, click Switching > Auto Recovery > Configuration in the navigation menu.

Figure 186: Port Auto Recovery Configuration

**Configuration**

**Port Auto Recovery Configuration**

---

**Auto Recovery Components**

All Components	<input type="checkbox"/>
ARP Inspection	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
Authentication Manager	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
BPDU Guard	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
BPDU Rate Limit	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
Broadcast Storm Control	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
CoA Disable Host Port	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
Denial Of Service	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
DHCP Rate Limit	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
Keepalive	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
Link Flap	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
MAC Locking	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
Multicast Storm Control	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
UDLD	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
Unicast Storm Control	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled

---

**Auto Recovery Parameters**

Recovery Time (Seconds)  (30 to 86400, 300 = Default)

---

**D-Disabled Interface Status**

Display  rows Showing 0 to 0 of 0 entries Filter:

Interface	Admin Mode	Port Status	Error Disable Reason	Auto Recovery Time
Table is Empty				

Table 172: Port Auto Recovery Configuration Fields

Field	Description
Auto Recovery Components	<p>This field lists all the components that support the Auto Recovery feature. For each component, you can enable or disable Auto Recovery.</p> <p>An interface in the diagnostic disabled state for the configured components is recovered (link up) when the recovery interval expires. If the interface continues to encounter errors (from any listed components), it may be placed back in the diagnostic disabled state, and the interface will be disabled (link down). Interfaces in the diagnostic disabled state may also be manually recovered by enabling them from the Port Summary page.</p> <p>Auto Recovery is available for the following components:</p> <ul style="list-style-type: none"> <li>• ARP Inspection</li> <li>• Authentication Manager</li> <li>• BPDU Guard</li> <li>• BPDU Rate Limit</li> <li>• Broadcast Storm Control</li> <li>• CoA Disable Host Port</li> <li>• Denial Of Service</li> <li>• DHCP Rate Limit</li> <li>• Keepalive</li> <li>• Link Flap</li> <li>• MAC Locking</li> <li>• Multicast Storm Control</li> <li>• UDLD</li> <li>• Unicast Storm Control</li> </ul>
Recovery Time	The auto recovery time interval. The auto recovery time interval is common for all components. The default value of the timer is 300 seconds and the range is from 30 to 86400.
D-Disabled Interface Status	This table displays the list of interfaces that are error disabled.
Interface	The interface which is error disabled.
Admin Mode	The administrative mode of the interface.
Port Status	Indicates whether the link is up or down. The link is the physical connection between the port or trunk and the interface on another device.

Table 172: Port Auto Recovery Configuration Fields (Continued)

Field	Description
Error Disable Reason	<p>If the device detects an error condition for an interface, then the device puts the interface in error disabled state by placing the interface in diagnostic disabled state. The interface can go into error disable state for one of the following reasons:</p> <ul style="list-style-type: none"> <li>• ARP Inspection</li> <li>• Authentication Manager</li> <li>• BPDU Guard</li> <li>• BPDU Storm</li> <li>• Broadcast Storm</li> <li>• CoA Disable Host Port</li> <li>• Denial Of Service</li> <li>• DHCP Rate Limit</li> <li>• Keepalive</li> <li>• Link Flap</li> <li>• MAC Locking</li> <li>• Multicast Storm</li> <li>• UDLD</li> <li>• Unicast Storm</li> </ul>
Auto Recovery Time Left	<p>When Auto Recovery is enabled and the interface is placed in diagnostic disabled state, then a recovery timer (in seconds) starts for that interface. Once this timer expires, the device checks if the interface is in diagnostic disabled state. If yes, then the device enables the diagnostic disabled interface.</p>

Click Submit to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.

## 5.12 Creating MAC Filters

Static MAC filtering allows you to associate a MAC address with a VLAN and set of source ports and destination ports. (The availability of source and destination port filters is subject to platform restrictions). Any packet with a static MAC address in a specific VLAN is admitted only if the ingress port is included in the set of source ports; otherwise the packet is dropped. If admitted, the packet is forwarded to all the ports in the destination list.

### 5.12.1 MAC Filter Configuration

Use the MAC Filter Configuration page to associate a MAC address with a VLAN and one or more source and/or destination ports

To access the MAC Filter Configuration page, click Switching > Filters > MAC Filters in the navigation menu.

Figure 187: MAC Filter Configuration

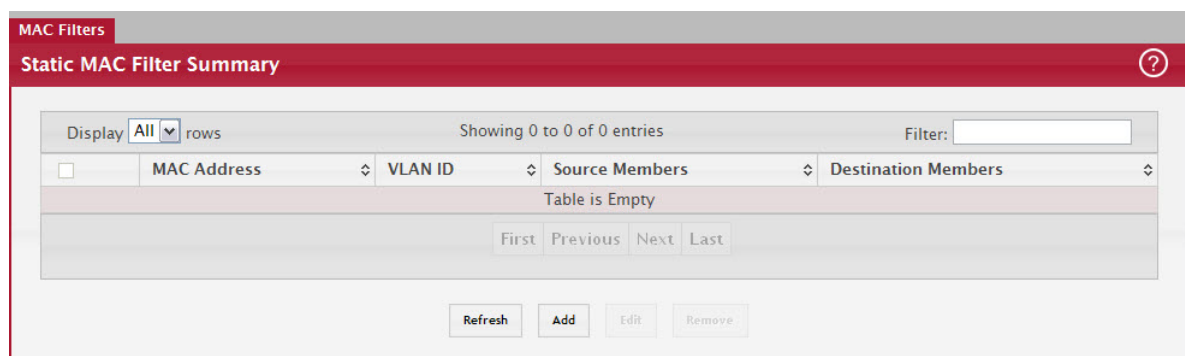


Table 173: MAC Filter Configuration Fields

Field	Description
MAC Address	The MAC address of the filter. The destination MAC address of an Ethernet frame must match this value to be considered for the filter. When adding or editing a filter, note that you cannot configure the following MAC addresses in this field: <ul style="list-style-type: none"> <li>• 00:00:00:00:00:00</li> <li>• 01:80:C2:00:00:00 to 01:80:C2:00:00:0F</li> <li>• 01:80:C2:00:00:20 to 01:80:C2:00:00:21</li> <li>• FF:FF:FF:FF:FF:FF</li> </ul>
VLAN ID	The VLAN ID associated with the filter. The VLAN ID is used with the MAC address to fully identify the frames to filter.
Source Port Mask	The port(s) included in the inbound filter. If a frame with the MAC address and VLAN ID combination specified in the filter is received on a port in the Source Members list, it is forwarded to a port in the Destination Members list. If the frame that meets the filter criteria is received on a port that is not in the Source Members list, it is dropped. To add source ports to the filter, select one or more ports from the Available Port List field (CTRL + click to select multiple ports). Then, use the appropriate arrow icon to move the selected ports to the Source Members field.
Destination Port Mask	The port(s) included in the outbound filter. A frame with the MAC address and VLAN ID combination specified in the filter is transmitted only out of ports in the list. To add destination ports to the filter, select one or more ports from the Available Port List field (CTRL + click to select multiple ports). Then, use the appropriate arrow icon to add the selected ports to the Source Members field.

### 5.12.1.1 Adding MAC Filters

1. To add a MAC filter, click Add from the MAC Filter summary page.
2. Enter a valid MAC address and select a VLAN ID from the drop-down menu.  
The VLAN ID drop-down menu only lists VLANs currently configured on the system.
3. Select one or more ports to include in the filter. Use CTRL + click to select multiple ports.
4. Click Submit to apply the changes to the system.

### 5.12.1.2 Modifying MAC Filters

To change the port mask(s) for an existing filter, select the entry from the MAC Filter field, and click Edit. When you have completed the changes, click Submit.

To change the MAC address or VLAN associated with a filter, you must remove and re-create the filter.

### 5.12.1.3 Removing MAC Filters

To remove a filter, select it from the MAC Filter drop-down menu and click Remove.

## 5.13 Configuring Dynamic ARP Inspection

Dynamic ARP Inspection (DAI) is a security feature that rejects invalid and malicious ARP packets. DAI prevents a class of man-in-the-middle attacks, where an unfriendly station intercepts traffic for other stations by poisoning the ARP caches of its unsuspecting neighbors. The miscreant sends ARP requests or responses mapping another station's IP address to its own MAC address.

DAI relies on DHCP snooping. DHCP snooping listens to DHCP message exchanges and builds a binding database of valid {MAC address, IP address, VLAN, and interface} tuples.

When DAI is enabled, the switch drops ARP packets whose sender MAC address and sender IP address do not match an entry in the DHCP snooping bindings database. You can optionally configure additional ARP packet validation.



### 5.13.1 DAI Configuration

Use the DAI Configuration page to configure global DAI settings.

To display the DAI Configuration page, click Switching > Dynamic ARP Inspection > Global in the navigation menu.

Figure 188: Dynamic ARP Inspection Global Configuration

Table 174: Dynamic ARP Inspection Global Configuration

Field	Description
Validate Source MAC	When this option is selected, DAI verifies that the sender hardware address in the ARP packet equals the source MAC address in the Ethernet header. If the addresses do not match, the ARP packet is dropped.
Validate Destination MAC	When this option is selected, DAI verifies that the target hardware address in the ARP packet equals the destination MAC address in the Ethernet header. If the addresses do not match, the ARP packet is dropped. This check applies only to ARP responses because the target MAC address is unspecified in ARP requests.
Validate IP	When this option is selected, DAI drops ARP packets with an invalid IP address. The following IP addresses are considered invalid: <ul style="list-style-type: none"> <li>• 0.0.0.0</li> <li>• 255.255.255.255</li> <li>• All IP multicast addresses</li> <li>• All class E addresses (240.0.0.0/4)</li> <li>• Loopback addresses (in the range 127.0.0.0/8)</li> </ul>

Click Submit to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.

### 5.13.2 DAI VLAN Configuration

Use the DAI VLAN Configuration page to select the DAI-capable VLANs for which information is to be displayed or configured.

Use the buttons to perform the following tasks:

- To enable DAI on a VLAN and to configure the optional DAI settings, click Add.
- To change the DAI settings on VLAN, select the VLAN with the settings to update and click Edit.
- To disable DAI on one or more VLANs, select each entry to disable and click Remove. After confirming the action, the entries are removed from the table.

To display the DAI Configuration page, click Switching > Dynamic ARP Inspection > VLAN in the navigation menu.

Figure 189: Dynamic ARP Inspection VLAN Configuration

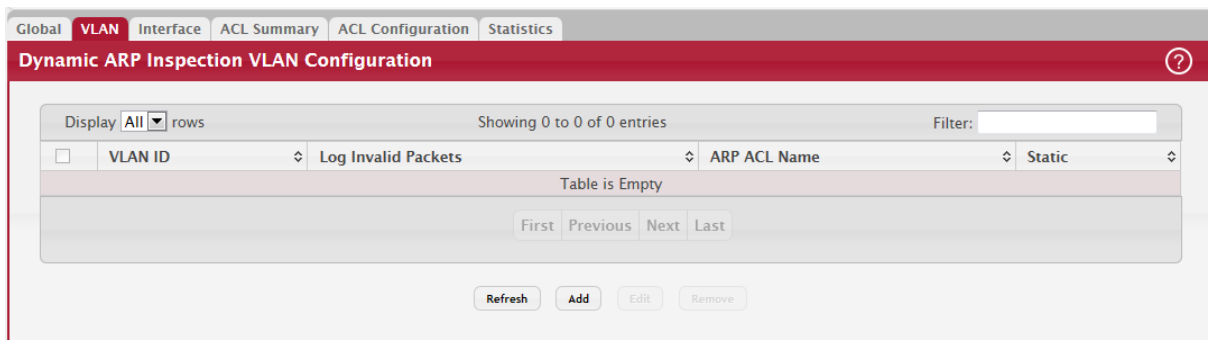


Table 175: Dynamic ARP Inspection VLAN Configuration

Field	Description
VLAN ID	Lists each VLAN that has been enabled for DAI. After you click Add, use the VLAN ID menu to select the VLAN on which to enable DAI. A VLAN does not need to exist on the system to be enabled for DAI.
Logging Invalid Packets	Indicates whether DAI logging is enabled on this VLAN. When logging is enabled, DAI generates a log message whenever an invalid ARP packet is discovered and dropped.
ARP ACL Name	The name of the of ARP access control list (ACL) that the VLAN uses as the filter for ARP packet validation. The ARP ACL must already exist on the system to associate it with a DAI-enabled VLAN. ARP ACLs include permit rules only.
Static	Determines whether to use the DHCP snooping database for ARP packet validation if the packet does not match any ARP ACL rules. The options are as follows: <ul style="list-style-type: none"> <li>• Enable – The ARP packet will be validated by the ARP ACL rules only. Packets that do not match any ARP ACL rules are dropped without consulting the DHCP snooping database.</li> <li>• Disable – The ARP packet needs further validation by using the entries in the DHCP Snooping database.</li> </ul>

- Click Submit to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click Refresh to refresh the page with the most current data from the switch.

### 5.13.3 DAI Interface Configuration

Use the DAI Interface Configuration page to select the DAI Interface for which information is to be displayed or configured.

To display the DAI Interface Configuration page, click Switching > Dynamic ARP Inspection > Interface Configuration in the navigation menu.

Figure 190: Dynamic ARP Inspection Interface Configuration

Interface	Trust State	Rate Limit	Burst Interval
1/0/1	Disabled	15	1
1/0/2	Disabled	15	1
1/0/3	Disabled	15	1
1/0/4	Disabled	15	1
1/0/5	Disabled	15	1
1/0/6	Disabled	15	1
1/0/7	Disabled	15	1
1/0/8	Disabled	15	1
1/0/9	Disabled	15	1
1/0/10	Disabled	15	1

Table 176: Dynamic ARP Inspection Interface Configuration

Field	Description
Interface	The interface associated with the rest of the data in the row. In the Edit Interface Configuration window, this field identifies the interface that is being configured.
Trust State	Indicates whether the DAI feature should check traffic on the interface for possible ARP packet violations. Trust state can be enabled or disabled after you select an interface and click Edit. This field has one of the following values: <ul style="list-style-type: none"> <li>Enabled – The interface is trusted. ARP packets arriving on this interface are forwarded without DAI validation.</li> <li>Disabled – The interface is not trusted. ARP packets arriving on this interface are subjected to ARP inspection.</li> </ul>
Rate Limit	The maximum rate for incoming ARP packets on the interface, in packets per second (pps). If the incoming rate exceeds the configured limit, the ARP packets are dropped. Rate limiting can be enabled or disabled after you select an interface and click Edit.
Burst Interval	The number of consecutive seconds the interface is monitored for incoming ARP packet rate limit violations.
Rate Limiting	Select this option to allow the interface to drop ARP packets if the rate at which they are received on the interface exceeds the configured Rate Limit for the Burst Interval duration. If this option is clear, rate limiting is disabled.

- Click Submit to apply the new configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click Refresh to refresh the page with the most current data from the switch.

#### 5.13.4 DAI ARP ACL Configuration

Use the DAI ARP ACL Configuration page to add or remove DAI ARP ACLs.

To display the DAI ARP ACL Configuration page, click Switching > Dynamic ARP Inspection > ACL Configuration in the navigation menu.

Figure 191: Dynamic ARP Inspection ACL Configuration

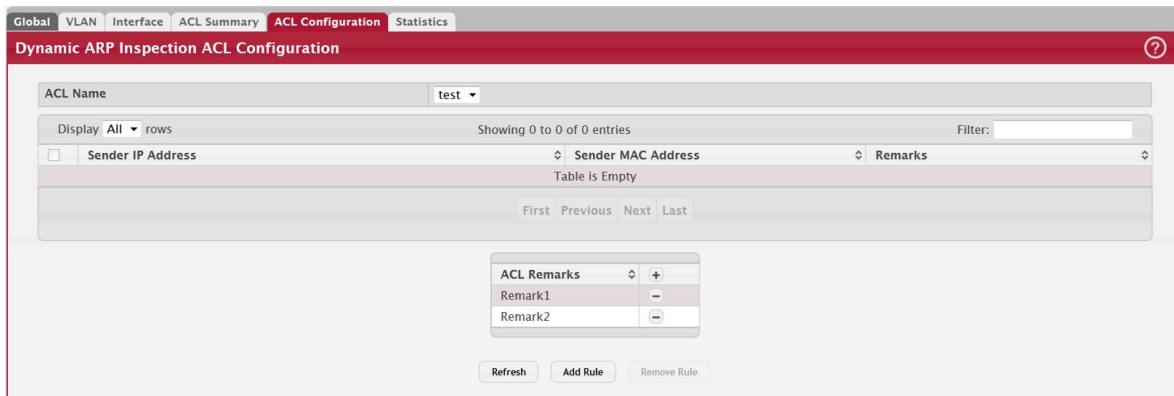


Table 177: Dynamic ARP Inspection ARP ACL Configuration

Field	Description
ACL Name	The menu contains the ARP ACL names that exist on the system.
Sender IP Address	The IP address of a system that is permitted to send ARP packets. The ARP packet must match on both the Sender IP Address and Sender MAC Address values in the rule to be considered valid.
Sender MAC Address	The MAC address of a system that is permitted to send ARP packets. The ARP packet must match on both the Sender IP Address and Sender MAC Address values in the rule to be considered valid.
Remarks	One or more remarks configured for the selected ACL and associated with the rule during rule creation.

Use the buttons to perform the following tasks:

- To create an ARP ACL and configure the first rule, click Add ACL.
- To add a new rule to an existing ACL, click Add Rule and select the name of the ACL to update from the ACL Name menu. Then, configure the rule.
- To remove one or more ARP ACLs, select each entry to delete and click Remove.
- Click Refresh to refresh the page with the most current data from the switch.

### 5.13.5 Add Access Control List

Use the Add Access Control List to create an ARP ACL and configure the first rule.

Figure 192: Add Access Control List

The Edit Access Control List page prevents the administrator from renaming a named IPv4, IPv6, or MAC ACL beginning with the case-insensitive names reserved for Dynamic ACLs, for example, IP-DACL-IN- and IPV6-DACL-IN-.

Figure 193: ACL Identifier Naming

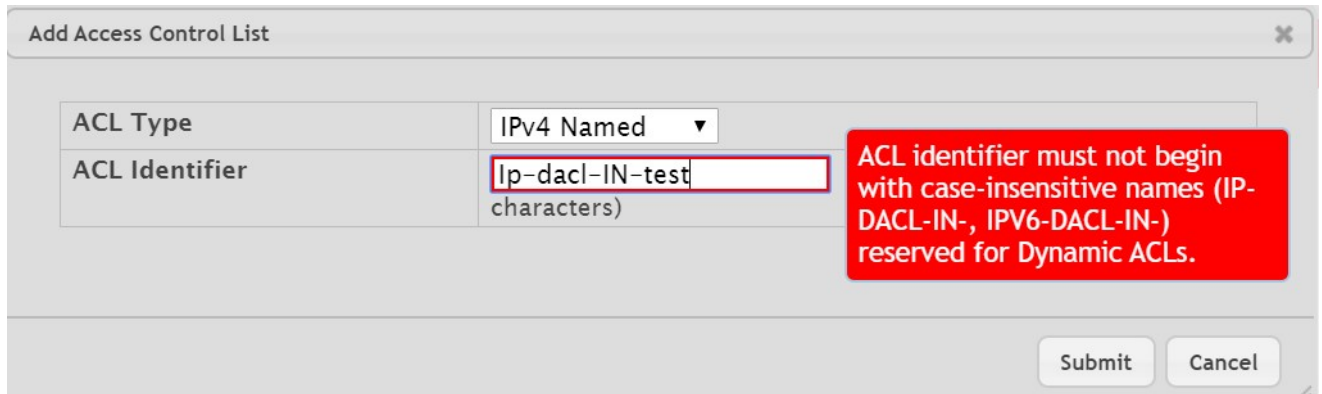


Table 178: Add Access Control List Fields

Field	Description
ACL Type	Possible values are IPv4 Named, IPv6 Named, and Extended MAC access lists.
ACL Identifier	255-character length ACL names are accepted for IPv4, IPv6, and MAC access-lists.

### 5.13.6 Add ACL Rule Configuration

Use the DAI ARP ACL Rule Configuration page to add or remove DAI ARP ACL Rules.

To display the DAI ARP ACL Rule Configuration page, click Add Rule from the Dynamic ARP Inspection ACL Configuration page.

Figure 194: Add ACL Rule



Table 179: Dynamic ARP Inspection ARP ACL Rule Configuration

Field	Description
Sender IP Address	To create a new rule for the selected ARP ACL, enter in this field the Sender IP Address match value for the ARP ACL.
Sender MAC Address	To create a new rule for the selected ARP ACL, enter in this field the Sender MAC Address match value for the ARP ACL.

Click Submit to add a new ARP ACL rule.

### 5.13.7 DAI ACL Summary

Use this page to configure ARP Access Control Lists (ACLs). An ARP ACL can contain one or more permit rules. Each rule contains the IP address and MAC address of a system allowed to send ARP packets. When an ARP ACL is associated with a DAI-enabled VLAN, and an ARP packet is received on an interface that is a member of that VLAN, DAI validates the address information in the ARP packet against the rules in the ACL. If the sender information in the ARP packet matches a rule in the ARP ACL, DAI considers the packet to be valid, and the packet is forwarded.

To display the DAI ARP ACL Configuration page, click Switching > Dynamic ARP Inspection > ACL Summary in the navigation menu.

Figure 195: Dynamic ARP Inspection ACL Summary

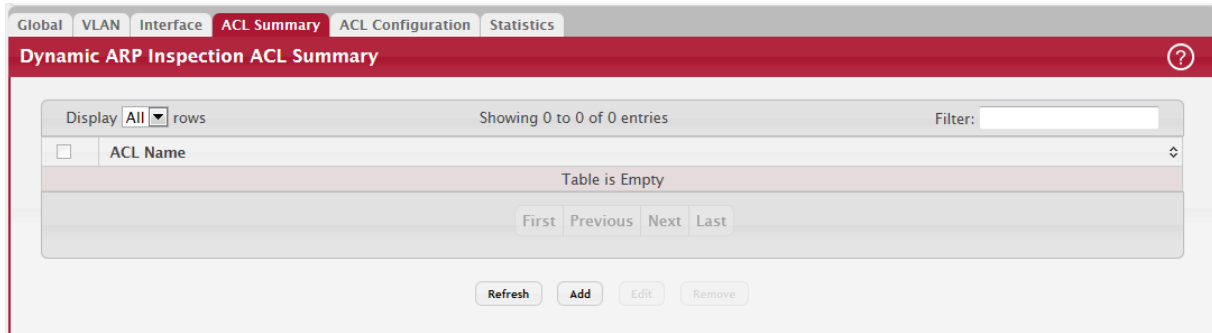


Table 180: Dynamic ARP Inspection ACL Summary Fields

Field	Description
ACL Name	The name of the ACL. Only the ACLs that appear in this column can be referenced by DNI-enabled VLANs.

Use the buttons to perform the following tasks:

- To add an ARP ACL, click Add and configure the ACL name.
- To configure rules for an ARP ACL, select the ACL to configure and click Edit. You are redirected to the Dynamic ARP Inspection ACL Configuration page for the selected ACL.
- To remove one or more ARP ACLs, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.
- Click Refresh to refresh the page with the most current data from the switch.

### 5.13.8 DAI Statistics

Use the DAI Statistics page to display the statistics per VLAN.

To display the DAI Statistics page, click Switching > Dynamic ARP Inspection > DAI Statistics in the navigation menu.

Figure 196: Dynamic ARP Inspection Statistics

Table 181: Dynamic ARP Inspection Statistics

Field	Description
VLAN ID	The DAI-enabled VLAN associated with the rest of the information in the row. When DAI is enabled on a VLAN, DAI is enabled on all interfaces that are members of that VLAN.
DHCP Drops	The number of ARP packets that have been dropped by DAI because no matching DHCP snooping binding entry was found in the DHCP snooping database.
ACL Drops	The number of ARP packets that have been dropped by DAI because the sender IP address and sender MAC address in the ARP packet did not match any rules in the ARP ACL associated with this VLAN. The static flag on this VLAN is enabled, which means ARP packets that fail to match an ARP ACL rule are dropped immediately and are not checked against the DHCP snooping database for further validation.
DHCP Permits	The number of ARP packets that were forwarded by DAI because a matching DHCP snooping binding entry was found in the DHCP snooping database.
ACL Permits	The number of ARP packets that were forwarded by DAI because the sender IP address and sender MAC address in the ARP packet matched a rule in the ARP ACL associated with this VLAN.
ACL Denials	The number of ARP packets that were dropped by DAI because the sender IP address and sender MAC address in the ARP packet matched a deny rule in the ARP ACL associated with this VLAN.
Bad Source MAC	The number of ARP packets that were dropped by DAI because the sender MAC address in ARP packet did not match the source MAC address in the Ethernet header.
Bad Dest MAC	The number of ARP packets that were dropped by DAI because the target MAC address in the ARP reply packet did not match the destination MAC address in the Ethernet header.
Invalid IP	The number of ARP packets that were dropped by DAI because the sender IP address in the ARP packet or target IP address in the ARP reply packet was invalid. The following IP addresses are considered invalid: <ul style="list-style-type: none"> <li>• 0.0.0.0</li> <li>• 255.255.255.255</li> <li>• All IP multicast addresses</li> <li>• All class E addresses (240.0.0.0/4)</li> <li>• Loopback addresses (in the range 127.0.0.0/8)</li> </ul>
Forwarded	The total number of valid ARP packets forwarded by DAI.
Dropped	The total number of invalid ARP packets dropped by DAI.

Click Refresh to refresh the page with the most current data from the switch.

## 5.14 GARP Configuration

Use this page to set the administrative mode for the features that use the Generic Attribute Registration Protocol (GARP), including GARP VLAN Registration Protocol (GVRP) and GARP Multicast Registration Protocol (GMRP). GARP is a general-purpose protocol that registers any network connectivity or membership-style information. GARP defines a set of switches interested in a given network attribute, such as VLAN ID or multicast address.

### 5.14.1 Switch Configuration

To access the GARP Switch Configuration page, click Switching > GARP > Switch in the navigation menu.

Figure 197: GARP Switch Configuration

Table 182: GARP Switch Configuration Fields

Field	Description
GVRP Mode	The administrative mode of GVRP on the system. When enabled, GVRP can help dynamically manage VLAN memberships on trunk ports.
GMRP Mode	The administrative mode of GMRP on the system. When enabled, GMRP can help control the flooding of multicast traffic by keeping track of group membership information. GMRP is similar to IGMP snooping in its purpose, but IGMP snooping is more widely used. GMRP must be running on both the host and the switch to function properly.

Click Refresh to refresh the page with the most current data from the switch.

### 5.14.2 Port Configuration

Use this page to set the per-interface administrative mode for GARP VLAN Registration Protocol (GVRP) and GARP Multicast Registration Protocol (GMRP). On this page, you can also set the GARP timers for each interface. GVRP and GMRP use the same set of GARP timers to specify the amount of time to wait before transmitting various GARP messages.

To access the GARP Port Configuration page, click Switching > GARP > Port in the navigation menu.



Figure 198: GARP Port Configuration

Interface	GVRP Mode	GMRP Mode	Join Timer (Centisecs)	Leave Timer (Centisecs)	Leave All Timer (Centisecs)
1/0/1	Disabled	Disabled	20	60	1000
1/0/2	Disabled	Disabled	20	60	1000
1/0/3	Disabled	Disabled	20	60	1000
1/0/4	Disabled	Disabled	20	60	1000
1/0/5	Disabled	Disabled	20	60	1000
1/0/6	Disabled	Disabled	20	60	1000
1/0/7	Disabled	Disabled	20	60	1000
1/0/8	Disabled	Disabled	20	60	1000
1/0/9	Disabled	Disabled	20	60	1000
1/0/10	Disabled	Disabled	20	60	1000

To change the GARP settings for one or more interfaces, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.

Table 183: GARP Port Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring one or more interfaces in the Edit GARP Port Configuration window, this field identifies the interfaces that are being configured.
GVRP Mode	The administrative mode of GVRP on the interface. When enabled, GVRP can help dynamically manage VLAN memberships on trunk ports. GVRP must also be enabled globally for the protocol to be active on the interface. When disabled, the protocol will not be active on the interface, and the GARP timers have no effect.
GMRP Mode	The administrative mode of GMRP on the interface. When enabled, GMRP can help control the flooding of multicast traffic by keeping track of group membership information. GMRP must also be enabled globally for the protocol to be active on the interface. When disabled, the protocol will not be active on the interface, and the GARP timers have no effect.
Join Timer (Centisecs)	The amount of time between the transmission of GARP PDUs registering (or re-registering) membership for a VLAN or multicast group.
Leave Timer (Centisecs)	The amount of time to wait after receiving an unregister request for a VLAN or multicast group before deleting the associated entry. This timer allows time for another station to assert registration for the same attribute to maintain uninterrupted service.
Leave All Timer (Centisecs)	The amount of time to wait before sending a LeaveAll PDU after the GARP application has been enabled on the interface or the last LeaveAll PDU was sent. A LeaveAll PDU indicates that all registrations will shortly be deregistered. Participants will need to rejoin to maintain registration.

Click Refresh to refresh the page with the most current data from the switch.

## 5.15 Configuring DHCP Snooping

DHCP snooping is a security feature that monitors DHCP messages between a DHCP client and DHCP servers to filter harmful DHCP messages and to build a bindings database of {MAC address, IP address, VLAN ID, port} tuples that are considered authorized. You can enable DHCP snooping globally and on specific VLANs, and configure ports within the VLAN to be trusted or untrusted. If a DHCP message arrives on an untrusted port, DHCP snooping filters messages that are not from authorized DHCP clients. DHCP server messages are forwarded only through trusted ports.

### 5.15.1 Global DHCP Snooping Configuration

Use this page to view and configure the global settings for DHCP Snooping.

To access the Global DHCP Snooping Configuration page, click Switching > DHCP Snooping > Base > Global in the navigation menu.

Figure 199: Global DHCP Snooping Configuration




Table 184: Global DHCP Snooping Configuration Fields

Field	Description
DHCP Snooping Mode	The administrative mode of DHCP snooping on the device.
MAC Address Validation	Enables or Disables the verification of the sender MAC address for DHCP snooping. When enabled, the device checks packets that are received on untrusted interface to verify that the MAC address and the DHCP client hardware address match. If the addresses do not match, the device drops the packet.

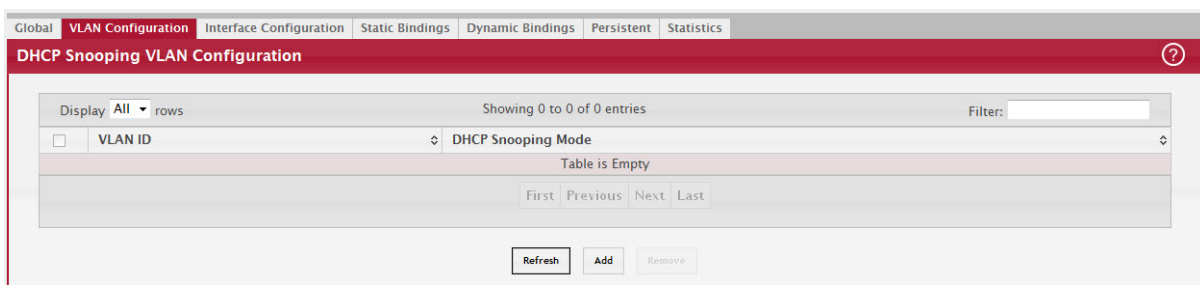
Click Refresh to refresh the page with the most current data from the switch.

### 5.15.2 DHCP Snooping VLAN Configuration

Use this page to view and configure the DHCP snooping settings on VLANs that exist on the device. DHCP snooping can be configured on switching VLANs and routing VLANs. For Layer 2 (non-routing) VLANs, DHCP snooping forwards valid DHCP client messages received on the VLANs. The message is forwarded on all trusted interfaces in the VLAN. When a DHCP packet is received on a routing VLAN, the DHCP snooping application applies its filtering rules and updates the bindings database. If a client message passes filtering rules, the message is placed into the software forwarding path, where it may be processed by the DHCP relay agent, the local DHCP server, or forwarded as an IP packet.

To access the DHCP Snooping VLAN Configuration page, click Switching > DHCP Snooping > Base > VLAN Configuration in the navigation menu.

Figure 200: DHCP Snooping VLAN Configuration



Use the buttons to perform the following tasks:

- To enable a VLAN for DHCP snooping, click Add and select the VLAN to administratively enable for DHCP snooping. To select multiple VLANs, CTRL + click each VLAN to select.
- To disable DHCP snooping on one or more VLANs, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

**Table 185: DHCP Snooping VLAN Configuration Fields**

Field	Description
VLAN ID	The VLAN ID that is enabled for DHCP snooping. In the Add DHCP Snooping VLAN Configuration window, this field lists the VLAN ID of all VLANs that exist on the device.
DHCP Snooping Mode	The current administration mode of DHCP snooping for the VLAN. Only VLANs that are enabled for DHCP snooping appear in the list.

Click Refresh to refresh the page with the most current data from the switch.

### 5.15.3 DHCP Snooping Interface Configuration

Use this page to view and configure the DHCP snooping settings for each interface. The DHCP snooping feature processes incoming DHCP messages. For DHCPRELEASE and DHCPDECLINE messages, the feature compares the receive interface and VLAN with the client's interface and VLAN in the binding database. If the interfaces do not match, the application logs the event (when logging of invalid packets is enabled) and drops the message. If MAC address validation is globally enabled, messages that pass the initial validation are checked to verify that the source MAC address and the DHCP client hardware address match. Where there is a mismatch, DHCP snooping logs the event (when logging of invalid packets is enabled) and drops the packet. To change the DHCP Snooping settings for one or more interfaces, select each entry to modify and click Edit. The same settings are applied to all selected interfaces.

To access the DHCP Snooping Interface Configuration page, click Switching > DHCP Snooping > Base > Interface Configuration in the navigation menu.

**Figure 201: DHCP Snooping Interface Configuration**

Interface	Trust State	Log Invalid Packets	Rate Limit (pps)	Burst Interval (Seconds)
0/1	Disabled	Disabled		
0/2	Disabled	Disabled		
0/3	Disabled	Disabled		
0/4	Disabled	Disabled		
0/5	Disabled	Disabled		
0/6	Disabled	Disabled		
0/7	Disabled	Disabled		
0/8	Disabled	Disabled		
1/1	Disabled	Disabled		
1/2	Disabled	Disabled		

Table 186: DHCP Snooping Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces, this field identifies each interface that is being configured.
Trust State	The trust state configured on the interface. The trust state is one of the following: <ul style="list-style-type: none"> <li>• Disabled – The interface is considered to be untrusted and could potentially be used to launch a network attack. DHCP server messages are checked against the bindings database. On untrusted ports, DHCP snooping enforces the following security rules: <ul style="list-style-type: none"> <li>- DHCP packets from a DHCP server (DHCP OFFER, DHCP ACK, DHCP NAK, DHCP RELEASE QUERY) are dropped.</li> <li>- DHCP RELEASE and DHCP DECLINE messages are dropped if the MAC address is in the snooping database but the binding's interface is other than the interface where the message was received.</li> <li>- DHCP packets are dropped when the source MAC address does not match the client hardware address if MAC Address Validation is globally enabled.</li> </ul> </li> <li>• Enabled – The interface is considered to be trusted and forwards DHCP server messages without validation.</li> </ul>
Log Invalid Packets	The administrative mode of invalid packet logging on the interface. When enabled, the DHCP snooping feature generates a log message when an invalid packet is received and dropped by the interface.
Rate Limit (pps)	The rate limit value for DHCP packets received on the interface. To prevent DHCP packets from being used as a DoS attack when DHCP snooping is enabled, the snooping application enforces a rate limit for DHCP packets received on untrusted interfaces. If the incoming rate of DHCP packets exceeds the value of this object during the amount of time specified for the burst interval, the port will be shutdown. You must administratively enable the port to allow it to resume traffic forwarding.
Burst Interval (Seconds)	The burst interval value for rate limiting on this interface. If the rate limit is unspecified, then burst interval has no meaning.

Click Refresh to refresh the page with the most current data from the switch.

#### 5.15.4 DHCP Snooping Static Bindings

Use this page to view, add, and remove static bindings in the DHCP snooping bindings database.

To access the DHCP Snooping Static Bindings page, click Switching > DHCP Snooping > Base > Static Bindings in the navigation menu.

Figure 202: DHCP Snooping Static Bindings

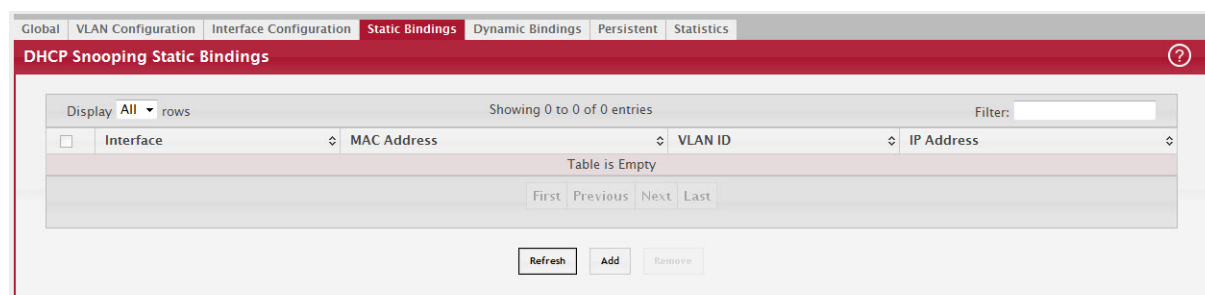


Table 187: DHCP Snooping Static Bindings Fields

Field	Description
Interface	The interface on which the DHCP client is authorized.
MAC Address	The MAC address associated with the DHCP client. This is the Key to the binding database.
VLAN ID	The ID of the VLAN the client is authorized to use.
IP Address	The IP address of the client.

Click Refresh to refresh the page with the most current data from the switch.

### 5.15.5 DHCP Snooping Dynamic Bindings

Use this page to view and clear dynamic bindings in the DHCP snooping bindings database. The DHCP snooping feature uses DHCP messages to build and maintain the bindings database. The bindings database includes data for clients only on untrusted ports. DHCP snooping creates a tentative binding from DHCP DISCOVER and REQUEST messages. Tentative bindings tie a client to an interface (the interface where the DHCP client message was received). Tentative bindings are completed when DHCP snooping learns the client's IP address from a DHCP ACK message on a trusted port. DHCP snooping removes bindings in response to DECLINE, RELEASE, and NACK messages. The DHCP snooping feature ignores the ACK messages as a reply to the DHCP Inform messages received on trusted ports.

To access the DHCP Snooping Dynamic Bindings page, click Switching > DHCP Snooping > Base > Dynamic Bindings in the navigation menu.

Figure 203: DHCP Snooping Dynamic Bindings

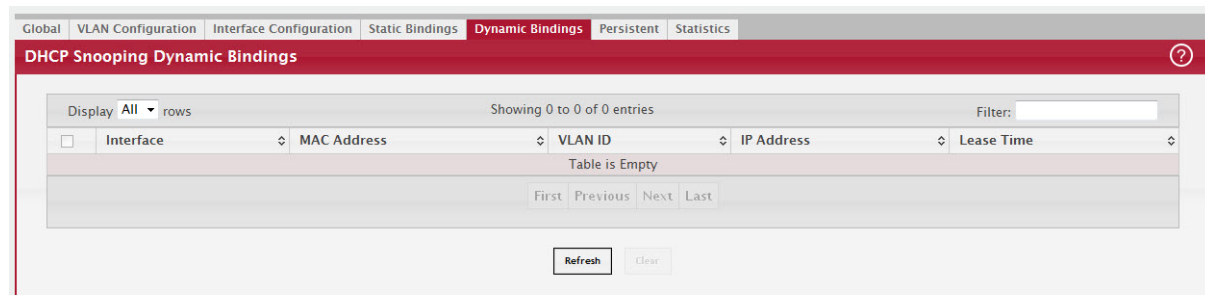


Table 188: DHCP Snooping Dynamic Bindings Fields

Field	Description
Interface	The interface on which the DHCP client message was received.
MAC Address	The MAC address associated with the DHCP client that sent the message. This is the Key to the binding database.
VLAN ID	The VLAN ID of the client interface.
IP Address	The IP address assigned to the client by the DHCP server.
Lease Time	The remaining IP address lease time for the client.
Clear (Button)	To remove one or more entries in the database, select each entry to delete and click Clear. You must confirm the action before the entry is deleted.

Click Refresh to refresh the page with the most current data from the switch.

### 5.15.6 DHCP Snooping Persistent Configuration

Use this page to configure the persistent location of the DHCP snooping bindings database. The bindings database can be stored locally on the device or on a remote system somewhere else in the network. The device must be able to reach the IP address of the remote system to send bindings to a remote database.

To access the DHCP Snooping Persistent Configuration page, click Switching > DHCP Snooping > Base > Persistent in the navigation menu.

Figure 204: DHCP Snooping Persistent Configuration

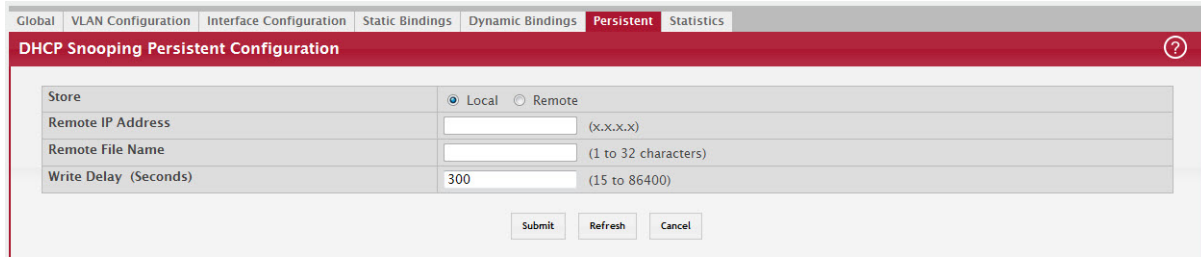


Table 189: DHCP Snooping Persistent Configuration Fields

Field	Description
Store	The location of the DHCP snooping bindings database, which is either locally on the device (Local) or on a remote system (Remote).
Remote IP Address	The IP address of the system on which the DHCP snooping bindings database will be stored. This field is available only if Remote is selected in the Store field.
Remote File Name	The file name of the DHCP snooping bindings database in which the bindings are stored. This field is available only if Remote is selected in the Store field.
Write Delay (Seconds)	The amount of time to wait between writing bindings information to persistent storage. This allows the device to collect as many entries as possible (new and removed) before writing them to the persistent file.

Click Refresh to refresh the page with the most current data from the switch.

### 5.15.7 DHCP Snooping Statistics

Use this page to configure the persistent location of the DHCP snooping bindings database. The bindings database can be stored locally on the device or on a remote system somewhere else in the network. The device must be able to reach the IP address of the remote system to send bindings to a remote database.

To access the DHCP Snooping Statistics page, click Switching > DHCP Snooping > Base > Statistics in the navigation menu.

Figure 205: DHCP Snooping Statistics

Interface	MAC Verify Failures	Client Ifc Mismatch	DHCP Server Msgs Received
0/5	0	0	0
0/6	0	0	0
0/7	0	0	0
0/8	0	0	0
1/1	0	0	0
1/2	0	0	0
1/3	0	0	0
1/4	0	0	0
1/5	0	0	0
1/6	0	0	0

Table 190: DHCP Snooping Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
MAC Verify Failures	The number of DHCP messages that were dropped because the source MAC address and client hardware address did not match. MAC address verification is performed only if it is globally enabled.
Client Ifc Mismatch	The number of packets that were dropped by DHCP snooping because the interface and VLAN on which the packet was received does not match the client's interface and VLAN information stored in the binding database.
DHCP Server Msgs Received	The number of DHCP server messages ((DHCP OFFER, DHCPACK, DHCPNAK, DHCPRELEASE, DHCPREQUEST)) that have been dropped on an untrusted port.
Clear Counters (Button)	To reset the statistics to zero for one or more interfaces, select each interface with the data to reset and click Clear Counters. You must confirm the action before the entry is deleted.

Click Refresh to refresh the page with the most current data from the switch.

### 5.15.8 DHCP L2 Relay Global Configuration

Use this page to control the administrative mode of DHCP Layer 2 Relay on the device. In Layer 2 switched networks, there may be one or more infrastructure devices (for example, a switch) between the client and the L3 Relay agent/DHCP server. In this instance, some of the client device information required by the L3 Relay agent may not be visible to it. When this happens, an L2 Relay agent can be used to add the information that the L3 Relay Agent and DHCP server need to perform their roles in IP address configuration and assignment.

To access the DHCP L2 Relay Global Configuration page, click Switching > DHCP Snooping > L2 Relay > Global in the navigation menu.

Figure 206: DHCP L2 Relay Global Configuration

Table 191: DHCP L2 Relay Global Configuration Fields

Field	Description
Admin Mode	The global mode of DHCP L2 relay on the device. When enabled, the device can act as a DHCP L2 relay agent. This functionality must also be enabled on each port you want this service to operate on.

Click Refresh to refresh the page with the most current data from the switch.

### 5.15.9 DHCP L2 Relay Interface Configuration

Use this page to enable L2 DHCP relay on individual ports. Note that L2 DHCP relay must also be enabled globally on the device. To change the DHCP L2 relay settings for one or more interfaces, select each entry to modify and click Edit. The same settings are applied to all selected interfaces.

To access the DHCP L2 Relay Interface Configuration page, click Switching > DHCP Snooping > L2 Relay > Interface Configuration in the navigation menu.

Figure 207: DHCP L2 Relay Interface Configuration

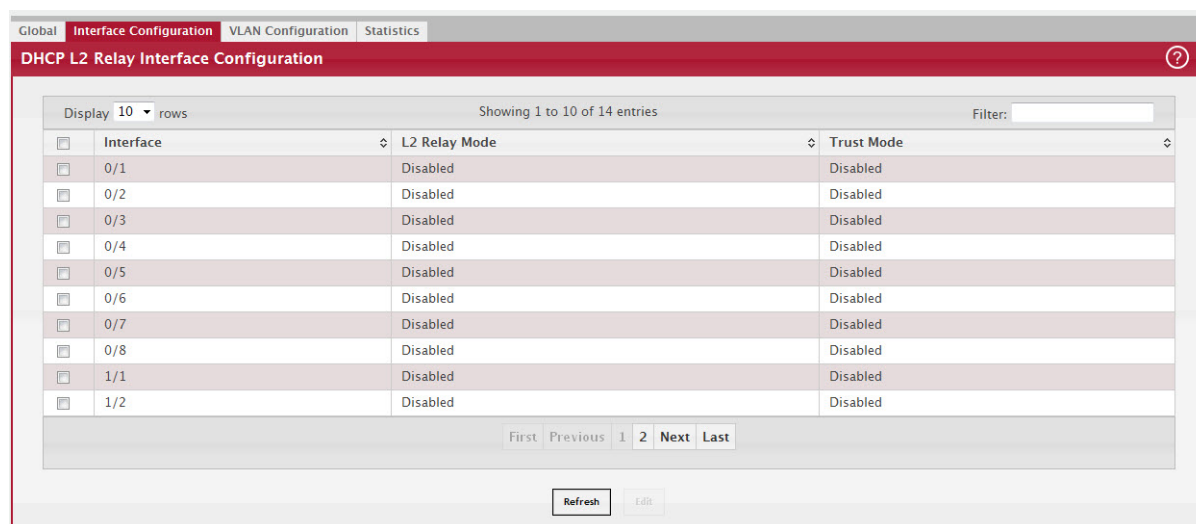


Table 192: DHCP L2 Relay Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces, this field identifies each interface that is being configured.
L2 Relay Mode	The administrative mode of L2 relay mode on the interface. When enabled, the interface can act as a DHCP relay agent and add information that the L3 relay agent and DHCP server need to perform their roles in IP address configuration and assignment.
Trust Mode	The L2 relay trust mode for the interface, which is one of the following: <ul style="list-style-type: none"> <li>Trusted – A trusted interface usually connects to other agents or servers participating in the DHCP interaction (e.g. other L2 or L3 relay agents or servers). An interface in this mode always expects to receive DHCP packets that include Option 82 information. If Option 82 information is not included, these packets are discarded.</li> <li>Untrusted – An untrusted interface is generally connected to clients. DHCP packets arriving on an untrusted interface are never expected to carry Option 82 and are discarded if they do.</li> </ul>



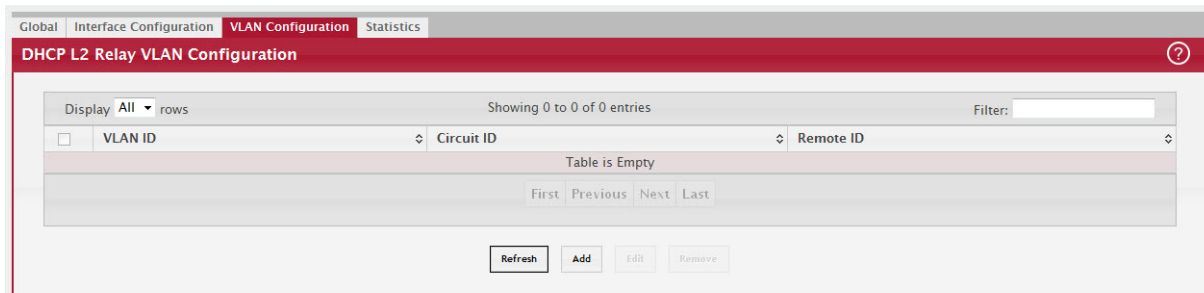
Click Refresh to refresh the page with the most current data from the switch.

### 5.15.10 DHCP L2 Relay VLAN Configuration

Use this page to control the DHCP L2 relay settings on a particular VLAN. The VLAN is identified by a service VLAN ID (S-VID), which a service provider uses to identify a customer's traffic while traversing the provider network to multiple remote sites. The device uses the VLAN membership of the switch port client (the customer VLAN ID, or C-VID) to perform a lookup on a corresponding S-VID.

To access the DHCP L2 Relay VLAN Configuration page, click Switching > DHCP Snooping > L2 Relay > VLAN Configuration in the navigation menu.

**Figure 208: DHCP L2 Relay VLAN Configuration**



Use the buttons to perform the following tasks:

- To enable a VLAN for DHCP L2 relay, click Add and select the VLAN from the available menu.
- To update the DHCP L2 relay settings for one or more VLANs, select each entry to update and click Edit. The same settings are applied to all selected VLANs.
- To disable one or more VLANs as DHCP L2 relay agents, select the appropriate VLANs and click Remove. You must confirm the action.

**Table 193: DHCP L2 Relay VLAN Configuration Fields**

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When configuring the settings for one or more VLANs, this field identifies each VLAN that is being configured.
Circuit ID	The administrative mode of the circuit ID. When enabled, if a client sends a DHCP request to the device and the client is in a VLAN that corresponds to the S-VID, the device adds the client's interface number to the Circuit ID sub-option of Option 82 in the DHCP request packet. This enables the device to reduce the broadcast domain to which the server replies are switched when the broadcast bit is set for DHCP packets. When this bit is set, the server is required to echo Option 82 in replies. Since the circuit-id field contains the client interface number, the L2 relay agent can forward the response to the requesting interface only, rather to all ports in the VLAN).
Remote ID	The DHCP remote identifier string. When a string is entered here, if a client sends a DHCP request to the device and the client is in a VLAN that corresponds to the S-VID, the device adds the string to the Remote-ID sub-option of Option 82 in the DHCP request packet. This sub-option can be used by the server for parameter assignment. The content of this option is vendor-specific.

Click Refresh to refresh the page with the most current data from the switch.

### 5.15.11 DHCP L2 Relay Interface Statistics

This page shows statistical information about the L2 DHCP Relay requests received on trusted and untrusted interfaces. An interface is untrusted when the DHCP L2 relay trust mode is disabled.

To access the DHCP L2 Relay Interface Statistics page, click Switching > DHCP Snooping > L2 Relay > Statistics in the navigation menu.

Figure 209: DHCP L2 Relay Interface Statistics

Interface	Untrusted Server Messages With Option-82	Untrusted Client Messages With Option-82	Trusted Server Messages With Option-82	Trusted Client Messages With Option-82
0/1	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
1/1	0	0	0	0
1/2	0	0	0	0

Table 194: DHCP L2 Relay Interface Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Untrusted Server Messages With Option-82	The number of messages received on an untrusted interface from a DHCP server that contained Option 82 data. These messages are dropped.
Untrusted Client Messages With Option-82	The number of messages received on an untrusted interface from a DHCP client that contained Option 82 data. These messages are dropped.
Trusted Server Messages With Option-82	The number of messages received on a trusted interface from a DHCP server that contained Option 82 data. These messages are forwarded.
Trusted Client Messages With Option-82	The number of messages received on a trusted interface from a DHCP client that contained Option 82 data. These messages are forwarded.
Clear (Button)	To reset the statistics to zero for one or more interfaces, select each interface with the data to reset and click Clear. You must confirm the action before the entry is deleted.

Click Refresh to refresh the page with the most current data from the switch.

### 5.15.12 DHCP Snooping IP Source Guard Interface Configuration

Use this page to configure IP Source Guard (IPSG) on each interface. IPSG is a security feature that filters IP packets based on source ID. This feature helps protect the network from attacks that use IP address spoofing to compromise or overwhelm the network. The source ID may be either the source IP address or a {source IP address, source MAC address} pair. The DHCP snooping bindings database, along with IPSG entries in the database, identify authorized source IDs. If you enable IPSG on a port where DHCP snooping is disabled or where DHCP snooping is enabled but the port is trusted, all IP traffic received on that port is dropped depending on the admin-configured IPSG entries. Additionally, IPSG interacts with port security, also known as port MAC locking, to enforce the source MAC address in received packets. Port security controls source MAC address learning in the Layer 2 forwarding database (MAC address table). When a frame is received with a previously unlearned source MAC address, port security queries the IPSG feature to determine whether the MAC address belongs to a valid binding. To change the IPSG configuration on one or more interfaces, select each entry to modify and click Edit. The same settings are applied to all selected interfaces.

To access the DHCP Snooping IP Source Guard Interface Configuration page, click Switching > DHCP Snooping > IP Source Guard > Interface Configuration in the navigation menu.

Figure 210: DHCP Snooping IP Source Guard Interface Configuration

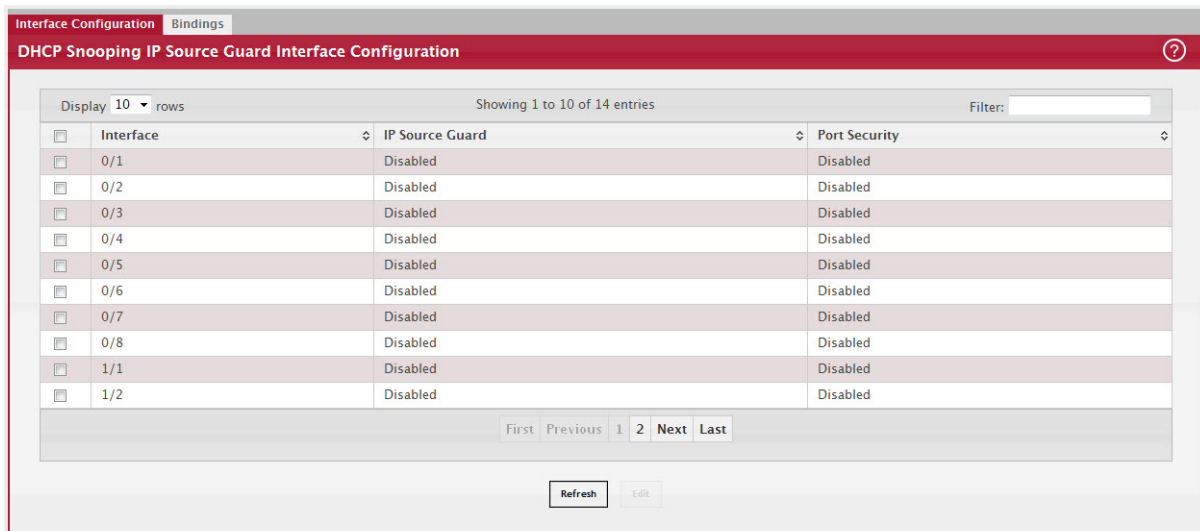


Table 195: DHCP Snooping IP Source Guard Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces in the Edit DHCP Snooping IP Source Guard Interface Configuration window, this field identifies each interface that is being configured.
IP Source Guard	The administrative mode of IPSG on the interface. When enabled, the source IP address is validated against the DHCP snooping bindings database, and DHCP packets will not be forwarded if the sender's IP address is not in the DHCP snooping bindings database.
Port Security	The administrative mode of IPSG Port Security on the interface. When IPSG Port Security is enabled, the packets will not be forwarded if the sender MAC address is not the in forwarding database table or the DHCP snooping bindings database. To enforce filtering based on MAC address, Port Security must be enabled globally and on the interface. IPSG Port Security cannot be enabled if IPSG is disabled.

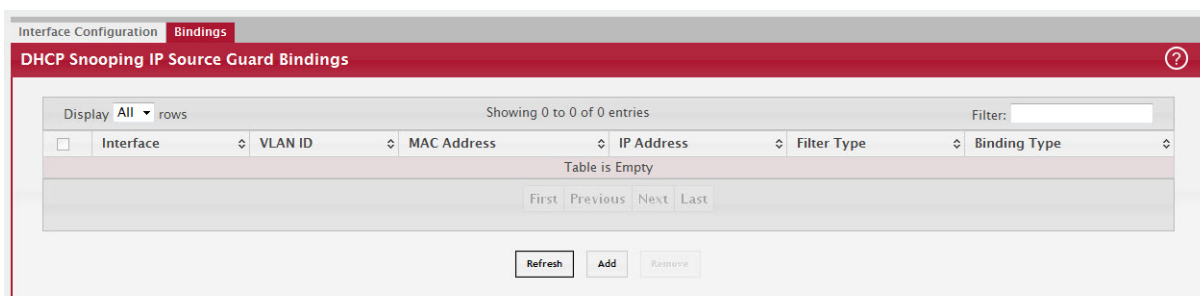
Click Refresh to refresh the page with the most current data from the switch.

### 5.15.13 DHCP Snooping IP Source Guard Bindings

Use this page to view IPSG bindings in the DHCP snooping IP Source Guard bindings database and to add or remove static bindings.

To access the DHCP Snooping IP Source Guard Bindings page, click Switching > DHCP Snooping > IP Source Guard > Bindings in the navigation menu.

Figure 211: DHCP Snooping IP Source Guard Bindings



Use the buttons to perform the following tasks:

- To add a static entry to the bindings database, click Add and specify the desired settings.
- To remove one or more entries, select each entry to delete and click Remove. You must confirm the action before the entry is deleted. Only static entries are selectable.

**Table 196: DHCP Snooping IP Source Guard Bindings Fields**

Field	Description
Interface	The interface on which the sender IP address is authorized.
VLAN ID	The authorized VLAN for the binding rule.
MAC Address	The authorized sender MAC address for the binding rule.
IP Address	The authorized source IP address for the binding rule.
Filter Type	The IPSG filter type.
Binding Type	The binding type, which is either dynamically learned or statically configured by an administrator.

Click Refresh to refresh the page with the most current data from the switch.

## 5.16 Configuring IGMP Snooping

Internet Group Management Protocol (IGMP) Snooping is a feature that allows a switch to forward multicast traffic intelligently on the switch. Multicast IP traffic is traffic that is destined to a host group. Host groups are identified by class D IP addresses, which range from 224.0.0.0 to 239.255.255.255. Based on the IGMP query and report messages, the switch forwards traffic only to the ports that request the multicast traffic. This prevents the switch from broadcasting the traffic to all ports and possibly affecting network performance.

A traditional Ethernet network may be separated into different network segments to prevent placing too many devices onto the same shared media. Bridges and switches connect these segments. When a packet with a broadcast or multicast destination address is received, the switch will forward a copy into each of the remaining network segments in accordance with the IEEE MAC Bridge standard. Eventually, the packet is made accessible to all nodes connected to the network.

This approach works well for broadcast packets that are intended to be seen or processed by all connected nodes. In the case of multicast packets, however, this approach could lead to less efficient use of network bandwidth, particularly when the packet is intended for only a small number of nodes. Packets will be flooded into network segments where no node has any interest in receiving the packet. While nodes will rarely incur any processing overhead to filter packets addressed to un-requested group addresses, they are unable to transmit new packets onto the shared media for the period of time that the multicast packet is flooded. The problem of wasting bandwidth is even worse when the LAN segment is not shared, for example in Full Duplex links.

Allowing switches to snoop IGMP packets is a creative effort to solve this problem. The switch uses the information in the IGMP packets as they are being forwarded throughout the network to determine which segments should receive packets directed to the group address.

### 5.16.1 Global Configuration and Status

Use the IGMP Snooping Global Configuration and Status page to enable IGMP snooping on the switch and view information about the current IGMP configuration.

To access the IGMP Snooping Configuration and Status page, click Switching > IGMP Snooping > Configuration in the navigation menu.

Figure 212: IGMP Snooping Global Configuration and Status




Table 197: IGMP Snooping Global Configuration and Status Fields

Field	Description
Admin Mode	Select the administrative mode for IGMP Snooping for the switch from the pull-down menu. The default is disable.
Multicast Control Frame Count	Shows the number of multicast control frames that have been processed by the CPU.
Interfaces Enabled for IGMP Snooping	Lists the interfaces currently enabled for IGMP Snooping. To enable interfaces for IGMP snooping, see <a href="#">Section 5.16.2: "Interface Configuration"</a> .
Data Frames Forwarded by the CPU	Shows the number of data frames forwarded by the CPU.

Select Enable or Disable the Admin Mode field and click Submit to turn the feature on or off. Perform a save if you want the changes to remain in effect over a power cycle.

### 5.16.2 Interface Configuration

Use the IGMP Snooping Interface Configuration page to configure IGMP snooping settings on specific interfaces.

To access the IGMP Snooping Interface Configuration page, click Switching > IGMP Snooping > Interface Configuration in the navigation menu.

Figure 213: IGMP Snooping Interface Configuration

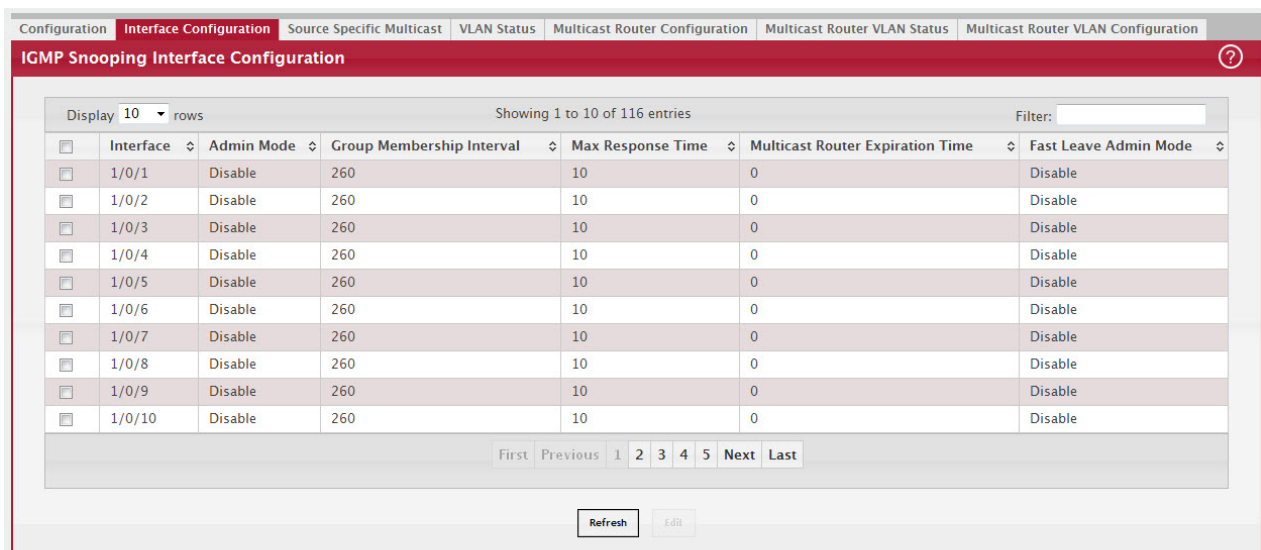


Table 198: IGMP Snooping Interface Configuration Fields

Field	Description
Interface	Select the physical or LAG interfaces to configure.
Admin Mode	Select the interface mode for the selected interface for IGMP Snooping for the switch from the pull-down menu. The default is disable.
Group Membership Interval	Specify the amount of time you want the switch to wait for a report for a particular group on a particular interface before it deletes that interface from the group. The valid range is from (2 to 3600) seconds. The default is 260 seconds.
Max Response Time	Specify the amount of time you want the switch to wait after sending a query on an interface because it did not receive a report for a particular group on that interface. Enter a value greater or equal to 1 and less than the Group Membership Interval in seconds. The default is 10 seconds. The configured value must be less than the Group Membership Interval.
Multicast Router Present Expiration Time	Specify the amount of time you want the switch to wait to receive a query on an interface before removing it from the list of interfaces with multicast routers attached. Enter a value between 0 and 3600 seconds. The default is 0 seconds. A value of zero indicates an infinite timeout; i.e., no expiration.
Fast Leave Admin Mode	Select the Fast Leave mode for the a particular interface from the pull-down menu. The default is Disable.

If you make any changes on the page, click Submit to apply the new settings to the switch.

### 5.16.3 Source Specific Multicast

This page displays information about multicast groups discovered by snooping IGMPv3 reports.

To access the Source Specific Multicast page, click Switching > IGMP Snooping > Source Specific Multicast in the navigation menu.

Figure 214: IGMP Snooping Source Specific Multicast

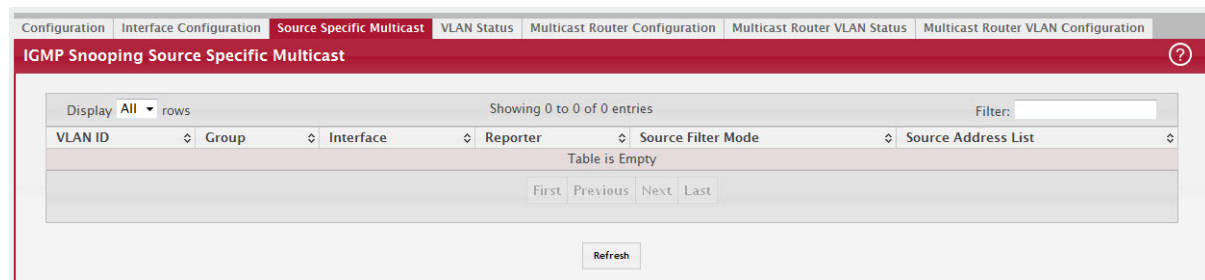


Table 199: IGMP Snooping Source Specific Multicast Fields

Field	Description
VLAN ID	VLAN on which the IGMP v3 report is received.
Group	The IPv4 multicast group address.
Interface	The interface on which the IGMP v3 report is received.
Reporter	The IPv4 address of the host that sent the IGMPv3 report.
Source Filter Mode	The source filter mode (Include/Exclude) for the specified group.
Source Address List	List of source IP addresses for which source filtering is requested.

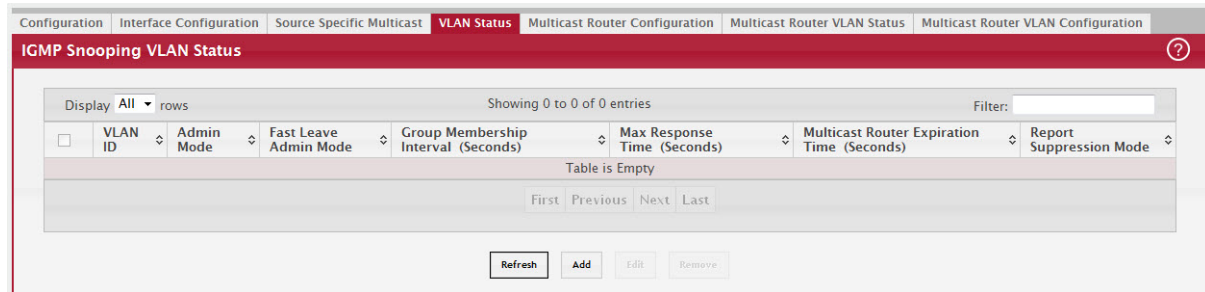
Click Refresh to refresh the page with the most current data from the switch.

### 5.16.4 VLAN Status

Use this page to enable or disable IGMP snooping on system VLANs and to view and configure per-VLAN IGMP snooping settings. Only VLANs that are enabled for IGMP snooping appear in the table.

To access the VLAN Status page, click Switching > IGMP Snooping > VLAN Status in the navigation menu.

Figure 215: IGMP Snooping VLAN Status



Use the buttons to perform the following tasks:

- To enable IGMP snooping on a VLAN, click Add and configure the settings in the available fields.
- To change the IGMP snooping settings for an IGMP-snooping enabled VLAN, select the entry with the settings to change and click Edit.
- To disable IGMP snooping on one or more VLANs, select each VLAN to modify and click Remove. You must confirm the action before IGMP snooping is disabled on the selected VLANs. When IGMP snooping is disabled, the VLAN entry is removed from the table, but the VLAN itself still exists on the system.

Table 200: IGMP Snooping VLAN Status Fields

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When enabling IGMP snooping on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for IGMP snooping appear in the menu. When modifying IGMP snooping settings, this field identifies the VLAN that is being configured.
Admin Mode	The administrative mode of IGMP snooping on the VLAN. IGMP snooping must be enabled globally and on an VLAN for the VLAN to be able to snoop IGMP packets to determine which network segments should receive multicast packets directed to the group address.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the VLAN. If Fast Leave is enabled, the VLAN can be immediately removed from the Layer 2 forwarding table entry upon receiving an IGMP leave message for a multicast group without first sending out MAC-based general queries.
Group Membership Interval (Seconds)	The number of seconds the VLAN should wait for a report for a particular group on the VLAN before the IGMP snooping feature deletes the VLAN from the group.
Max Response Time (Seconds)	The number of seconds the VLAN should wait after sending a query if does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.

Table 200: IGMP Snooping VLAN Status Fields (Continued)

Field	Description
Multicast Router Expiration Time (Seconds)	The number of seconds the VLAN should wait to receive a query before it is removed from the list of VLANs with multicast routers attached.
Report Suppression Mode	The IGMPv1 and IGMPv2 report suppression mode. The device uses IGMP report suppression to limit the membership report traffic sent to multicast-capable routers. When this mode is enabled, the device does not send duplicate reports to the multicast router. Note that this mode is supported only when the multicast query has IGMPv1 and IGMPv2 reports. This feature is not supported when the query includes IGMPv3 reports. The options are as follows: <ul style="list-style-type: none"> <li>• Enabled – Only the first IGMP report from all hosts for a group IGMP report is forwarded to the multicast routers.</li> <li>• Disabled – The device forwards all IGMP reports from all hosts in a multicast group to the multicast routers.</li> </ul>

Click Refresh to refresh the page with the most current data from the switch.

### 5.16.5 Multicast Router Configuration

If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure a switch port as a multicast router interface. Use the Multicast Snooping Multicast Router Configuration page to manually configure an interface as a static multicast router interface.

To access the IGMP Snooping Multicast Router Configuration page, click Switching > IGMP Snooping > Multicast Router Configuration in the navigation menu.

Figure 216: Multicast Router Configuration

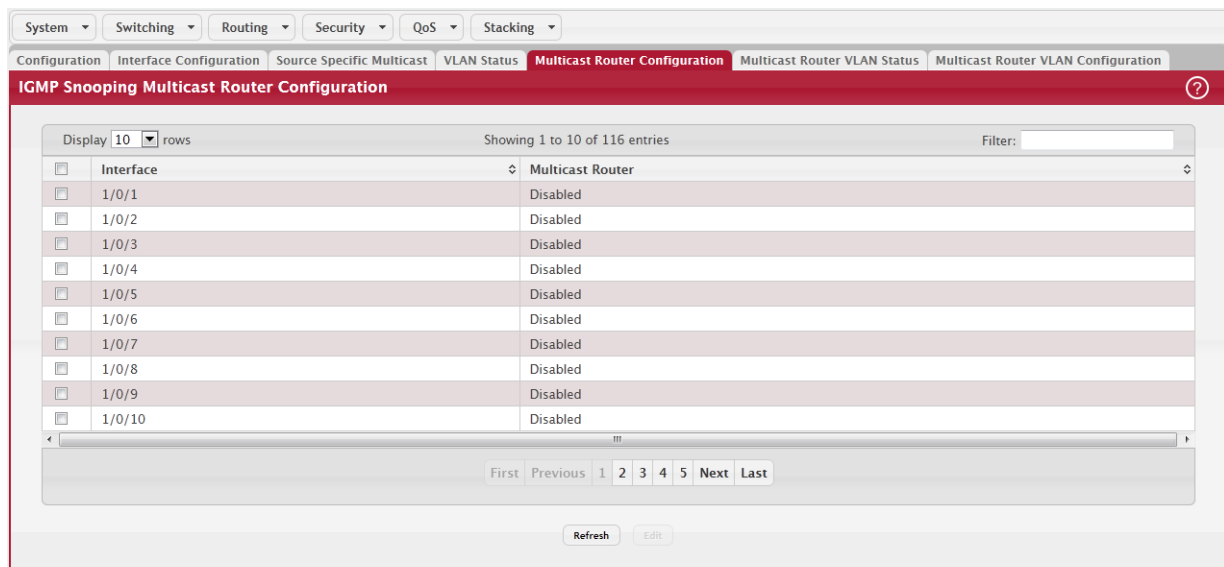


Table 201: Multicast Router Configuration Fields

Field	Description
Interface	Select the physical or LAG interface to display.
Multicast Router	Set the multicast router status: <ul style="list-style-type: none"> <li>• Enabled: The port is a multicast router interface.</li> <li>• Disabled: The port does not have a multicast router configured.</li> </ul>



If you enable or disable multicast router configuration on an interface, click Submit to apply the new settings to the switch.

### 5.16.6 Multicast Router VLAN Status

If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure one or more VLANs on each interface to act as a multicast router interface, which is an interface that faces a multicast router or IGMP querier and receives multicast traffic.

To access the Multicast Router VLAN Status page, click Switching > IGMP Snooping > Multicast Router VLAN Status in the navigation menu.

Figure 217: IGMP Snooping Multicast Router VLAN Status

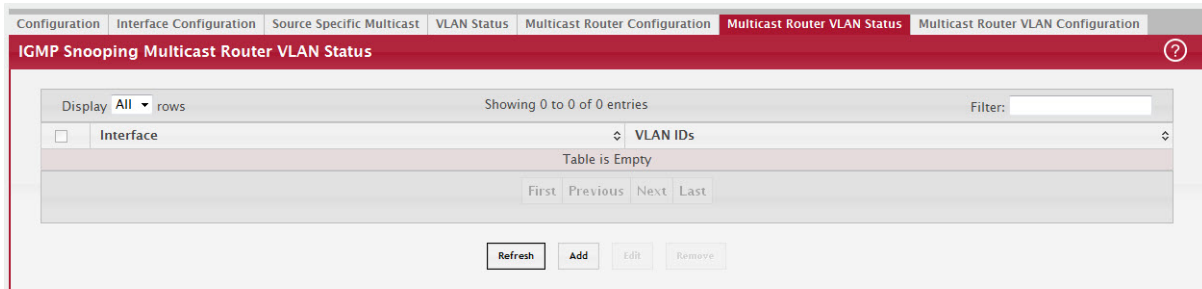


Table 202: IGMP Snooping Multicast Router VLAN Status Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that are configured with multicast router VLANs appear in the table.
VLAN IDs	The ID of the VLAN configured as enabled for multicast routing on the associated interface.

Use the buttons as follows:

- Click Refresh to refresh the page with the most current data from the switch.
- To disable all VLANs as multicast router interfaces for one or more physical ports or LAGs, select each entry to modify and click Remove.
- To enable or disable specific VLANs as multicast router interfaces for a physical port or LAG, use the Add and Edit buttons. A multicast router interface faces a multicast router or IGMP querier and receives multicast traffic.

### 5.16.7 Multicast Router VLAN Configuration

Use this page to enable or disable specific VLANs as multicast router interfaces for a physical port or LAG. A multicast router interface faces a multicast router or IGMP querier and receives multicast traffic.

To access the IGMP Snooping Multicast Router VLAN Configuration page, click Switching > IGMP Snooping > Multicast Router VLAN Configuration in the navigation menu.

Figure 218: IGMP Snooping Multicast Router VLAN Configuration

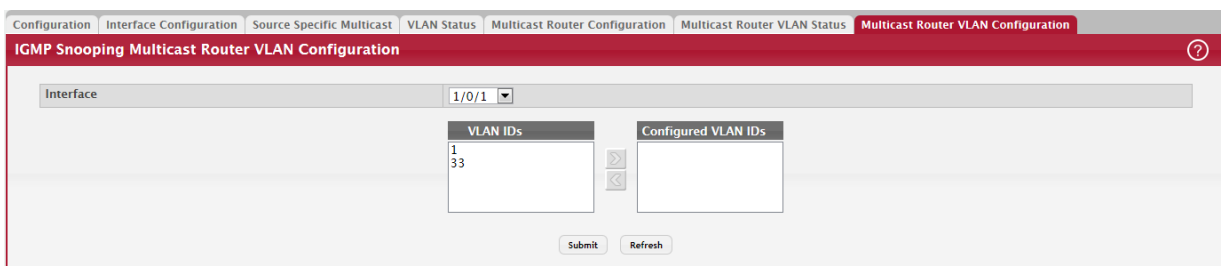


Table 203: IGMP Snooping Multicast Router VLAN Configuration Fields

Field	Description
Interface	Select the port or LAG on which to enable or disable a VLAN multicast routing interface.
VLAN IDs	The VLANs configured on the system that are not currently enabled as multicast router interfaces on the selected port or LAG. To enable a VLAN as a multicast router interface, click the VLAN ID to select it (or CTRL + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the Configured VLAN IDs window.
Configured VLAN IDs	The VLANs that are enabled as multicast router interfaces on the selected port or LAG. To disable a VLAN as a multicast router interface, click the VLAN ID to select it (or CTRL + click to select multiple VLAN IDs). Then, click the appropriate arrow to move the selected VLAN or VLANs to the VLAN IDs window.

Click Refresh to refresh the page with the most current data from the switch.

## 5.17 Configuring IGMP Snooping Querier

Use this page to configure the global IGMP snooping querier settings on the device. IGMP snooping requires that one central switch or router periodically query all end-devices on the network to announce their multicast memberships. This central device is the IGMP querier. When Layer 3 IP multicast routing protocols are enabled in a network with IP multicast routing, the IP multicast router acts as the IGMP querier. However, if the IP-multicast traffic in a VLAN needs to be Layer 2 switched only, an IP-multicast router is not required. The IGMP snooping querier can perform the IGMP snooping functions on the VLAN.

### 5.17.1 Configuration

To access the IGMP Snooping Querier Configuration page, click Switching > IGMP Snooping Querier > Configuration in the navigation menu.

Figure 219: IGMP Snooping Querier Configuration

Table 204: IGMP Snooping Querier Configuration Fields

Field	Description
Admin Mode	The administrative mode for the IGMP snooping querier on the device. When enabled, the IGMP snooping querier sends out periodic IGMP queries that trigger IGMP report messages from the switches that want to receive IP multicast traffic. The IGMP snooping feature listens to these IGMP reports to establish appropriate forwarding.
IP Address	The snooping querier address to be used as source address in periodic IGMP queries. This address is used when no IP address is configured on the VLAN on which the query is being sent.
IGMP Version	The IGMP protocol version used in periodic IGMP queries.

Table 204: IGMP Snooping Querier Configuration Fields (Continued)

Field	Description
Query Interval (Seconds)	The amount of time the IGMP snooping querier on the device should wait between sending periodic IGMP queries.
Querier Expiry Interval (Seconds)	The amount of time the device remains in non-querier mode after it has discovered that there is a multicast querier in the network.

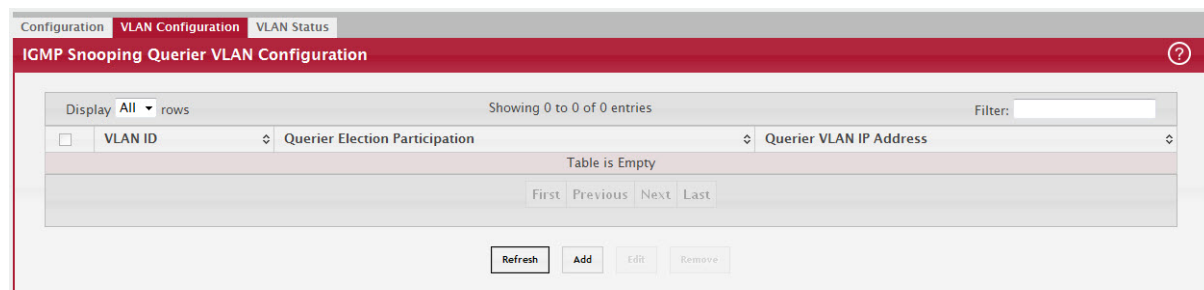
- If you make any changes to this page, click Submit to apply the changes to the system.
- Click Refresh to refresh the page with the most current data from the switch.

## 5.17.2 VLAN Configuration

Use this page to enable the IGMP snooping querier feature on one or more VLANs and to configure per-VLAN IGMP snooping querier settings. Only VLANs that have the IGMP snooping querier feature enabled appear in the table.

To access the IGMP Snooping Querier VLAN Configuration page, click Switching > IGMP Snooping Querier > VLAN Configuration in the navigation menu.

Figure 220: IGMP Snooping Querier VLAN Configuration



Use the buttons to perform the following tasks:

- To enable the IGMP snooping querier feature on a VLAN, click Add and specify the desired settings.
- To change the IGMP snooping querier settings for a VLAN, select the entry to modify and click Edit.
- To disable the IGMP snooping querier feature on one or more VLANs, select each entry to change and click Remove. You must confirm the action before the entry is deleted. Clicking this button does not remove the VLAN from the system.

Table 205: IGMP Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN on which the IGMP snooping querier is enabled. When enabling the IGMP snooping querier on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for the IGMP snooping querier appear in the menu. When modifying IGMP snooping querier settings, this field identifies the VLAN that is being configured.
Querier Election Participation	The participation mode for the IGMP snooping querier election process: <ul style="list-style-type: none"> <li>• Enabled – The IGMP snooping querier on this VLAN participates in the querier election process when it discovers the presence of another querier in the VLAN. If the snooping querier finds that the other querier source IP address is lower than its own address, it stops sending periodic queries. If the snooping querier wins the election (because it has the lowest IP address), then it continues sending periodic queries.</li> <li>• Disabled – When the IGMP snooping querier on this VLAN sees other queriers of the same version in the VLAN, the snooping querier moves to the non-querier state and stops sending periodic queries.</li> </ul>
Querier VLAN IP Address	The IGMP snooping querier address the VLAN uses as the source IP address in periodic IGMP queries sent on the VLAN. If this value is not configured, the VLAN uses the global IGMP snooping querier IP address.

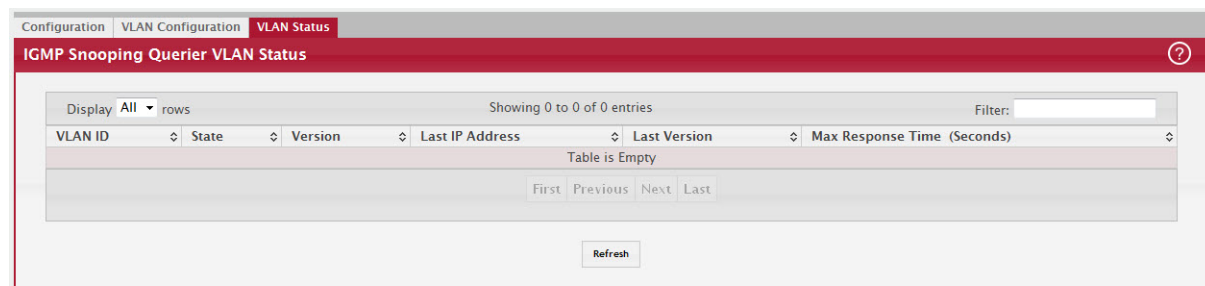
Click Refresh to refresh the page with the most current data from the switch.

### 5.17.3 VLAN Status

Use this page to view information about the IGMP snooping querier status for all VLANs that have the snooping querier enabled.

To access the IGMP Snooping Querier VLAN Status page, click Switching > IGMP Snooping Querier > VLAN Status in the navigation menu.

**Figure 221: IGMP Snooping Querier VLAN Status**



**Table 206: IGMP Snooping Querier VLAN Configuration Fields**

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row. The table includes only VLANs that have the snooping querier enabled.
State	The operational state of the IGMP snooping querier on the VLAN, which is one of the following: <ul style="list-style-type: none"> <li>Querier – The snooping switch is the querier in the VLAN. The snooping switch will send out periodic queries with a time interval equal to the configured querier query interval. If the snooping switch sees a better querier (numerically lower) in the VLAN, it moves to non-querier mode.</li> <li>Non-Querier – The snooping switch is in non-querier mode in the VLAN. If the querier expiry interval timer expires, the snooping switch moves into querier mode.</li> <li>Disabled – The snooping querier is not operational on the VLAN. The snooping querier moves to the disabled mode when IGMP snooping is not operational on the VLAN, when the querier address is not configured, or the network management address is not configured.</li> </ul>
Version	The operational IGMP protocol version of the querier.
Last IP Address	The IP address of the last querier from which a query was snooped on the VLAN.
Last Version	The IGMP protocol version of the last querier from which a query was snooped on the VLAN.
Max Response Time (Seconds)	The maximum response time to be used in the queries that are sent by the snooping querier.

Click Refresh to refresh the page with the most current data from the switch.

## 5.18 Configuring MLD Snooping

Use this page to enable Multicast Listener Discovery (MLD) snooping on the device and to view global status information. In IPv4, Layer 2 switches can use IGMP snooping to limit the flooding of multicast traffic by dynamically configuring Layer 2 interfaces so that multicast traffic is forwarded to only those interfaces associated with IP multicast address. In IPv6 networks, MLD snooping performs a similar function. With MLD snooping, IPv6 multicast data is selectively forwarded to a list of ports intended to receive the data (instead of being flooded to all of the ports in a VLAN). This list is constructed by snooping IPv6 multicast control packets.

## 5.18.1 Global Configuration and Status

To access the MLD Snooping Configuration and Status page, click Switching > MLD Snooping > Configuration in the navigation menu.

Figure 222: MLD Snooping Configuration and Status

Field	Value
MLD Snooping Admin Mode	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Multicast Control Frame Count	0
Interfaces Enabled for MLD Snooping	0/1, 0/2, 0/74
VLANs Enabled for MLD Snooping	10

Submit Refresh Cancel

Table 207: MLD Snooping Configuration and Status Fields

Field	Description
MLD Snooping Admin Mode	The administrative mode of MLD snooping on the device.
Multicast Control Frame Count	The number of multicast control frames that have been processed by the CPU.
Interfaces Enabled for MLD Snooping	One or more interfaces on which MLD snooping is administratively enabled. MLD snooping must be enabled globally and on an interface for the interface to be able to snoop MLD packets to determine which segments should receive multicast packets directed to the group address.
VLANs Enabled for MLD Snooping	One or more VLANs on which MLD snooping is administratively enabled.

Select *Enable* or *Disable* for the MLD Snooping Admin Mode field and click *Submit* to turn the feature on or off. Perform a save if you want the changes to remain in effect over a power cycle.

## 5.18.2 Interface Configuration

Use this page to configure MLD snooping settings on specific interfaces. To configure the settings for one or more interfaces, select each entry to modify and click *Edit*. The same MLD snooping settings are applied to all selected interfaces.

To access the MLD Snooping Interface Configuration page, click Switching > MLD Snooping > Interface Configuration in the navigation menu.

Figure 223: MLD Snooping Interface Configuration

Interface	Admin Mode	Group Membership Interval	Max Response Time	Multicast Router Expiration Time	Fast Leave Admin Mode
1/0/1	Disabled	260	10	0	Disabled
1/0/2	Disabled	260	10	0	Disabled
1/0/3	Disabled	260	10	0	Disabled
1/0/4	Disabled	260	10	0	Disabled
1/0/5	Disabled	260	10	0	Disabled
1/0/6	Disabled	260	10	0	Disabled
1/0/7	Disabled	260	10	0	Disabled
1/0/8	Disabled	260	10	0	Disabled
1/0/9	Disabled	260	10	0	Disabled
1/0/10	Disabled	260	10	0	Disabled

Table 208: MLD Snooping Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring MLD snooping settings, this field identifies each interface that is being configured.
Admin Mode	The administrative mode of MLD snooping on the interface. MLD snooping must be enabled globally and on an interface for the interface to be able to snoop MLD packets to determine which segments should receive multicast packets directed to the group address.
Group Membership Interval	The number of seconds the interface should wait for a report for a particular group on the interface before the MLD snooping feature deletes the interface from the group.
Max Response Time	The number of seconds the interface should wait after sending a query if it does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.
Multicast Router Present Expiration Time	The number of seconds the interface should wait to receive a query before it is removed from the list of interfaces with multicast routers attached.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the interface. If Fast Leave is enabled, the interface can be immediately removed from the Layer 2 forwarding table entry upon receiving an MLD leave message for a multicast group without first sending out MAC-based general queries.

If you make any changes on the page, click Submit to apply the new settings to the switch.

### 5.18.3 Source Specific Multicast

This page displays Source Specific Multicast (SSM) information learned by snooping MLDv2 reports. MLDv2 includes support for SSM, in which a receiver can request to receive multicast packets from one or more specific source address or from all addresses except one or more specified source addresses. If a host sends an MLDv2 report, the MLD snooping feature records the information and adds an entry to the table on this page.

To access the Source Specific Multicast page, click Switching > MLD Snooping > Source Specific Multicast in the navigation menu.

Figure 224: MLD Snooping Source Specific Multicast

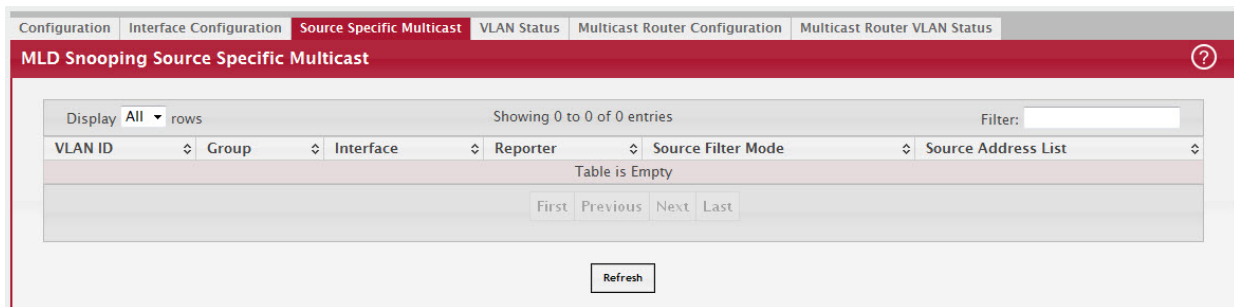


Table 209: MLD Snooping Source Specific Multicast Fields

Field	Description
VLAN ID	The VLAN on which the MLDv2 report is received.
Group	The IPv6 multicast group address of the multicast group the host belongs to.
Interface	The interface on which the MLD v2 report is received.
Reporter	The IPv6 address of the host that sent the MLDv2 report.
Source Filter Mode	The source filter mode for the specified group, which is one of the following: <ul style="list-style-type: none"> <li>• Include – The receiver has expressed interest in receiving multicast traffic for the multicast group from the source or sources in the Source Address List.</li> <li>• Exclude – The receiver has expressed interest in receiving multicast traffic for the multicast group from any source except the source or sources in the Source Address List.</li> </ul>
Source Address List	The source IPv6 address or addresses for which source filtering is requested.

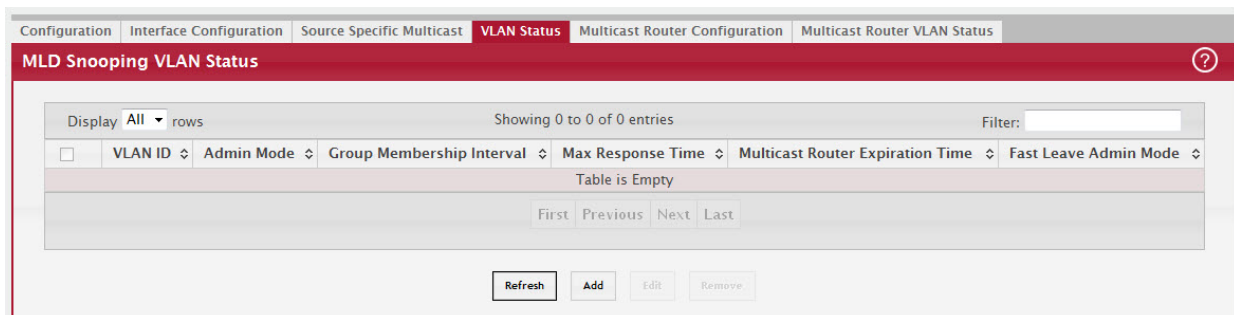
Click Refresh to refresh the page with the most current data from the switch.

### 5.18.4 VLAN Status

Use this page to enable or disable MLD snooping on system VLANs and to view and configure per-VLAN MLD snooping settings. Only VLANs that are enabled for MLD snooping appear in the table.

To access the VLAN Status page, click Switching > MLD Snooping > VLAN Status in the navigation menu.

Figure 225: MLD Snooping VLAN Status



Use the buttons to perform the following tasks:

- To enable MLD snooping on a VLAN, click Add and configure the settings in the available fields.
- To change the MLD snooping settings for an MLD-snooping enabled VLAN, select the entry with the settings to change and click Edit.
- To disable MLD snooping on one or more VLANs, select each VLAN to modify and click Remove. You must confirm the action before MLD snooping is disabled on the selected VLANs. When MLD snooping is disabled, the VLAN entry is removed from the table, but the VLAN itself still exists on the system.

Table 210: MLD Snooping VLAN Status Fields

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row. When enabling MLD snooping on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for MLD snooping appear in the menu. When modifying MLD snooping settings, this field identifies the VLAN that is being configured.
Admin Mode	The administrative mode of MLD snooping on the VLAN. MLD snooping must be enabled globally and on a VLAN for the VLAN to be able to snoop MLD packets and determine which network segments should receive multicast packets directed to the group address.
Group Membership Interval	The number of seconds the VLAN should wait for a report for a particular group on the VLAN before the MLD snooping feature deletes the VLAN from the group.
Max Response Time	The number of seconds the VLAN should wait after sending a query if does not receive a report for a particular group. The specified value should be less than the Group Membership Interval.
Multicast Router Expiration Time	The number of seconds the VLAN should wait to receive a query before it is removed from the list of VLANs with multicast routers attached.
Fast Leave Admin Mode	The administrative mode of Fast Leave on the VLAN. If Fast Leave is enabled, the VLAN can be immediately removed from the Layer 2 forwarding table entry upon receiving an MLD leave message for a multicast group without first sending out MAC-based general queries.

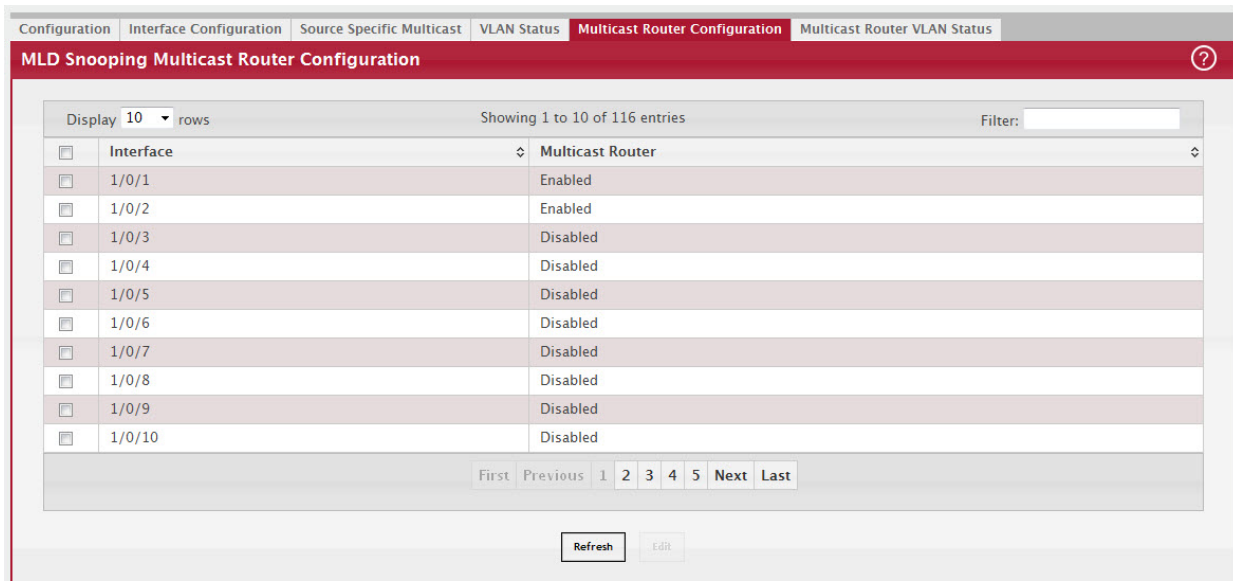
Click Refresh to refresh the page with the most current data from the switch.

### 5.18.5 Multicast Router Configuration

Use this page to manually configure an interface as a static MLD snooping multicast router interface. If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure an interface as a multicast router interface, which is an interface that faces a multicast router or MLD querier and receives multicast traffic.

To access the MLD Snooping Multicast Router Configuration page, click Switching > MLD Snooping > Multicast Router Configuration in the navigation menu.

Figure 226: Multicast Router Configuration





Use the buttons to perform the following tasks:

- To change the multicast router mode for one or more interfaces, select each entry to modify and click Edit.

**Table 211: Multicast Router Configuration Fields**

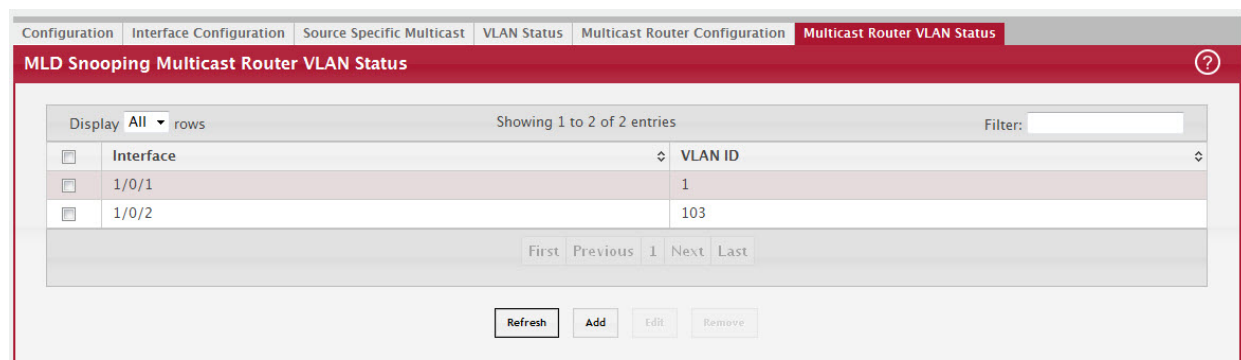
Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring the MLD snooping multicast router settings, this field identifies each interface that is being configured.
Multicast Router	Indicates whether the interface is enabled or disabled as a multicast router interface.

### 5.18.6 Multicast Router VLAN Status

Use this page to enable or disable specific VLANs as static multicast router interfaces for a physical port or LAG and to view the multicast router VLAN status for each interface. A multicast router interface faces a multicast router or MLD querier and receives multicast traffic. If a multicast router is attached to the switch, its existence can be learned dynamically. You can also statically configure one or more VLANs on each interface to act as multicast router interfaces.

To access the Multicast Router VLAN Status page, click Switching > MLD Snooping > Multicast Router VLAN Status in the navigation menu.

**Figure 227: MLD Snooping Multicast Router VLAN Status**



Use the buttons to perform the following tasks:

- To enable one or more VLANs as multicast router interfaces on a port or LAG, click Add and configure the settings in the available fields.
- To change the VLANs that are enabled as multicast router interfaces for a port or LAG, select the entry with the settings to change and click Edit.
- To disable all VLAN multicast routing interfaces for a port or LAG, select each entry to modify and click Remove. You must confirm the action.

**Table 212: MLD Snooping Multicast Router VLAN Status Fields**

Field	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that are configured with multicast router VLANs appear in the table. When adding multicast router VLAN information for an interface, use the Interface menu to select the interface on which to enable one or more multicast router VLAN interfaces. When editing multicast router VLAN information, this field shows the interface that is being configured.
VLAN IDs	The ID of each VLAN configured as enabled as a multicast router interface on the associated interface. When changing the multicast routing VLAN interfaces that are associated with an interface, click the VLAN ID to select it (or CTRL + click to select multiple VLAN IDs).

Click Refresh to refresh the page with the most current data from the switch.

## 5.19 Configuring MLD Snooping Querier

Use this page to configure the global MLD snooping querier settings on the device. MLD snooping requires that one central switch or router periodically query all end-devices on the network to announce their multicast memberships. This central device is the MLD querier. When Layer 3 IP multicast routing protocols are enabled in a network with IP multicast routing, the IP multicast router acts as the MLD querier. However, if the IP-multicast traffic in a VLAN needs to be Layer 2 switched only, an IP-multicast router is not required. The MLD snooping querier can perform the MLD snooping functions on the VLAN.

### 5.19.1 Configuration

To access the MLD Snooping Querier Configuration page, click Switching > MLD Snooping Querier > Configuration in the navigation menu.

Figure 228: MLD Snooping Querier Configuration

Table 213: MLD Snooping Querier Configuration Fields

Field	Description
Admin Mode	The administrative mode for the MLD snooping querier on the device. When enabled, the MLD snooping querier sends out periodic MLD queries that trigger MLD report messages from the switches that want to receive IP multicast traffic. The MLD snooping feature listens to these MLD reports to establish appropriate forwarding.
IPv6 Address	The snooping querier unicast link-local IPv6 address to be used as the source address in periodic MLD queries. This address is used when no IPv6 address is configured on the VLAN on which the query is being sent.
MLD Version	The MLD protocol version used in periodic MLD queries.
Query Interval (Seconds)	The amount of time the MLD snooping querier should wait between sending periodic MLD queries.
Querier Expiry Interval (Seconds)	The amount of time the device remains in non-querier mode after it has discovered that there is a multicast querier in the network.

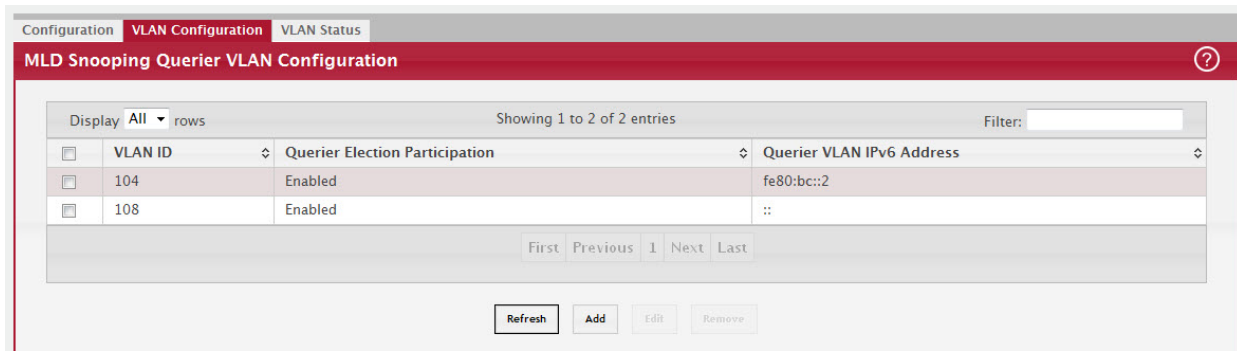
- If you make any changes to this page, click Submit to apply the changes to the system.
- Click Refresh to refresh the page with the most current data from the switch.

### 5.19.2 VLAN Configuration

Use this page to enable the MLD snooping querier feature on one or more VLANs and to configure per-VLAN MLD snooping querier settings. Only VLANs that have the MLD snooping querier feature enabled appear in the table.

To access the MLD Snooping Querier VLAN Configuration page, click Switching > MLD Snooping Querier > VLAN Configuration in the navigation menu.

Figure 229: MLD Snooping Querier VLAN Configuration



Use the buttons to perform the following tasks:

- To enable the MLD snooping querier feature on a VLAN, click Add and specify the desired settings.
- To change the MLD snooping querier settings for a VLAN, select the entry to modify and click Edit.
- To disable the MLD snooping querier feature on one or more VLANs, select each entry to change and click Remove. You must confirm the action before the entry is deleted. Clicking this button does not remove the VLAN from the system.

Table 214: MLD Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN on which the MLD snooping querier is enabled. When enabling the MLD snooping querier on a VLAN, use this menu to select the desired VLAN. Only VLANs that have been configured on the system and are not already enabled for the MLD snooping querier appear in the menu. When modifying MLD snooping querier settings, this field identifies the VLAN that is being configured.
Querier Election Participation	The participation mode for the MLD snooping querier election process: <ul style="list-style-type: none"> <li>• Enabled – The MLD snooping querier on this VLAN participates in the querier election process when it discovers the presence of another querier in the VLAN. If the snooping querier finds that the other querier source IPv6 address is lower than its own address, it stops sending periodic queries. If the snooping querier wins the election (because it has the lowest IPv6 address), then it continues sending periodic queries.</li> <li>• Disabled – When the MLD snooping querier on this VLAN sees other queriers of the same version in the VLAN, the snooping querier moves to the non-querier state and stops sending periodic queries.</li> </ul>
Querier VLAN IPv6 Address	The MLD snooping querier unicast link-local IPv6 address the VLAN uses as the source address in periodic MLD queries sent on the VLAN. If this value is not configured, the VLAN uses the global MLD snooping querier IPv6 address.

Click Refresh to refresh the page with the most current data from the switch.

### 5.19.3 VLAN Status

Use this page to view information about the MLD snooping querier status for all VLANs that have the snooping querier enabled.

To access the MLD Snooping Querier VLAN Status page, click Switching > MLD Snooping Querier > VLAN Status in the navigation menu.

Figure 230: MLD Snooping Querier VLAN Status

VLAN ID	State	Version	Last IPv6 Address	Last Version	Max Response Time (Seconds)
104	Disabled	1	::		
108	Disabled	1	::		

Table 215: MLD Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row. The table includes only VLANs that have the snooping querier enabled.
State	The operational state of the MLD Snooping Querier on a VLAN, which is one of the following: <ul style="list-style-type: none"> <li>Querier – The snooping switch is the querier in the VLAN. The snooping switch will send out periodic queries with a time interval equal to the configured querier query interval. If the snooping switch sees a better querier (numerically lower) in the VLAN, it moves to non-querier mode.</li> <li>Non-Querier – The snooping switch is in non-querier mode in the VLAN. If the querier expiry interval timer expires, the snooping switch moves into querier mode.</li> <li>Disabled – The snooping querier is not operational on the VLAN. The snooping querier moves to the disabled mode when MLD snooping is not operational on the VLAN, when the querier address is not configured, or the network management address is not configured.</li> </ul>
Version	The operational MLD protocol version of the querier.
Last IPv6 Address	The IPv6 address of the last querier from which a query was snooped on the VLAN.
Last Version	The MLD protocol version of the last querier from which a query was snooped on the VLAN.
Max Response Time (Seconds)	The maximum response time to be used in the queries that are sent by the snooping querier.

Click Refresh to refresh the page with the most current data from the switch.

## 5.20 Creating Port Channels

Port-channels, which are also known as link aggregation groups (LAGs), allow you to combine multiple full-duplex Ethernet links into a single logical link. Network devices treat the aggregation as if it were a single link, which increases fault tolerance and provides load sharing. You assign the port-channel (LAG) VLAN membership after you create a port-channel. The port channel by default becomes a member of the management VLAN.

A port-channel (LAG) interface can be either static or dynamic, but not both. All members of a port channel must participate in the same protocols. A static port-channel interface does not require a partner system to be able to aggregate its member ports.

### NOTICE

If you configure the maximum number of dynamic port-channels (LAGs) that your platform supports, additional port-channels that you configure are automatically static.

Static LAGs are supported. When a port is added to a LAG as a static member, it neither transmits nor receives LACPDUs.

### 5.20.1 Port Channel Summary

Use the Port Channel Summary page to group one or more full duplex Ethernet links to be aggregated together to form a port-channel, which is also known as a link aggregation group (LAG). The switch can treat the port-channel as if it were a single link.

To access the Port Channel Summary page, click Switching > Port Channel > Summary in the navigation menu.

Figure 231: Port Channel Summary

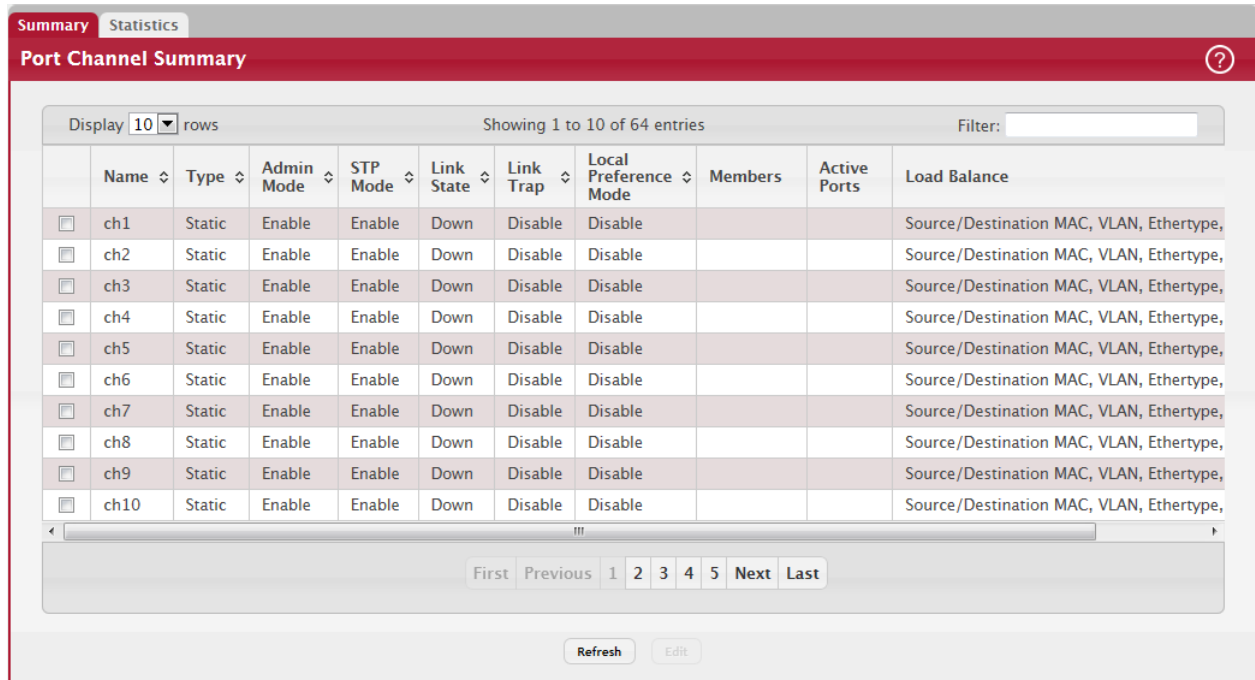


Table 216: Port Channel Summary Fields

Field	Description
Name	Identifies the user-configured text name of the port channel.
Type	The type of port channel: <ul style="list-style-type: none"> <li>• Dynamic – Uses Link Aggregation Control Protocol (LACP) Protocol Data Units (PDUs) to exchange information with the link partners to help maintain the link state. To utilize Dynamic link aggregation on this port channel, the link partner must also support LACP.</li> <li>• Static – Does not require a partner system to be able to aggregate its member ports. When a port is added to a port channel as a static member, it neither transmits nor receives LACP PDUs.</li> </ul> When configuring a port channel, use the Static Mode field to set the port channel type. If the Static Mode is disabled, the port channel type is Dynamic.
Admin Mode	Select enable or disable from the pull-down menu. When the Port Channel is disabled no traffic will flow and LACPDUs will be dropped, but the links that form the Port Channel will not be released. The factory default is enable.
STP Mode	Shows whether the Spanning Tree Protocol (STP) Administrative Mode is enabled or disabled on the port channel
Link State	Indicates whether the link is Up or Down.
Link Trap	Shows whether to send traps when link status changes. If the status is Enabled, traps are sent.
Members	Lists the ports that are members of the Port Channel, in Slot/Port notation (Unit/Slot/Port for stackable systems). There can be a maximum of 8 ports assigned to a Port Channel.

Table 216: Port Channel Summary Fields (Continued)

Field	Description
Active Ports	Lists the ports that are actively participating members of this Port Channel, in Slot/Port notation (Unit/Slot/Port for stackable systems).
Load Balance	<p>The algorithm used to distribute traffic load among the physical ports of the port channel while preserving the per-flow packet order. The packet attributes the load-balancing algorithm can use to determine the outgoing physical port include the following:</p> <ul style="list-style-type: none"> <li>• Source MAC, VLAN, EtherType, Incoming Port</li> <li>• Destination MAC, VLAN, EtherType, Incoming Port</li> <li>• Source/Destination MAC, VLAN, EtherType, Incoming Port</li> <li>• Source IP and Source TCP/UDP Port Fields</li> <li>• Destination IP and Destination TCP/UDP Port Fields</li> <li>• Source/Destination IP and TCP/UDP Port Fields</li> <li>• Enhanced Hashing Mode</li> </ul>

### 5.20.2 Port Channel Configuration

Use the Port Channel Configuration page to group one or more full duplex Ethernet links to be aggregated together to form a port-channel, which is also known as a link aggregation group (LAG). The switch treats the port-channel as if it were a single link.

To access the Port Channel Configuration page, click Switching > Port Channel > Summary in the navigation menu. Select a port and click Edit.

Figure 232: Port Channel Configuration

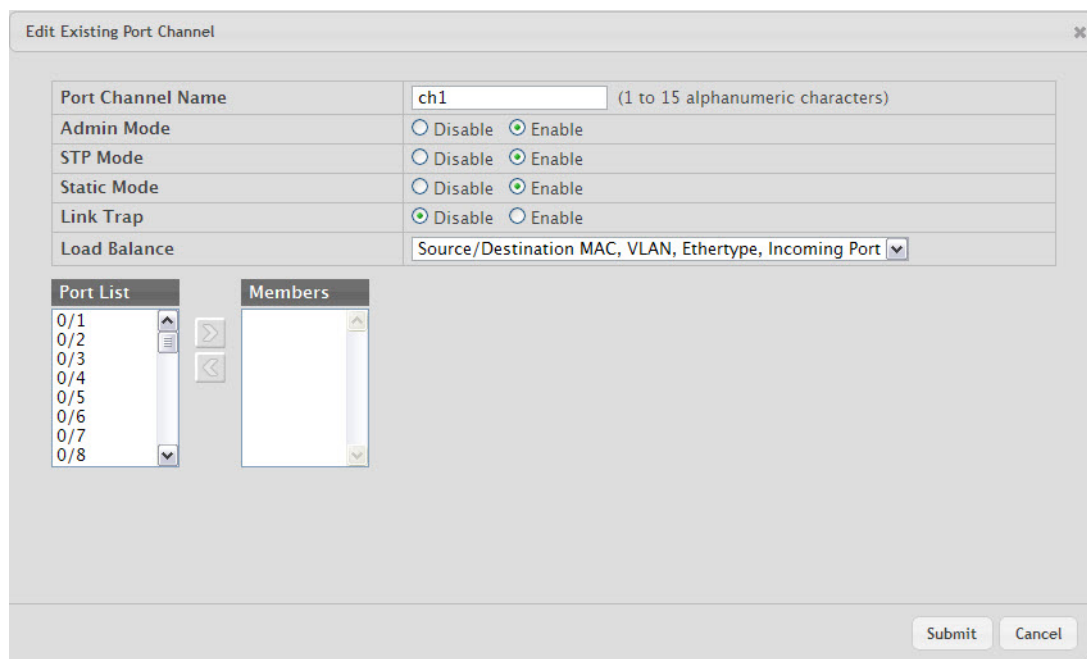


Table 217: Port Channel Configuration Fields

Field	Description
Port Channel Inter- face	Select the port channel to configure. The port channel follows a Slot/Port (or Unit/Slot/Port for stacking platforms) interface naming convention, where the slot is 3.
Port Channel Name	Enter the name you want assigned to the Port Channel. You may enter any string of up to 15 alphanumeric characters. You must specify a valid name to create the Port Channel.

Table 217: Port Channel Configuration Fields (Continued)

Field	Description
Link Trap	Specify whether you want to have a trap sent when link status changes. The factory default is enable, which will cause the trap to be sent.
Administrative Mode	Select enable or disable from the pull-down menu. When the Port Channel is disabled no traffic will flow and LACPDU's will be dropped, but the links that form the Port Channel will not be released. The factory default is enable.
Link Status	Indicates whether the link is Up or Down.
STP Mode	Select the Spanning Tree Protocol (STP) Administrative Mode associated with the Port Channel: <ul style="list-style-type: none"> <li>• Disable: Spanning tree is disabled for this Port Channel.</li> <li>• Enable: Spanning tree is enabled for this Port Channel.</li> </ul>
Static Mode	Select enable or disable from the pull-down menu. The factory default is Disable. <ul style="list-style-type: none"> <li>• Enable: The port channel is statically maintained, which means it does not transmit or process received LAGPDUs. The member ports do not transmit LAGPDUs and all the LAGPDUs it may receive are dropped. A static port-channel interface does not require a partner system to be able to aggregate its member ports.</li> <li>• Disable: The port channel is dynamically maintained. The interface transmits and processes LAGPDUs and requires a partner system.</li> </ul>
Local Preference Mode	This field is available only on systems that support stacking. When this option is enabled, the LAG-destined unicast traffic egresses only out of members of the LAG interface on the local unit. This feature makes sure that the LAG-destined unicast traffic does not cross the external stack link when the LAG has members on the local unit.
Load Balance	Select the hashing algorithm used to distribute the traffic load among available physical ports in the LAG. The range of possible values may vary with the type of switch. The possible values are: <ul style="list-style-type: none"> <li>• Source MAC, VLAN, EtherType, and source port</li> <li>• Destination MAC, VLAN, EtherType and source port</li> <li>• Source/Destination MAC, VLAN, EtherType, and source port</li> <li>• Source IP and Source TCP/UDP Port</li> <li>• Destination IP and Destination TCP/UDP Port</li> <li>• Source/Destination IP and source/destination TCP/UDP Port</li> <li>• Enhanced hashing mode</li> </ul>
Port Channel Members	After you create one or more port channel, this field lists the members of the Port Channel. If there are no port channels on the system, this field is not present.
Unit/Slot/Port Participation	This column lists the physical ports available on the system. Select each port's membership status for the Port Channel you are configuring. There can be a maximum of 8 ports assigned to a Port Channel. <ul style="list-style-type: none"> <li>• Include: The port participates in the port channel.</li> <li>• Exclude: The port does not participate in the port channel, which is the default.</li> </ul>
Membership Conflicts	Shows ports that are already members of other Port Channels. A port may only be a member of one Port Channel at a time. If the entry is blank, the port is not currently a member of any Port Channel.

- If you make any changes to this page, click Submit to apply the changes to the system.
- To remove a port channel, select it from the Port Channel Name drop-down menu and click delete. All ports that were members of this Port Channel are removed from the Port Channel and included in the default VLAN. This field will not appear when a new Port Channel is being created.

### 5.20.3 Port Channel Statistics

This page displays the flap count for each port channel and their member ports. A flap occurs when a port-channel interface or port-channel member port goes down.

To access the Port Channel Statistics page, click Switching > Port Channel > Statistics in the navigation menu.

Figure 233: Port Channel Statistics

Interface	Channel Name	Type	Flap Count
0/1/1	ch1	Port Channel	0
0/1/2	ch2	Port Channel	0
0/1/3	ch3	Port Channel	0
0/1/4	ch4	Port Channel	0
0/1/5	ch5	Port Channel	0
0/1/6	ch6	Port Channel	0

Table 218: Port Channel Statistics Fields

Field	Description
Interface	The port channel or member port (physical port) associated with the rest of the data in the row.
Channel Name	The port channel name associated with the port channel. For a physical port, this field identifies the name of the port channel of which the port is a member.
Type	The interface type, which is either Port Channel (logical link-aggregation group) or Member Port (physical port).
Flap Count	The number of times the interface has gone down. The counter for a member port is incremented when the physical port is either manually shut down by the administrator or when its link state is down. When a port channel is administratively shut down, the flap counter for the port channel is incremented, but the flap counters for its member ports are not affected. When all active member ports for a port channel are inactive (either administratively down or link down), then the port channel flap counter is incremented.
Clear Counters (Button)	Click this button to reset the flap counters for all port channels and member ports to 0.

Click Refresh to display the latest information from the router.

## 5.21 Viewing Multicast Forwarding Database Information

The Layer 2 Multicast Forwarding Database (MFDB) is used by the switch to make forwarding decisions for packets that arrive with a multicast destination MAC address. By limiting multicasts to only certain ports in the switch, traffic is prevented from going to parts of the network where that traffic is unnecessary.

When a packet enters the switch, the destination MAC address is combined with the VLAN ID and a search is performed in the Layer 2 Multicast Forwarding Database. If no match is found, the packet is either flooded to all ports in the VLAN or discarded, depending on the switch configuration. If a match is found, the packet is forwarded only to the ports that are members of that multicast group.

### 5.21.1 MFDB Table

Use the MFDB Table page to view the port membership information for all active multicast address entries. The key for an entry consists of a VLAN ID and MAC address pair. Entries may contain data for more than one protocol.

To access the MFDB Table page, click Switching > Multicast Forwarding Database > Summary in the navigation menu.



Figure 234: MFDB Table

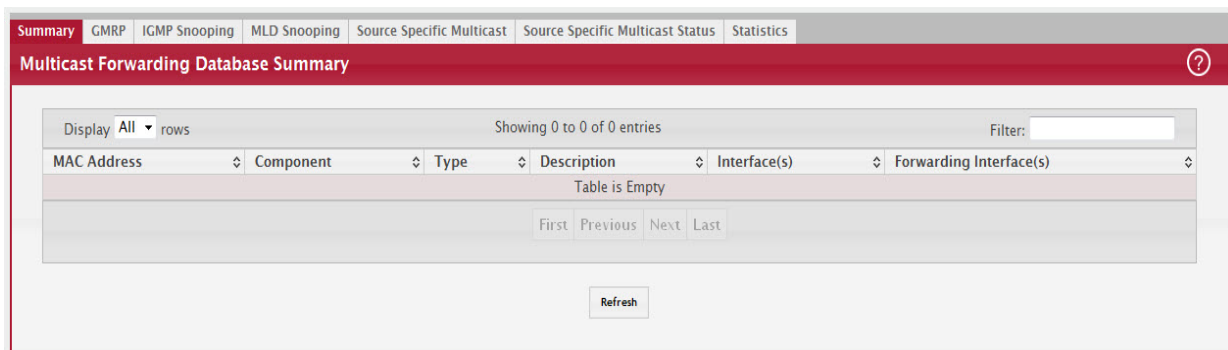


Table 219: MFDB Summary Fields

Field	Description
MAC Address	The VLAN ID (the first two groups of hexadecimal digits) and multicast MAC address (the last six groups of hexadecimal digits) that has been added to the MFDB.
Component	The feature on the device that was responsible for adding the entry to the multicast forwarding database, which is one of the following: <ul style="list-style-type: none"> <li>• IGMP Snooping – A Layer 2 feature that allows the device to dynamically add or remove ports from IPv4 multicast groups by listening to IGMP join and leave requests.</li> <li>• MLD Snooping – A Layer 2 feature that allows the device to dynamically add or remove ports from IPv6 multicast groups by listening to MLD join and leave requests.</li> <li>• GMRP – Generic Address Resolution Protocol (GARP) Multicast Registration Protocol, which helps help control the flooding of multicast traffic by keeping track of group membership information.</li> <li>• Static Filtering – A static MAC filter that was manually added to the address table by an administrator.</li> </ul>
Type	The type of entry, which is one of the following: <ul style="list-style-type: none"> <li>• Static – The entry has been manually added to the MFDB by an administrator.</li> <li>• Dynamic – The entry has been added to the MFDB as a result of a learning process or protocol.</li> </ul>
Description	A text description of this multicast table entry.
Interfaces	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.
Forwarding Inter- faces	The list of forwarding interfaces. This list does not include any interfaces that are listed as static filtering interfaces.

- To search for a MAC address if the list is too long to scan, enter the MAC address in hex format and click Search.
- Click Refresh to update the information on the screen with the most current data.

### 5.21.2 GMRP Table

Use the GMRP Table page to display the entries in the multicast forwarding database (MFDB) that were added by using the GARP Multicast Registration Protocol (GMRP).

To access the MFDB Table page, click Switching > Multicast Forwarding Database > GMRP in the navigation menu.

Figure 235: GMRP Table

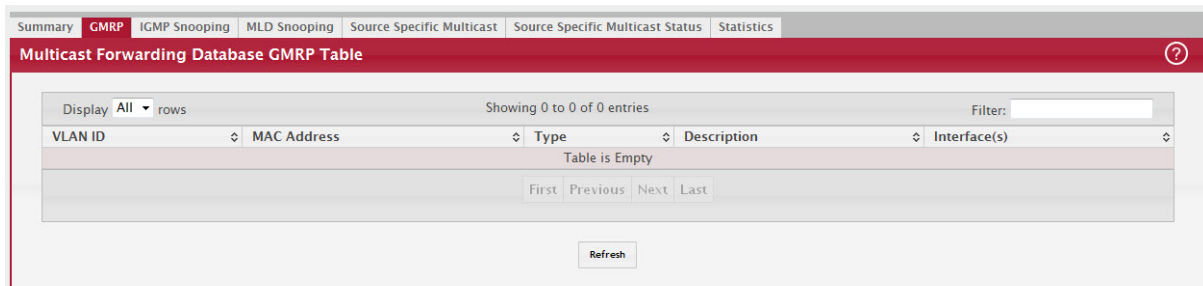


Table 220: GMRP Fields

Field	Description
VLAN ID	The VLAN ID associated with the entry in the MFDB.
MAC Address	The multicast MAC address associated with the entry in the MFDB.
Type	The type of entry, which is one of the following: <ul style="list-style-type: none"> <li>• Static – The entry has been manually added to the MFDB by an administrator.</li> <li>• Dynamic – The entry has been added to the MFDB as a result of a learning process or protocol. Entries that appear on this page have been added by using GARP.</li> </ul>
Description	A text description of this multicast table entry.
Interface(s)	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.

- Click Refresh to update the information on the screen with the most current data.

### 5.2.1.3 IGMP Snooping Table

This page displays the entries in the multicast forwarding database (MFDB) that were added because they were discovered by the IGMP snooping feature. IGMP snooping allows the device to dynamically add or remove ports from IPv4 multicast groups by listening to IGMP join and leave requests.

To access the MFDB Table page, click Switching > Multicast Forwarding Database > IGMP Snooping in the navigation menu.

Figure 236: IGMP Snooping Table

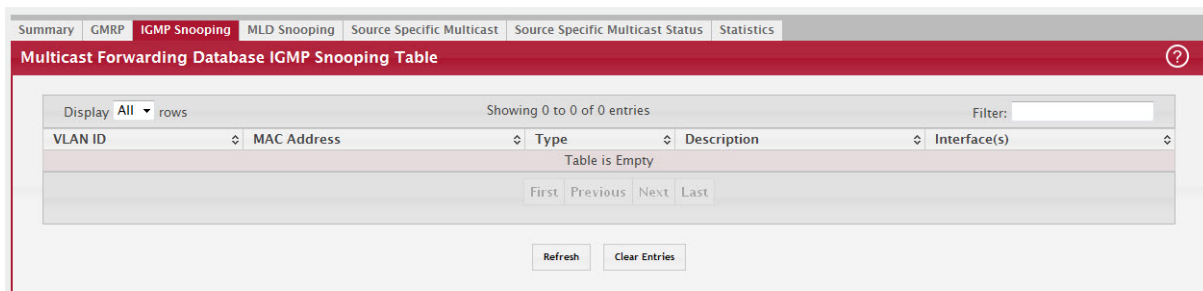


Table 221: IGMP Snooping Fields

Field	Description
VLAN ID	The VLAN ID associated with the entry in the MFDB.
MAC Address	The multicast MAC address associated with the entry in the MFDB.

Table 221: IGMP Snooping Fields (Continued)

Field	Description
Type	The type of entry, which is one of the following: <ul style="list-style-type: none"> <li>Static – The entry has been manually added to the MFDB by an administrator.</li> <li>Dynamic – The entry has been added to the MFDB as a result of a learning process or protocol. Entries that appear on this page have been learned by examining IGMP messages.</li> </ul>
Description	A text description of this multicast table entry.
Interfaces	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.

- Click Refresh to update the information on the screen with the most current data.

### 5.21.4 MLD Snooping Table

This page displays the entries in the multicast forwarding database (MFDB) that were added because they were discovered by the MLD snooping feature. MLD snooping allows the device to dynamically add or remove ports from IPv6 multicast groups by listening to MLD join and leave requests.

To access the MFDB Table page, click Switching > Multicast Forwarding Database > MLD Snooping in the navigation menu.

Figure 237: MLD Snooping Table

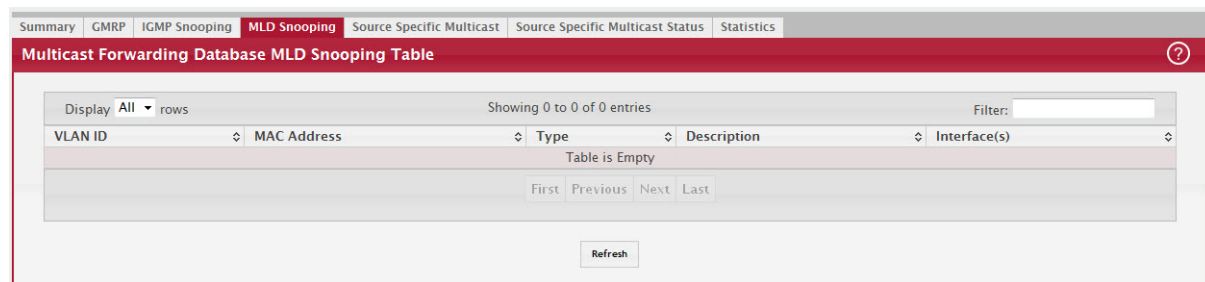


Table 222: MLD Snooping Fields

Field	Description
VLAN ID	The VLAN ID associated with the entry in the MFDB.
MAC Address	The multicast MAC address associated with the entry in the MFDB.
Type	The type of entry, which is one of the following: <ul style="list-style-type: none"> <li>Static – The entry has been manually added to the MFDB by an administrator.</li> <li>Dynamic – The entry has been added to the MFDB as a result of a learning process or protocol. Entries that appear on this page have been learned by examining MLD messages.</li> </ul>
Description	A text description of this multicast table entry.
Interfaces	The list of interfaces that will forward or filter traffic sent to the multicast MAC address.

- Click Refresh to update the information on the screen with the most current data.

### 5.21.5 Source Specific Multicast

This page displays the entries in the multicast forwarding database (MFDB) for source specific multicast, those were added because they were discovered by the IGMP Snooping or MLD Snooping feature.

To access the MFDB Table page, click Switching > Multicast Forwarding Database > Source Specific Multicast in the navigation menu.

Figure 238: Source Specific Multicast

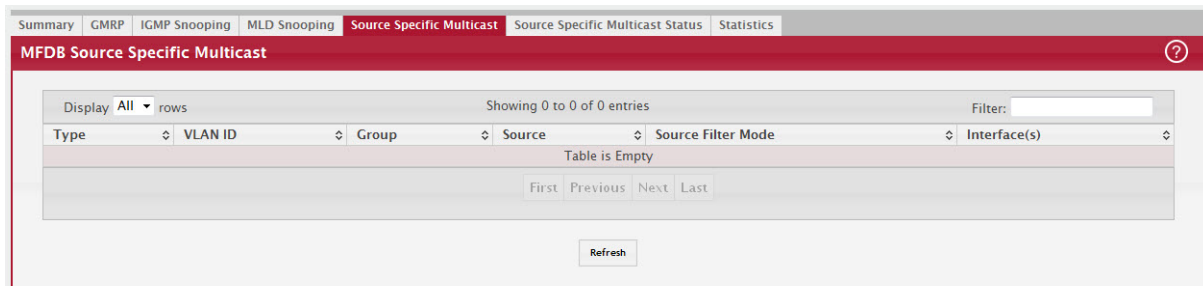


Table 223: Source Specific Multicast Fields

Field	Description
Type	Type of snooping. The values can be either IGMP Snooping or MLD Snooping.
VLAN ID	VLAN on which the entry is learned.
Group	The multicast group address.
Source	The source address.
Source Filter Mode	The source filter mode (Include/Exclude) for the specified group.
Interface(s)	Specifies the list of interfaces on which a incoming packet is forwarded.

Click Refresh to update the information on the screen with the most current data.

### 5.21.6 Source Specific Multicast Status

This page displays the entries in the multicast forwarding database (MFDB) for source specific multicast, those were added because they were discovered by the IGMP Snooping or MLD Snooping feature.

To access the MFDB Table page, click Switching > Multicast Forwarding Database > Source Specific Multicast Status in the navigation menu.

Figure 239: Source Specific Multicast Status

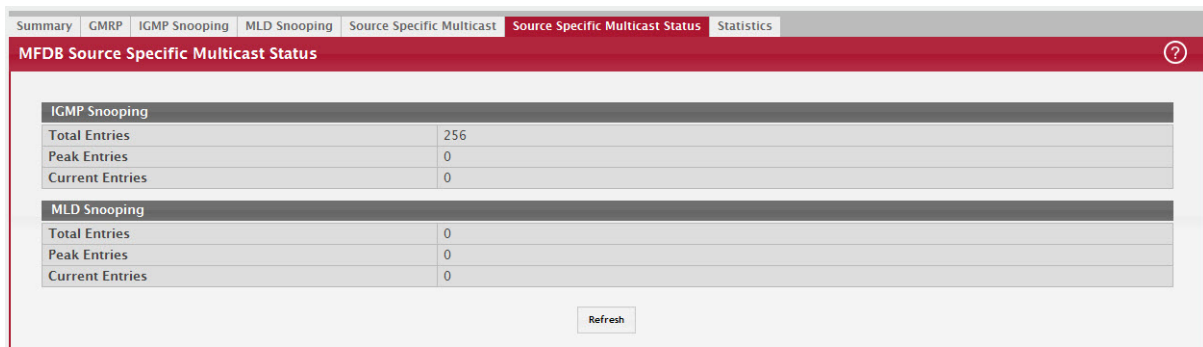


Table 224: Source Specific Multicast Status Fields

Field	Description
IGMP Snooping	
Total Entries	The total number of entries that can possibly be in IGMP snooping's SSMFDB.
Peak Entries	The largest number of entries that have been present in the IGMP snooping's SSMFDB.
Current Entries	The current number of entries in the IGMP snooping's SSMFDB.

Table 224: Source Specific Multicast Status Fields (Continued)

Field	Description
MLD Snooping	
Total Entries	The total number of entries that can possibly be in MLD snooping's SSMFDB.
Peak Entries	The largest number of entries that have been present in the MLD snooping's SSMFDB.
Current Entries	The current number of entries in the MLD snooping's SSMFDB.

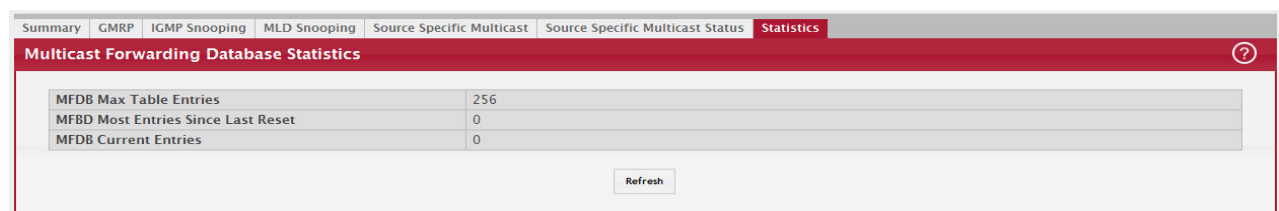
Click Refresh to update the information on the screen with the most current data.

### 5.21.7 MFDB Statistics

Use the multicast forwarding database Stats page to view statistical information about the MFDB table.

To access the Stats page, click Switching > Multicast Forwarding Database > Statistics in the navigation menu.

Figure 240: Multicast Forwarding Database Statistics



Field	Description
MFDB Max Table Entries	256
MFDB Most Entries Since Last Reset	0
MFDB Current Entries	0

Refresh

Table 225: Multicast Forwarding Database Statistics Fields

Field	Description
MFDB Max Table Entries	The maximum number of entries that the multicast forwarding database can hold.
MFDB Most Entries Since Last Reset	The largest number of entries that have been present in the multicast forwarding database since the device was last reset. This value is also known as the MFDB high-water mark.
MFDB Current Entries	The current number of entries in the multicast forwarding database.

Click Refresh to update the information on the screen with the most current data.

## 5.22 Multicast VLAN Registration

Multicast VLAN Registration (MVR) allows the switch to listen to the Internet Group Management Protocol (IGMP) frames. Both protocols operate independently from each other and can be enabled on the switch interfaces. In such case, MVR listens to the Join and Report messages only for the statically configured groups. All other groups are managed by IGMP snooping. MVR uses the multicast VLAN, a dedicated VLAN used to transfer multicast traffic over the network avoiding duplication of multicast streams for clients in different VLANs.

### 5.22.1 MVR Global Configuration

Use this page to view and configure the global settings for MVR. To access the MVR Global Configuration page, click Switching > MVR > Global.

Figure 241: MVR Global Configuration

Table 226: MVR Global Configuration Fields

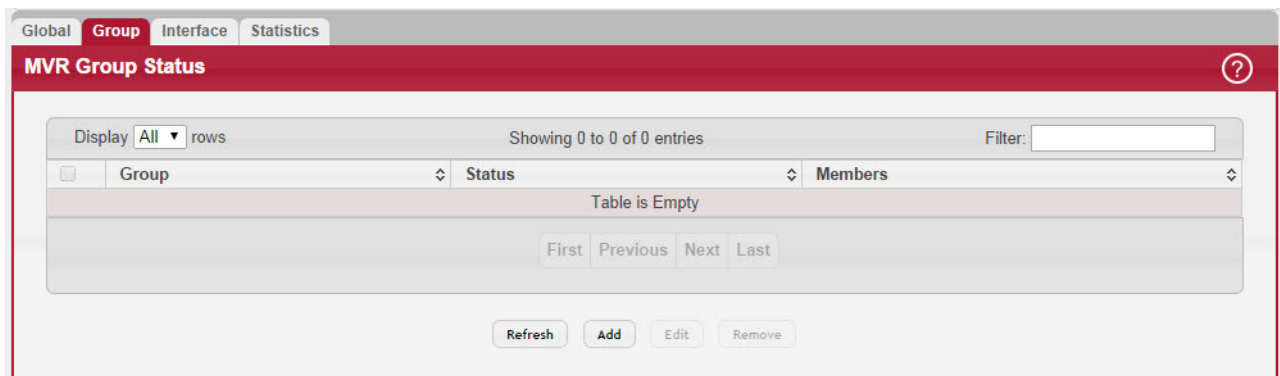
Field	Description
Admin Mode	The administrative mode of MVR on the device.
MVR Mode	The MVR learning mode, which can be one of the following: <ul style="list-style-type: none"> <li>Compatible – MVR does not learn source ports membership; instead, all source ports are members of all groups by default. MVR does not forward IGMP Joins and Leaves from the hosts to the router.</li> <li>Dynamic – MVR learns source ports membership from IGMP queries. MVR forwards the IGMP Joins and Leaves from the hosts to the router.</li> </ul> The multicast traffic is forwarded only to the receiver ports that joined the group, either by IGMP Joins or MVR static configuration.
Multicast VLAN	A dedicated VLAN used to transfer multicast traffic over the network, avoiding duplication of multicast streams for clients in different VLANs.
Maximum Multicast Groups	The maximum number of membership groups that can be statically configured in the MVR database.
Current Multicast Groups	The current number of membership groups that are statically configured in the MVR database.
Query Response Time	The maximum time to wait for an IGMP membership report on a receiver port before removing the port from the multicast group. The query time is specified in tenths of a second.

- If you make any configuration changes, click Submit to apply the new settings to the switch.
- Click Refresh to update the information on the screen with the most current data.

### 5.22.2 MVR Group Status

Use this page to view or configure MVR groups. MVR maintains two types of group entries in its database, Static and Dynamic. Static entries are configured by the administrator and Dynamic entries are learned by MVR on the source ports. To access the MVR Group Status page, click Switching > MVR > Group.

Figure 242: MVR Group Status



Use the buttons to perform the following tasks:

- To add a group, click Add and specify a group address in the available field.
- To edit a configured group, select the entry to modify and click Edit. Then, configure which interfaces should be members of that group.
- To remove one or more configured groups, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 227: MVR Group Status Fields

Field	Description
Group	The multicast group address.
Status	The status of the group, which can be one of the following: <ul style="list-style-type: none"> <li>• Active – Group has one or more MVR ports participating.</li> <li>• Inactive – Group has no MVR ports participating.</li> </ul>
Members	The list of interfaces which participate in the MVR group. In the compatible mode, all source ports are members of all groups by default.
Contiguous Group Count	This field is available in the Add Group dialog. Specify the desired number of groups to be created starting with the entered group address. The default contiguous group count is 1.
Available Interfaces	This field is available in the Edit Group Configuration dialog. The interfaces that can be added to the group. To move an interface between the Available Interfaces and Group Interfaces fields, click the interface (or CTRL + click to select multiple interfaces), and then click the appropriate arrow to move the selected interfaces to the desired field.
Group Interfaces	This field is available in the Edit Group Configuration dialog. The interfaces that are members of the MVR group.

Click Refresh to update the information on the screen with the most current data.

### 5.22.3 MVR Interface Status

Use this page to configure MVR settings on specific interfaces. To configure the settings for one or more interfaces, select each entry to modify and click Edit. The same MVR settings are applied to all selected interfaces. To access the MVR Interface Status page, click Switching > MVR > Interface.

Figure 243: MVR Interface Status

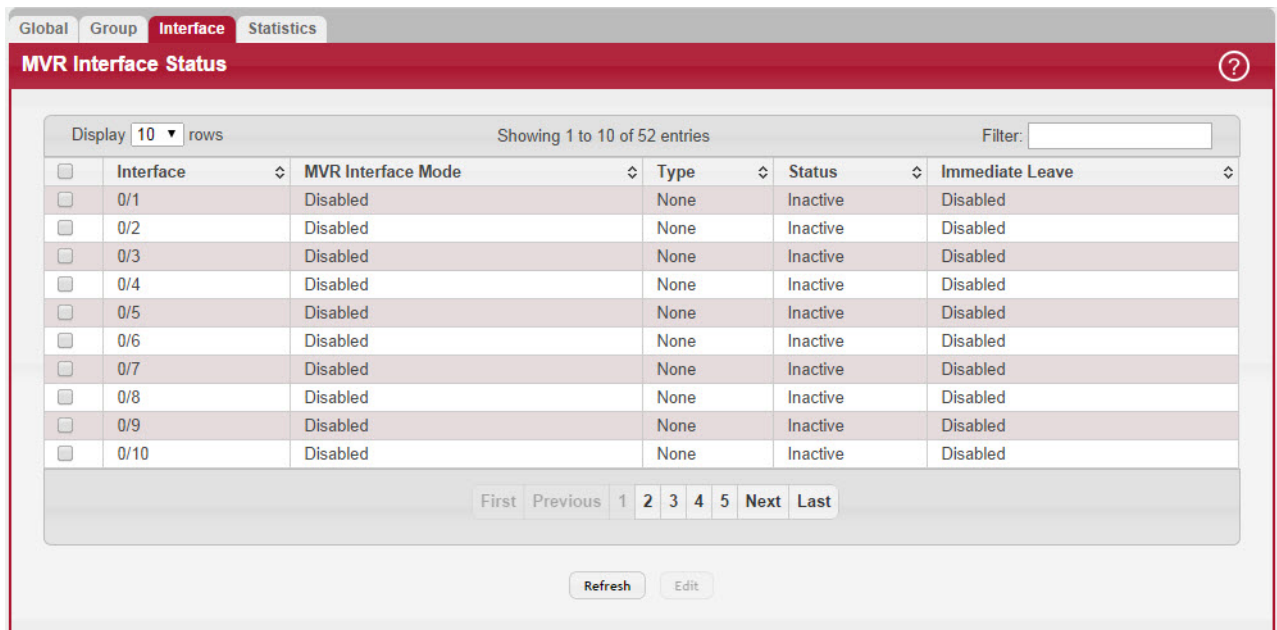


Table 228: MVR Interface Status Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring MVR settings, this field identifies the interface(s) that are being configured.
MVR Interface Mode	The administrative mode of MVR on the interface. MVR must be enabled globally and on an interface to listen to the Join and Report messages for the configured groups.
Type	The type of interface, which can be one of the following: <ul style="list-style-type: none"> <li>• Source – The port where multicast traffic is flowing to. It must be a member of the multicast VLAN.</li> <li>• Receiver – The port where listening host is connected to the switch. It must not be a member of the multicast VLAN.</li> <li>• None – The port is not an MVR port.</li> </ul>
Status	The active state of the interface, which can be one of the following: <ul style="list-style-type: none"> <li>• Active – The port has link up and is in the forwarding state.</li> <li>• Inactive – The port may not have link up, not be in the forwarding state, or both.</li> </ul>
Immediate Leave	The MVR immediate leave mode on the interface. It can only be configured on the receiver ports. MVR handles IGMP Leaves in Normal or Immediate leave mode. When a Leave message is received, in the normal mode a general IGMP query is sent from the switch to the receiver port, whereas in the immediate leave mode the switch is immediately reconfigured not to forward specific multicast stream to the receiver port. The immediate leave mode is used for ports where only one client may be connected.

Click Refresh to update the information on the screen with the most current data.

### 5.22.4 MVR Statistics

Use this page to view statistical information about IGMP packets intercepted by MVR. To access the MVR Statistics page, click Switching > MVR > Statistics.



Figure 244: MVR Statistics

Statistics	Transmit	Receive
IGMP Queries	0	0
IGMPv1 Reports	0	0
IGMPv2 Reports	0	0
IGMP Leaves	0	0
Packet Failures	0	0

[Refresh](#)

Table 229: MVR Statistics Fields

Field	Description
IGMP Queries	The total number of IGMP Queries successfully transmitted or received by the processor.
IGMPv1 Reports	The total number of IGMPv1 Reports successfully transmitted or received by the processor.
IGMPv2 Reports	The total number of IGMPv2 Reports successfully transmitted or received by the processor.
IGMP Leaves	The total number of IGMP Leaves successfully transmitted or received by the processor.
Packet Failures	The total number of packets which failed to get transmitted or received by the processor.

Click Refresh to update the information on the screen with the most current data.

### 5.23 Configuring Protected Ports

Use this page to configure and view protected ports groups. A port that is a member of a protected ports group is a protected port. A port that is not a member of any protected ports group is an unprotected port. Each port can be a member of only one protected ports group. Ports in the same protected ports group cannot forward traffic to other protected ports within the group, even if they are members of the same VLAN. However, a port in a protected ports group can forward traffic to ports that are in a different protected ports group. A protected port can also forward traffic to unprotected ports. Unprotected ports can forward traffic to both protected and unprotected ports.

To access the Protected Ports Configuration page, click Switching > Protected Ports > Configuration in the navigation menu.

Figure 245: Protected Ports Configuration

Configuration

Protected Ports Configuration

Display **All** rows      Showing 0 to 0 of 0 entries      Filter:

<input type="checkbox"/> Group Name	Protected Ports
Table is Empty	

First Previous Next Last

[Refresh](#) [Add](#) [Edit](#) [Remove](#)

Use the buttons to perform the following tasks:

- To create a protected ports group and add ports to the group, click Add and configure the settings in the available fields.
- To change the name or the port members for an existing group, select the group to update and click Edit.
- To remove one or more protected ports groups, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

**Table 230: Protected Ports Configuration Fields**

Field	Description
Group Name	This is the configured name of the protected ports group.
Protected Ports	The ports that are members of the protected ports group. When adding a port to a protected ports group, the Available Interfaces field lists the ports that are not already members of a protected ports group. To move an interface between the Available Interfaces and Selected Interfaces fields, click the port (or CTRL + click to select multiple ports), and then click the appropriate arrow to move the port(s) to the desired field.

Click Refresh to update the information on the screen with the most current data.

## 5.24 Priority Flow Control

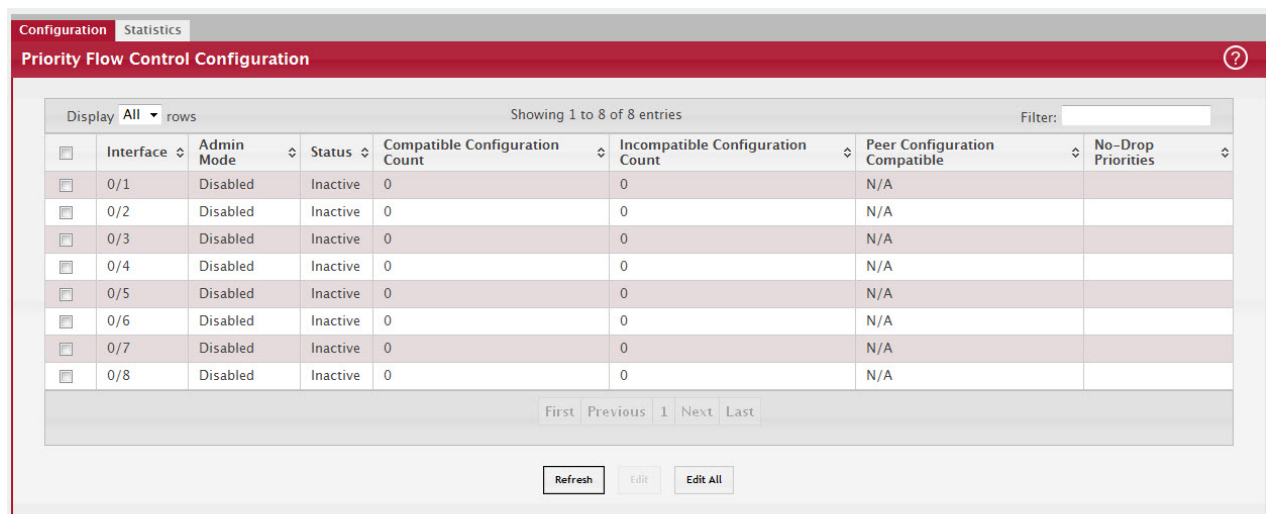
Priority Flow Control (PFC) allows a physical link to pause traffic based on the 802.1p priority field of the 802.1Q VLAN header within the frame. Ordinarily, when flow control is enabled on a physical link, it applies to all traffic on the link. When congestion occurs, the hardware sends pause frames that temporarily suspend all traffic flow on the port to help prevent buffer overflow and dropped frames. PFC allows ports to pause traffic based on its 802.1p priority value. Because PFC-enabled ports can pause the congested priority or priorities independently, protocols that are highly loss-sensitive can share the same link with traffic that has different loss tolerances.

### 5.24.1 Priority Flow Control Configuration

Use this page to configure per-port Priority-based Flow Control (PFC) settings.

To display the Priority Flow Control Configuration page, click Switching > Priority Flow Control > Configuration in the navigation menu.

**Figure 246: Priority Flow Control Configuration**



Use the buttons to perform the following tasks:

- To configure PFC settings for one or more interfaces, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.
- To apply the same PFC settings to all interfaces, click Edit All and configure the desired settings.

**Table 231: Priority Flow Control Configuration Fields**

Field	Description
Interface	The physical port associated with the rest of the data in the row. When configuring one or more ports, the Interface field in the Priority Flow Control Mapping Action window identifies the ports that are being configured.
Admin Mode	The administrative mode of PFC on the interface: <ul style="list-style-type: none"> <li>• Enabled – In times of high congestion on the link, the port will pause traffic based on the 802.1p priority of the frame and the configured priority action.</li> <li>• Disabled – If 802.3x flow control is enabled and the link is congested, the port will pause all traffic regardless of priority. If flow control is disabled, the port drops traffic during periods of high congestion.</li> </ul>
Status	The operational status of PFC on the interface.
Compatible Configuration Count	The number of compatible PFC configurations the interface has accepted from peer devices. The count does not include duplicate configurations. The PFC configuration is considered to be compatible if the no-drop priority vector matches exactly with that of the configuration source. Upstream devices should be configured so that all such devices advertise the same PFC configuration.
Incompatible Configuration Count	The number of incompatible PFC configurations the interface has received from peer devices. A PFC configuration is incompatible if the sum of no-drop priorities on all ports for the peer configuration is greater than the local system limit.
Peer Configuration Compatible	Indicates whether the local system has accepted a compatible configuration from a peer switch.
No-Drop Priorities	The 802.1p priorities that are configured as no-drop. If traffic with an 802.1p priority that is designated as no-drop is congested, the traffic is paused to prevent loss. Drop priorities do not participate in the traffic pause.
After you click Edit or Edit All, a window opens and allows you to configure the PFC settings for the selected interfaces (or for all interfaces). The following information describes the additional fields that appear in the window used for configuring per-interface PFC settings.	
Priority Action (0–7)	The action to take for each 802.1p priority value. A frame with a higher priority value is considered to be more time-sensitive than a frame with a lower priority value. If congestion occurs on the link and PFC is enabled on the port, traffic with an 802.1p priority value configured with a no-drop action is paused. Traffic with an 802.1p priority value configured with a drop action is not paused and may experience loss.

- Click Refresh to update the information on the screen with the most current data.

## 5.24.2 Priority Flow Control Statistics

This page displays information about the Priority-based Flow Control (PFC) frames transmitted and received by each interface on the device. A PFC-enabled interface transmits PFC frames during periods of congestion. The PFC frame includes information about which 802.1p priority values are configured with a no-drop (pause-enabled) action.

To display the Priority Flow Control Statistics page, click Switching > Priority Flow Control > Statistics in the navigation menu.

Figure 247: Priority Flow Control Statistics

Interface	Tx Total	Rx Total	Rx Priority 0	Rx Priority 1	Rx Priority 2	Rx Priority 3	Rx Priority 4	Rx Priority 5	Rx Priority 6	Rx Priority 7
0/1	0	0	0	0	0	0	0	0	0	0
0/2	0	0	0	0	0	0	0	0	0	0
0/3	0	0	0	0	0	0	0	0	0	0
0/4	0	0	0	0	0	0	0	0	0	0
0/5	0	0	0	0	0	0	0	0	0	0
0/6	0	0	0	0	0	0	0	0	0	0
0/7	0	0	0	0	0	0	0	0	0	0
0/8	0	0	0	0	0	0	0	0	0	0

Table 232: Priority Flow Control Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Tx Total	The total number of PFC frames the interface has sent to its link partner.
Rx Total	The total number of PFC frames the interface has received from its link partner.
Rx Priority 0 to 7	The number of PFC frames received from the link partner that specified a no-drop (pause) action for the 802.1p priority value.

- Click Refresh to update the information on the screen with the most current data.

## 5.25 Configuring Spanning Tree Protocol

The Spanning Tree Protocol (STP) provides a tree topology for any arrangement of bridges. STP also provides one path between end stations on a network, eliminating loops. Spanning tree versions supported include Common STP, Multiple STP, and Rapid STP.

Classic STP provides a single path between end stations, avoiding and eliminating loops. For information on configuring Common STP, see [Section 5.25.3: "CST Port Configuration"](#).

Multiple Spanning Tree Protocol (MSTP) supports multiple instances of Spanning Tree to efficiently channel VLAN traffic over different interfaces. Each instance of the Spanning Tree behaves in the manner specified in IEEE 802.1w, Rapid Spanning Tree (RSTP), with slight modifications in the working but not the end effect (chief among the effects, is the rapid transitioning of the port to 'Forwarding'). The difference between the RSTP and the traditional STP (IEEE 802.1D) is the ability to configure and recognize full duplex connectivity and ports which are connected to end stations, resulting in rapid transitioning of the port to 'Forwarding' state and the suppression of Topology Change Notification. These features are represented by the parameters 'pointtopoint' and 'edgeport'. MSTP is compatible to both RSTP and STP. It behaves appropriately to STP and RSTP bridges. A MSTP bridge can be configured to behave entirely as a RSTP bridge or a STP bridge.

### NOTICE

For two bridges to be in the same region, the force version should be 802.1S and their configuration name, digest key, and revision level should match. For more information about regions and their effect on network topology, refer to the IEEE 802.1Q standard.

### 5.25.1 Switch Configuration/Status

The Spanning Tree Switch Configuration/Status page contains fields for enabling STP on the switch.

To display the Spanning Tree Switch Configuration/Status page, click Switching > Spanning Tree > Switch in the navigation menu.

Figure 248: Spanning Tree Switch Configuration

Table 233: Spanning Tree Switch Configuration Fields

Field	Description
Spanning Tree Admin Mode	The administrative mode of STP on the device. When enabled, the device participates in the root bridge election process and exchanges Bridge Protocol Data Units (BPDUs) with other switches in the spanning tree to determine the root path costs and maintain topology information.
Force Protocol Version	The STP version the device uses, which is one of the following: <ul style="list-style-type: none"> <li>• IEEE 802.1d – Classic STP provides a single path between end stations, avoiding and eliminating loops.</li> <li>• IEEE 802.1w – Rapid Spanning Tree Protocol (RSTP) behaves like classic STP but also has the ability to configure and recognize full-duplex connectivity and ports that are connected to end stations, resulting in rapid transitioning of the port to the Forwarding state and the suppression of Topology Change Notifications.</li> <li>• IEEE 802.1s – Multiple Spanning Tree Protocol (MSTP) includes all the advantages of RSTP and also supports multiple spanning tree instances to efficiently channel VLAN traffic over different interfaces. MSTP is compatible with both RSTP and STP.</li> </ul>
Configuration Name	The name of the MSTP region. Each switch that participates in the same MSTP region must share the same Configuration Name, Configuration Revision Level, and MST-to-VLAN mappings.
Configuration Revision Level	The revision number of the MSTP region. This number must be the same on all switches that participate in the MSTP region.
Configuration Digest Key	The 16 byte signature of type HMAC-MD5 created from the MST Configuration Table (a VLAN ID-to-MST ID mapping).
Configuration Format Selector	The version of the configuration format being used in the exchange of BPDUs.

- If you make any configuration changes, click Submit to apply the new settings to the switch.
- Click Refresh to update the information on the screen with the most current data.

## 5.25.2 CST Configuration

Use the CST Configuration page to configure the Common Spanning Tree (CST) settings. The settings and information on this page define the device within the spanning tree topology that connects all STP/RSTP bridges and MSTP regions.

To display the CST Configuration page, click Switching > Spanning Tree > CST in the navigation menu.

Figure 249: Spanning Tree CST

Field	Value	Range/Unit
Bridge Priority	8000	(0 to F000 hex)
Bridge Max Age	20	(6 to 40)
Bridge Hello Time	2	
Bridge Forward Delay	15	(4 to 30)
Spanning Tree Maximum Hops	20	(6 to 40)
BPDUs Guard	<input type="checkbox"/>	
BPDUs Filter	<input type="checkbox"/>	
Spanning Tree Tx Hold Count	6	(1 to 10)
Bridge Identifier	80:00:00:0C:29:D3:80:EA	
Time Since Topology Change	124d:18:09:29	
Topology Change Count	0	
Topology Change	False	
Designated Root	80:00:00:0C:29:D3:80:EA	
Root Path Cost	0	
Root Port	00:00	
Max Age	20	
Forward Delay	15	
Hold Time	6	
CST Regional Root	80:00:00:0C:29:D3:80:EA	
CST Path Cost	0	

Table 234: Spanning Tree CST Fields

Field	Description
Bridge Priority	The value that helps determine which bridge in the spanning tree is elected as the root bridge during STP convergence. A lower value increases the probability that the bridge becomes the root bridge.
Bridge Max Age	The amount of time a bridge waits before implementing a topological change.
Bridge Hello Time	The amount of time the root bridge waits between sending hello BPDUs.
Bridge Forward Delay	The amount of time a bridge remains in a listening and learning state before forwarding packets.
Spanning Tree Maximum Hops	The maximum number of hops a Bridge Protocol Data Unit (BPDU) is allowed to traverse within the spanning tree region before it is discarded.
BPDUs Guard	When enabled, BPDUs Guard can disable edge ports that receive BPDU packets. This prevents a new device from entering the existing STP topology. Thus devices that were originally not a part of STP are not allowed to influence the STP topology.
BPDUs Filter	When enabled, this feature filters the BPDU traffic on the edge ports. When spanning tree is disabled on a port, BPDU filtering allows BPDU packets received on that port to be dropped.
Spanning Tree Tx Hold Count	The maximum number of BPDUs that a bridge is allowed to send within a hello time window.
Bridge Identifier	A unique value that is automatically generated based on the bridge priority value and the base MAC address of the bridge. When electing the root bridge for the spanning tree, if the bridge priorities for multiple bridges are equal, the bridge with the lowest MAC address is elected as the root bridge.

Table 234: Spanning Tree CST Fields (Continued)

Field	Description
Time Since Topology Change	The amount of time that has passed since the topology of the spanning tree has changed since the device was last reset.
Topology Change Count	The number of times the topology of the spanning tree has changed.
Topology Change	Indicates whether a topology change is in progress on any port assigned to the CST. If a change is in progress the value is True; otherwise, it is False.
Designated Root	The bridge identifier of the root bridge for the CST. The identifier is made up of the bridge priority and the base MAC address.
Root Path Cost	The path cost to the designated root for the CST. Traffic from a connected device to the root bridge takes the least-cost path to the bridge. If the value is 0, the cost is automatically calculated based on port speed.
Root Port	The port on the bridge with the least-cost path to the designated root for the CST.
Max Age	The amount of time a bridge waits before implementing a topological change.
Forward Delay	The forward delay value for the root port bridge.
Hold Time	The minimum amount of time between transmissions of Configuration BPDUs.
CST Regional Root	The bridge identifier of the CST regional root. The identifier is made up of the priority value and the base MAC address of the regional root bridge.
CST Path Cost	The path cost to the CST tree regional root.

- If you make any configuration changes, click Submit to apply the new settings to the switch.
- Click Force to force the port to send out 802.1w or 802.1D BPDUs.
- Click Refresh to update the screen with most recent data.

### 5.25.3 CST Port Configuration

Use the CST Port page to view and configure the Common Spanning Tree (CST) settings for each interface on the device. To configure CST settings for an interface and to view additional information about the interface's role in the CST topology, select the interface to view or configure and click Edit.

To display the Spanning Tree CST Port Configuration/Status page, click Switching > Spanning Tree > CST Port in the navigation menu.

Figure 250: Spanning Tree CST Port

Interface	Port Role	Port Forwarding State	Port Priority	Port Path Cost	Description
1/0/1	Designated	Forwarding	0x0080	200000	
1/0/2	Designated	Forwarding	0x0080	200000	
1/0/3	Disabled	Disabled	0x0080	0	
1/0/4	Designated	Forwarding	0x0080	200000	
1/0/5	Designated	Forwarding	0x0080	200000	
1/0/6	Designated	Forwarding	0x0080	200000	
1/0/7	Designated	Forwarding	0x0080	200000	
1/0/8	Designated	Forwarding	0x0080	200000	
0/1/1	Disabled	Disabled	0x0060	0	
0/1/2	Disabled	Disabled	0x0060	0	

Table 235: Spanning Tree CST Port Fields

Field	Description
Interface	The port or link aggregation group (LAG) associated with the rest of the data in the row. When configuring CST settings for an interface, this field identifies the interface being configured.
Port Role	The role of the port within the CST, which is one of the following: <ul style="list-style-type: none"> <li>• Root – A port on the non-root bridge that has the least-cost path to the root bridge.</li> <li>• Designated – A port that has the least-cost path to the root bridge on its segment.</li> <li>• Alternate – A blocked port that has an alternate path to the root bridge.</li> <li>• Backup – A blocked port that has a redundant path to the same network segment as another port on the bridge.</li> <li>• Master – The port on a bridge within an MST instance that links the MST instance to other STP regions.</li> <li>• Disabled – The port is administratively disabled and is not part of the spanning tree.</li> </ul>
Port Forwarding State	<ul style="list-style-type: none"> <li>• Blocking – The port discards user traffic and receives, but does not send, BPDUs. During the election process, all ports are in the blocking state. The port is blocked to prevent network loops.</li> <li>• Listening – The port sends and receives BPDUs and evaluates information to provide a loop-free topology. This state occurs during network convergence and is the first state in transitioning to the forwarding state.</li> <li>• Learning – The port learns the MAC addresses of frames it receives and begins to populate the MAC address table. This state occurs during network convergence and is the second state in transitioning to the forwarding state.</li> <li>• Forwarding – The port sends and receives user traffic.</li> <li>• Disabled – The port is administratively disabled and is not part of the spanning tree.</li> </ul>
Port Priority	The priority for the port within the CST. This value is used in determining which port on a switch becomes the root port when two ports have the same least-cost path to the root. The port with the lower priority value becomes the root port. If the priority values are the same, the port with the lower interface index becomes the root port.
Port Path Cost	The path cost from the port to the root bridge.
Description	A user-configured description of the port. After you select an interface and click Edit, a window opens and allows you to edit the CST port settings and view additional CST information for the interface. The following information describes the additional fields available in the Edit CST Port Entry window.
Admin Edge Port	Select this option administratively configure the interface as an edge port. An edge port is an interface that is directly connected to a host and is not at risk of causing a loop.
Auto-calculate Port Path Cost	Shows whether the path cost from the port to the root bridge is automatically determined by the speed of the interface (Enabled) or configured manually (Disabled).
Hello Timer	The amount of time the port waits between sending hello BPDUs.
External Port Path Cost	The cost of the path from the port to the CIST root. This value becomes important when the network includes multiple regions.
Auto-calculate External Port Path Cost	Shows whether the path cost from the port to the CIST root is automatically determined by the speed of the interface (Enabled) or configured manually (Disabled).
BPDU Flood	This option determines the behavior of the interface if STP is disabled on the port and the port receives a BPDU. If BPDU flooding is enabled, the port will flood the received BPDU to all the ports on the switch that are similarly disabled for spanning tree.
BPDU Guard Effect	Shows the status of BPDU Guard Effect on the interface. When enabled, BPDU Guard Effect can disable edge ports that receive BPDU packets. This prevents a new device from entering the existing STP topology. Thus devices that were originally not a part of STP are not allowed to influence the STP topology.
Port ID	A unique value that is automatically generated based on the port priority value and the interface index.



Table 235: Spanning Tree CST Port Fields (Continued)

Field	Description
Port Up Time Since Counters Last Cleared	The amount of time that the port has been up since the counters were cleared.
Port Mode	The administrative mode of spanning tree on the port.
Designated Root	The bridge ID of the root bridge for the CST.
Designated Cost	The path cost offered to the LAN by the designated port.
Designated Bridge	The bridge ID of the bridge with the designated port.
Designated Port	The port ID of the designated port.
Topology Change Acknowledge	Indicates whether the next BPDU to be transmitted for this port will have the topology change acknowledgment flag set.
Auto Edge	When enabled, Auto Edge allows the interface to become an edge port if it does not receive any BPDUs within a given amount of time.
Edge Port	Indicates whether the interface is configured as an edge port (Enabled).
Point-to-point MAC	Indicates whether the link type for the interface is a point-to-point link.
Root Guard	When enabled, Root Guard allows the interface to discard any superior information it receives to protect the root of the device from changing. The port gets put into discarding state and does not forward any frames.
Loop Guard	When enabled, Loop Guard prevents an interface from erroneously transitioning from blocking state to forwarding when the interface stops receiving BPDUs. The port is marked as being in loop-inconsistent state. In this state, the interface does not forward frames.
TCN Guard	When enabled, TCN Guard restricts the interface from propagating any topology change information received through that interface.
CST Regional Root	The bridge ID of the bridge that has been elected as the root bridge of the CST region.
CST Path Cost	The path cost from the interface to the CST regional root.
Loop Inconsistent State	Identifies whether the interface is currently in a loop inconsistent state. An interface transitions to a loop inconsistent state if loop guard is enabled and the port stops receiving BPDUs. In this state, the interface does not transmit frames.
Transitions Into LoopInconsistent State	The number of times this interface has transitioned into loop inconsistent state.
Transitions Out Of LoopInconsistent State	The number of times this interface has transitioned out of loop inconsistent state.

- If you make any configuration changes, click Submit to apply the new settings to the switch.
- Click Cancel to cancel the change.
- Click Refresh to update the screen with most recent data.

#### 5.25.4 MST Configuration

Use the MST Configuration page to view and configure the Multiple Spanning Tree Instances (MSTIs) on the device. Multiple Spanning Tree Protocol (MSTP) allows the creation of MSTIs based upon a VLAN or groups of VLANs. Configuring MSTIs creates an active topology with a better distribution of network traffic and an increase in available bandwidth when compared to classic STP.

- Use the buttons to perform the following tasks:
- To configure a new MSTI, click Add and specify the desired settings.
- To change the Priority or the VLAN associations for an existing MSTI, select the entry to modify and click Edit.
- To remove one or more MSTIs, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

To display the Spanning Tree MST Summary page, click Switching > Spanning Tree > MST in the navigation menu.

Figure 251: Spanning Tree MST Summary

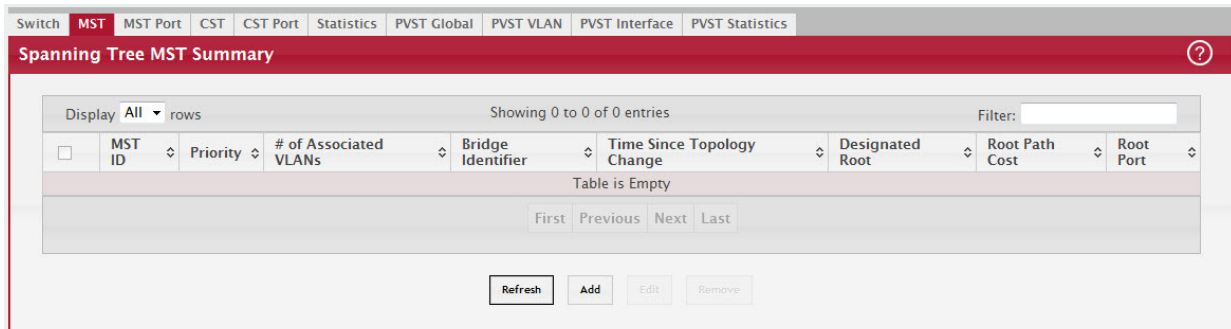


Table 236: Spanning Tree MST Summary Fields

Field	Description
MST ID	The number that identifies the MST instance.
Priority	The bridge priority for the spanning-tree instance. This value affects the likelihood that the bridge is selected as the root bridge. A lower value increases the probability that the bridge is selected as the root bridge.
# of Associated VLANs	The number of VLANs that are mapped to the MSTI. This number does not contain any information about the VLAN IDs that are mapped to the instance.
Bridge Identifier	A unique value that is automatically generated based on the bridge priority value of the MSTI and the base MAC address of the bridge. When electing the root bridge for an MST instance, if the bridge priorities for multiple bridges are equal, the bridge with the lowest MAC address is elected as the root bridge.
Time Since Topology Change	The amount of time that has passed since the topology of the MSTI has changed.
Designated Root	The bridge identifier of the root bridge for the MST instance. The identifier is made up of the bridge priority and the base MAC address.
Root Path Cost	The path cost to the designated root for this MST instance. Traffic from a connected device to the root bridge takes the least-cost path to the bridge. If the value is 0, the cost is automatically calculated based on port speed.
Root Port	The port on the bridge with the least-cost path to the designated root for the MST instance.

- If you make any configuration changes, click Submit to apply the new settings to the switch.
- Click Refresh to update the screen with most recent data.

### 5.25.5 MST Port Configuration

Use this page to view and configure the Multiple Spanning Tree (MST) settings for each interface on the device. To configure MST settings for an interface and to view additional information about the interface's role in the MST topology, first select the appropriate MST instance from the MST ID menu. Then, select the interface to view or configure and click Edit. To display the Spanning Tree MST Port Summary page, click Switching > Spanning Tree > MST Port in the navigation menu.

**NOTICE**

If no MST instances have been configured on the switch, the page displays a No MSTs Available message and does not display the fields shown in [Figure 252: "Spanning Tree MST Port Configuration," on page 285.](#)

Figure 252: Spanning Tree MST Port Configuration

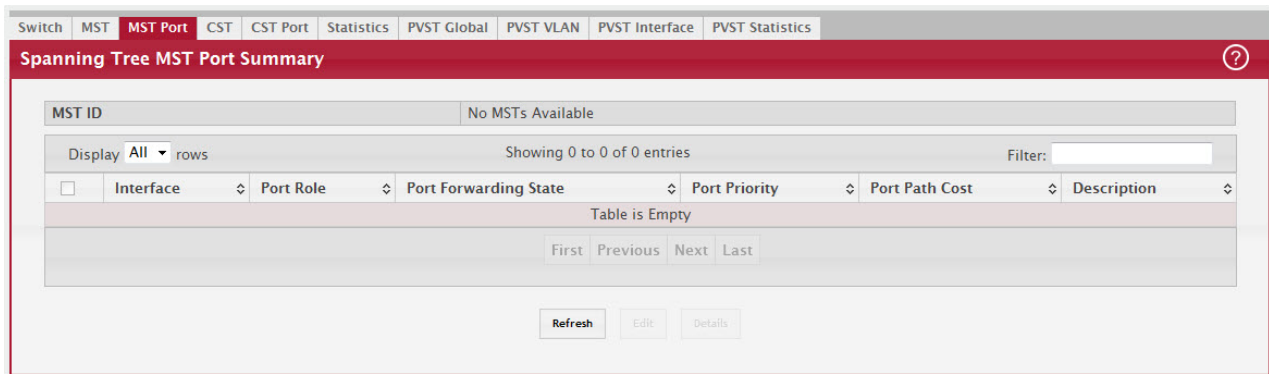


Table 237: Spanning Tree MST Port Configuration Fields

Field	Description
MST ID	The menu contains the ID of each MST instance that has been created on the device.
Interface	The port or link aggregation group (LAG) associated with the rest of the data in the row. When configuring MST settings for an interface, this field identifies the interface being configured.
Port Role	The role of the port within the MST, which is one of the following: <ul style="list-style-type: none"> <li>• Root – A port on the non-root bridge that has the least-cost path to the root bridge.</li> <li>• Designated – A port that has the least-cost path to the root bridge on its segment.</li> <li>• Alternate – A blocked port that has an alternate path to the root bridge.</li> <li>• Backup – A blocked port that has a redundant path to the same network segment as another port on the bridge.</li> <li>• Master – The port on a bridge within an MST instance that links the MST instance to other STP regions.</li> <li>• Disabled – The port is administratively disabled and is not part of the spanning tree.</li> </ul>
Port Forwarding State	<ul style="list-style-type: none"> <li>• Blocking – The port discards user traffic and receives, but does not send, BPDUs. During the election process, all ports are in the blocking state. The port is blocked to prevent network loops.</li> <li>• Listening – The port sends and receives BPDUs and evaluates information to provide a loop-free topology. This state occurs during network convergence and is the first state in transitioning to the forwarding state.</li> <li>• Learning – The port learns the MAC addresses of frames it receives and begins to populate the MAC address table. This state occurs during network convergence and is the second state in transitioning to the forwarding state.</li> <li>• Forwarding – The port sends and receives user traffic.</li> <li>• Disabled – The port is administratively disabled and is not part of the spanning tree.</li> </ul>
Port Priority	The priority for the port within the MSTI. This value is used in determining which port on a switch becomes the root port when two ports have the same least-cost path to the root. The port with the lower priority value becomes the root port. If the priority values are the same, the port with the lower interface index becomes the root port.
Port Path Cost	The path cost from the port to the root bridge.
Description	A user-configured description of the port. After you select an interface and click Edit, a window opens and allows you to edit the MST port settings and view additional MST information for the interface. The following information describes the additional fields available in this window.
Auto-calculate Port Path Cost	Shows whether the path cost from the port to the root bridge is automatically determined by the speed of the interface (Enabled) or configured manually (Disabled).
Port ID	A unique value that is automatically generated based on the port priority value and the interface index.

**Table 237: Spanning Tree MST Port Configuration Fields (Continued)**

Field	Description
Port Up Time Since Counters Last Cleared	The amount of time that the port has been up since the counters were cleared.
Port Mode	The administrative mode of spanning tree on the port.
Designated Root	The bridge ID of the root bridge for the MST instance.
Designated Cost	The path cost offered to the LAN by the designated port.
Designated Bridge	The bridge ID of the bridge with the designated port.
Designated Port	The port ID of the designated port.
Loop Inconsistent State	Identifies whether the interface is currently in a loop inconsistent state. An interface transitions to a loop inconsistent state if loop guard is enabled and the port stops receiving BPDUs. In this state, the interface does not transmit frames.
Transitions Into LoopInconsistent State	The number of times this interface has transitioned into loop inconsistent state.
Transitions Out Of LoopInconsistent State	The number of times this interface has transitioned out of loop inconsistent state.

- If you make any configuration changes, click Submit to apply the new settings to the switch.
- Click Refresh to update the screen with most recent data.

### 5.25.6 Spanning Tree Statistics

Use the Spanning Tree Statistics page to view information about the number and type of bridge protocol data units (BPDUs) transmitted and received on each port.

To display the Spanning Tree Statistics page, click Switching > Spanning Tree > Statistics in the navigation menu.

**Figure 253: Spanning Tree Statistics**

Interface	STP BPDUs Rx	STP BPDUs Tx	RSTP BPDUs Rx	RSTP BPDUs Tx	MSTP BPDUs Rx	MSTP BPDUs Tx
1/0/1	0	0	0	0	0	256164
1/0/2	0	0	0	0	0	256164
1/0/3	0	0	0	0	0	0
1/0/4	0	0	0	0	0	256164
1/0/5	0	0	0	0	0	256164
1/0/6	0	0	0	0	0	256164
1/0/7	0	0	0	0	0	256164
1/0/8	0	0	0	0	0	256164
0/1/1	0	0	0	0	0	0
0/1/2	0	0	0	0	0	0

**Table 238: Spanning Tree Statistics Fields**

Field	Description
Interface	The port or link aggregation group (LAG) associated with the rest of the data in the row.
STP BPDUs Rx	The number of classic STP (IEEE 802.1d) BPDUs received by the interface.
STP BPDUs Tx	The number of classic STP BPDUs sent by the interface.
RSTP BPDUs Rx	The number of RSTP (IEEE 802.1w) BPDUs received by the interface.
RSTP BPDUs Tx	The number of RSTP BPDUs sent by the interface.
MSTP BPDUs Rx	The number of MSTP (IEEE 802.1s) BPDUs received by the interface.
MSTP BPDUs Tx	The number of MSTP BPDUs sent by the interface.

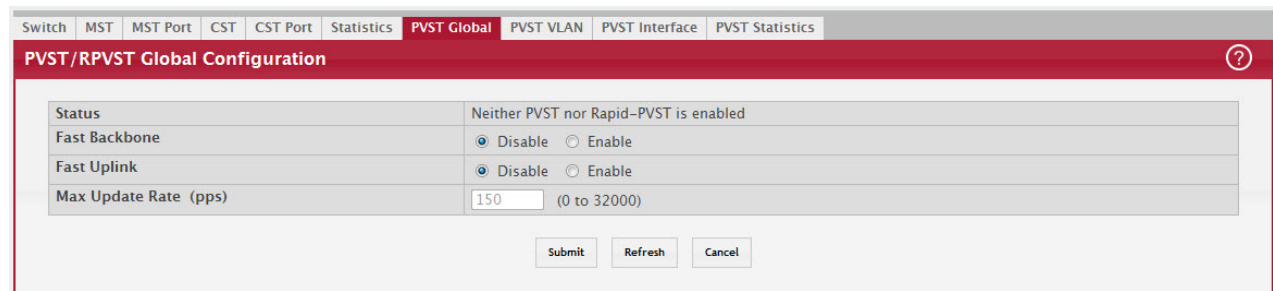
- Click Refresh to update the screen with most recent data.

### 5.25.7 PVST Global

Use this page to view and configure Per VLAN Spanning Tree Protocol (PVSTP)/Per VLAN Rapid Spanning Tree Protocol (PVRSTP) Global settings for the device.

To display the PVST Global page, click Switching > Spanning Tree > PVST Global in the navigation menu.

**Figure 254: PVST Global**



**Table 239: PVSTP/PVRSTP Global Fields**

Field	Description
Status	PVSTP/PVRSTP configuration operational mode.
Fast Backbone	Configures Fast Backbone mode. When enabled, the switch detects the indirect link failures and accelerates the spanning tree convergence.
Fast Uplink	Configures Fast Uplink mode.
Max Update Rate (pps)	Configures Fast Uplink's Maximum Update Rate.

- If you make any configuration changes, click Submit to apply the new settings to the switch.
- Click Refresh to update the screen with most recent data.

### 5.25.8 PVST VLAN

Use this page to view and configure Per VLAN Spanning Tree Protocol (PVSTP)/Per VLAN Rapid Spanning Tree Protocol (PVRSTP) VLAN settings for the device.

To display the PVST VLAN page, click Switching > Spanning Tree > PVST VLAN in the navigation menu.

Figure 255: PVST VLAN

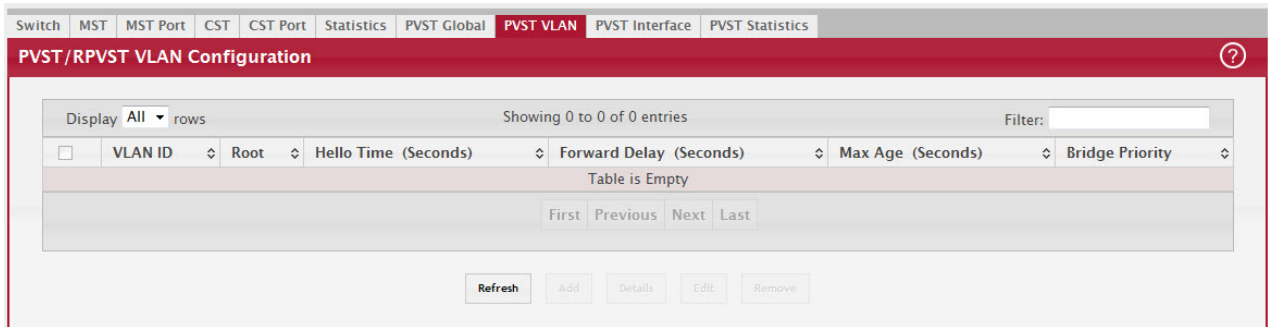


Table 240: PVSTP/PVRSTP VLAN Details Fields

Field	Description
VLAN ID	The unique VLAN identifier (VID).
Root	Configures the switch to become the root bridge or standby root bridge by modifying the bridge priority to a lower value to ensure the bridge is the root (or standby) bridge.
Hello Time (Seconds)	The interval between sending successive BDPUs. Configures the spanning tree hello time interval for the specified VLAN.
Forward Delay (Seconds)	Configures the spanning tree forward delay time for a specified VLAN. This interval is the time spent in listening and learning states before transitioning a port to the forwarding states.
Max Age (Seconds)	The maximum age time before a bridge port saves its configuration information.
Bridge Priority	Configures the bridge priority of a VLAN. The value will be automatically adjusted (rounded) as needed by the system. If the value configured is not among the specified values, then it will be rounded off to the nearest valid value.
To view details of any VLAN, the entry needs to be selected and Details button need to be pressed.	
Root ID	
Priority	The root ID priority for the specified VLAN.
Address	The root ID MAC address for the specified VLAN.
Cost	The root ID cost for the specified VLAN.
Port	The root ID port for the specified VLAN.
Hello Time (Seconds)	The root ID hello time for the specified VLAN.
Max Age (Seconds)	The maximum age for the specified VLAN.
Forward Delay (Seconds)	The root ID forward delay for the specified VLAN.
Bridge ID	
Priority	The bridge ID priority for the specified VLAN.
Address	The bridge ID MAC address for the specified VLAN.
Hello Time (Seconds)	The bridge ID hello time for the specified VLAN.
Max Age (Seconds)	The bridge ID maximum age for the specified VLAN.
Forward Delay (Seconds)	The bridge ID forward delay for the specified VLAN.
Aging Time (Seconds)	The bridge ID aging time for the specified VLAN.
Interface Details	
Interface	Interface which participates in the specified VLAN.
Role	The role of the interface.
Status	The status of the interface.

Table 240: PVSTP/PVRSTP VLAN Details Fields (Continued)

Field	Description
Cost	The cost value of the interface.
Prio.Nbr	The priority and neighbor of the interface.

Table 241: PVSTP/PVRSTP VLAN Add/Edit Fields

Field	Description
Root	Configures the switch to become the root bridge or standby root bridge by modifying the bridge priority to a lower value to ensure the bridge is the root (or standby) bridge.
Hello Time (Seconds)	The interval between sending successive BDPUs. Configures the spanning tree hello time interval for the specified VLAN.
Forward Delay (Seconds)	Configures the spanning tree forward delay time for a specified VLAN. This interval is the time spent in listening and learning states before transitioning a port to the forwarding states.
Max Age (Seconds)	The maximum age time before a bridge port saves its configuration information.
Bridge Priority	Configures the bridge priority of a VLAN. The value will be automatically adjusted (rounded) as needed by the system. If the value configured is not among the specified values, then it will be rounded off to the nearest valid value.

- Click Refresh to update the screen with most recent data.
- Click Add to add a new row to the VLAN configuration
- Select an entry and then click Edit to change the PVST configuration on the VLAN.
- Select an entry and then click Remove to remove the PVST row from the VLAN configuration.

## 5.25.9 PVST Interface

Use this page to view and configure Per VLAN Spanning Tree Protocol (PVSTP)/Per VLAN Rapid Spanning Tree Protocol (PVRSTP) Interface settings for the device.

To display the PVST Interface page, click Switching > Spanning Tree > PVST Interface in the navigation menu.

Figure 256: PVST Interface

The screenshot displays the 'PVST/RPVST Interface Configuration' page. At the top, there are navigation tabs: Switch, MST, MST Port, CST, CST Port, Statistics, PVST Global, PVST VLAN, **PVST Interface**, and PVST Statistics. Below the tabs, the page title 'PVST/RPVST Interface Configuration' is shown with a help icon. The main configuration area includes:

- Interface:** A dropdown menu set to '1/0/1'.
- Priority:** An input field with a refresh icon.
- Per VLAN Configuration:** A section with a 'Display All rows' dropdown, 'Showing 0 to 0 of 0 entries', and a 'Filter:' input field.
- Table:** A table with columns 'VLAN ID', 'Priority', and 'Cost'. The table is currently empty, with the text 'Table is Empty' centered below the columns. Navigation buttons 'First', 'Previous', 'Next', and 'Last' are located below the table.
- Buttons:** 'Refresh' and 'Edit' buttons are located at the bottom of the configuration area.

Table 242: PVSTP/PVRSTP Interface Fields

Field	Description
Interface	List of physical interfaces and LAGs.
Priority	The port priority configuration is used to allow the operator to select the relative importance of the port in the selection process for forwarding. Set this value to a lower number to prefer a port for forwarding of frames. This field is available for configuration only when PVSTP/PVRSTP is enabled.
Per VLAN Configuration	Configuration of each VLAN.
VLAN ID	The unique VLAN identifier (VID).
Priority	The per VLAN priority value configuration of the port is the priority used to allow the operator to select the relative importance of the port in the selection process for forwarding. Set this value to a lower number to prefer a port for forwarding of frames. This priority configuration is used when the port is configured as a point-to-point link type.
Cost	The path cost from the port to the root bridge.

Table 243: PVSTP/PVRSTP Interface Edit Fields

Field	Description
Interface	List of physical interfaces and LAGs.
Priority	The port priority configuration is used to allow the operator to select the relative importance of the port in the selection process for forwarding. Set this value to a lower number to prefer a port for forwarding of frames. This field is available for configuration only when PVSTP/PVRSTP is enabled.
Cost	The path cost from the port to the root bridge.

- Click Refresh to update the screen with most recent data.
- Select an entry and then click Edit to change the PVST interface configuration.

### 5.25.10 PVST Statistics

Use this page to view and configure Per VLAN Spanning Tree Protocol (PVSTP)/Per VLAN Rapid Spanning Tree Protocol (PVRSTP) Statistics settings for the device.

To display the PVST Statistics page, click Switching > Spanning Tree > PVST Statistics in the navigation menu.

Figure 257: PVST Statistics

Fast Backbone	
Transition via Fast Backbone	0
Inferior BPDUs Received	0
RLQ Request PDUs Received	0
RLQ Response PDUs Received	0
RLQ Request PDUs Sent	0
RLQ Response PDUs Sent	0

Fast Uplink	
Fast Uplink Transitions	0
Proxy Multicast Addresses Transmitted	0

Refresh



Table 244: PVSTP/PVRSTP Statistics Fields

Field	Description
Fast Backbone Transition via Fast Backbone	Number of fast backbone transitions.
Inferior BPDUs Received	Number of the received inferior BPDUs.
RLQ Request PDUs Received	Number of the received RLQ request PDUs.
RLQ Response PDUs Received	Number of the received RLQ response PDUs.
RLQ Request PDUs Sent	Number of the sent RLQ request PDUs.
RLQ Response PDUs Sent	Number of the sent RLQ response PDUs.
Fast Uplink	
Fast Uplink Transitions	Number of the fast uplink transitions.
Proxy Multicast Addresses Transmitted	Number of the transmitted proxy multicast addresses.

- Click Refresh to update the screen with most recent data.

## 5.26 Mapping 802.1p Priority

The IEEE 802.1p feature allows traffic prioritization at the MAC level. The switch can prioritize traffic based on the 802.1p tag attached to the L2 frame. Each port on the switch has multiple queues to give preference to certain packets over others based on the class of service (CoS) criteria you specify. When a packet is queued for transmission in a port, the rate at which it is serviced depends on how the queue is configured and possibly the amount of traffic present in the other queues of the port. If a delay is necessary, packets get held in the queue until the scheduler authorizes the queue for transmission.

Use the 802.1p Priority Mapping page in the Class of Service folder to assign 802.1p priority values to various traffic classes on one or more interfaces.

To display the page, click Switching > Class of Service > 802.1p in the navigation menu.

Figure 258: 802.1p Priority Mapping

Interface	Priority 0	Priority 1	Priority 2	Priority 3	Priority 4	Priority 5	Priority 6	Priority 7
Global	1	0	0	1	2	2	3	3
0/1	1	0	0	1	2	2	3	3
0/2	1	0	0	1	2	2	3	3
0/3	1	0	0	1	2	2	3	3
0/4	1	0	0	1	2	2	3	3
0/5	1	0	0	1	2	2	3	3
0/6	1	0	0	1	2	2	3	3
0/7	1	0	0	1	2	2	3	3
0/8	1	0	0	1	2	2	3	3
1/1	1	0	0	1	2	2	3	3

Table 245: 802.1p Priority Mapping

Field	Description
Interface	The interface associated with the rest of the data in the row. The Global entry represents the common settings for all interfaces, unless specifically overridden individually.
Priority	The heading row lists each 802.1p priority value (0–7), and the data in the table shows which traffic class is mapped to the priority value. Incoming frames containing the designated 802.1p priority value are mapped to the corresponding traffic class in the device.
802.1p Priority	The 802.1p priority value to be mapped.
Traffic Class	The internal traffic class to which the corresponding 802.1p priority value is mapped. The default value for each 802.1p priority level is displayed for reference.

## 5.27 Configuring Port Security

Port Security can be enabled on a per-port basis. When a port is locked, only packets with allowable source MAC addresses can be forwarded. All other packets are discarded. A MAC address can be defined as allowable by one of two methods: dynamically or statically.

Both methods are used concurrently when a port is locked.

### NOTICE

Dynamic locking implements a first arrival mechanism for Port Security. You specify how many addresses can be learned on the locked port. If the limit has not been reached, a packet with an unknown source MAC address is learned and forwarded normally. Once the limit is reached, no more addresses are learned on the port. Any packets with source MAC addresses that were not already learned are discarded. Note that you can effectively disable dynamic locking by setting the number of allowable dynamic entries to zero.

Static locking allows you to specify a list of MAC addresses that are allowed on a port. The behavior of packets is the same as for dynamic locking: only packets with an allowable source MAC address can be forwarded.


To see the MAC addresses learned on a specific port, see [Section 4.8: "Configuring and Searching the Forwarding Database"](#).

Disabled ports can only be activated from the Configuring Ports page.

### 5.27.1 Port Security Administration

Use the Port Security Administration page to enable or disable the port security feature on your switch. To access the Port Security Administration page, click Switching > Port Security > Global in the navigation menu.

Figure 259: Port Security Administration



Select Enable or Disable from the Port Security Mode list and click Submit.

### 5.27.2 Port Security Interface Configuration

Use this page to configure the port security feature on a selected interface. To access the Port Security Interface Configuration page, click Switching > Port Security > Interface in the navigation menu.

Figure 260: Port Security Interface Configuration

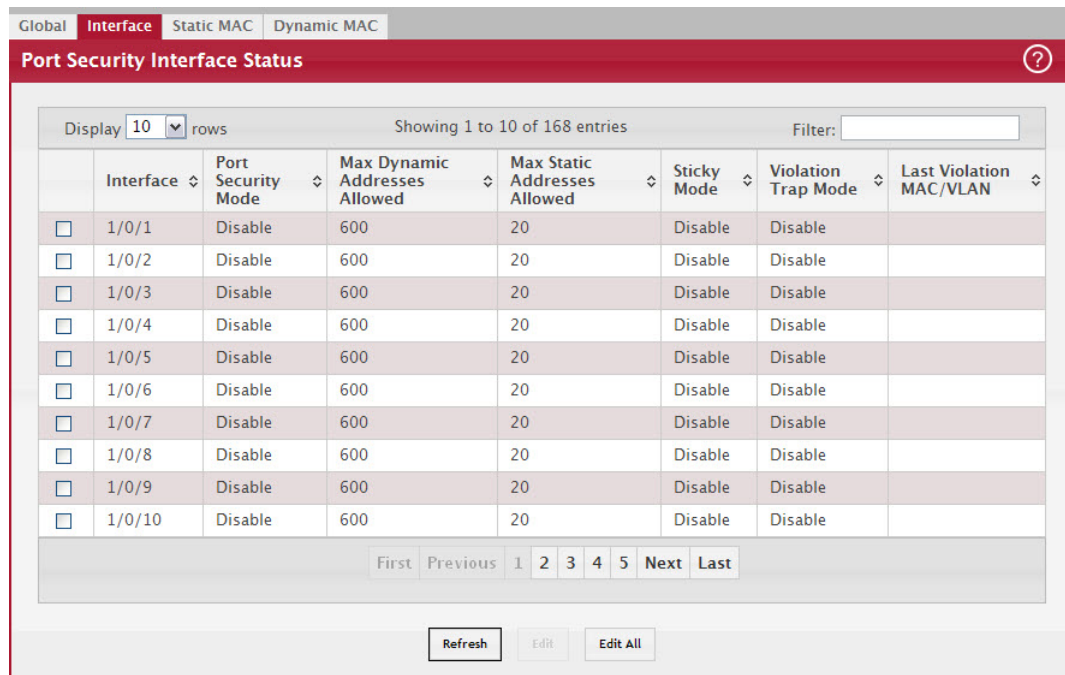


Table 246: Port Security Interface Configuration Fields

Field	Description
Interface	Select the physical interface or the LAG on which to configure port security information.
Port Security	Determines whether port security is enabled. The default mode is Disable. <ul style="list-style-type: none"> <li>• Enable: Locks the port so that only packets with allowable source MAC addresses can be forwarded. All other packets are discarded.</li> <li>• Disable: The port is not locked, so no port security restrictions are applied.</li> </ul>

**Table 246: Port Security Interface Configuration Fields (Continued)**

Field	Description
Maximum Number of Dynamically Learned MAC Addresses Allowed	Sets the maximum number of dynamically learned MAC addresses on the selected interface. Once the limit is reached, no more addresses are learned on the port. Any packets with source MAC addresses that were not already learned are discarded. You can effectively disable dynamic locking by setting the number of allowable dynamic entries to zero.
Maximum Number of Statically Locked MAC Addresses Allowed	Sets the maximum number of statically locked MAC addresses on the selected interface.
Add a Static MAC Address	Adds a MAC address to the list of statically locked MAC addresses for the selected interface. Only packets with an allowable source MAC address can be forwarded.
VLAN ID	Adds a corresponding VLAN ID for the MAC Address being added to the list of statically locked MAC addresses for the selected interface.
Enable Violation Traps	Enables or disables the sending of new violation traps designating when a packet with a disallowed MAC address is received on a locked port. Value is No by default.
Convert dynamically learned address to static locked	When you click Move, all the dynamically learned entries on this interface are added to the static MAC address list for this interface. After moving them, you can view them in the Port Security Static page.

If you make any changes to the page, click Submit to apply the new settings to the system.

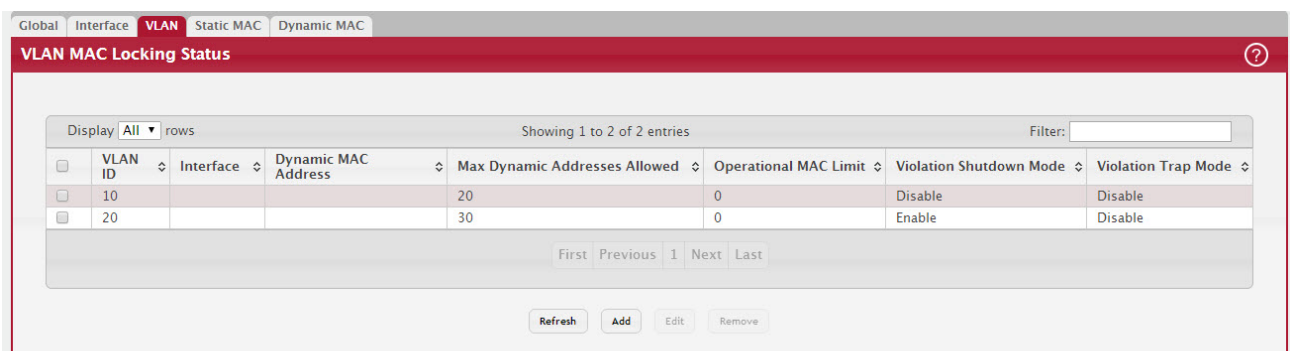
### 5.27.3 VLAN MAC Locking

Use this page to configure VLAN MAC Locking. VLAN MAC locking allows you to secure the network by locking down allowable MAC addresses on a given VLAN. Packets with a matching source MAC address can be forwarded normally. All other packets will be discarded. VLAN MAC locking will lock the dynamic MAC entries.

If VLAN and port MAC locking are enabled, VLAN MAC locking will be given precedence over port MAC locking.

To access the VLAN MAC Locking Status page, click Switching > Port Security > VLAN in the navigation menu.

**Figure 261: VLAN MAC Locking Status Configuration**



**Table 247: Port Security Interface Configuration Fields**

Field	Description
VLAN ID	The VLAN ID specified in the Ethernet frame received by the interface.
Interface	The interface associated with the rest of the data in the row.
Dynamic MAC Address	The MAC address that was learned on the device. An address is dynamically learned when a frame arrives on the interface and the source MAC address in the frame is added to the MAC address table.

**Table 247: Port Security Interface Configuration Fields (Continued)**

Field	Description
Max Dynamic Addresses Allowed	The number of source MAC addresses that can be dynamically learned on an interface. If an interface reaches the configured limit, any other addresses beyond that limit are not learned, and the frames are discarded. Frames with a source MAC address that has already been learned will be forwarded. A dynamically-learned MAC address is removed from the MAC address table if the entry ages out, the link goes down, or the system resets. Note that the behavior of a dynamically-learned address changes if the sticky mode for the interface is enabled or the address is converted to a static MAC address.
Operational MAC Limit	The number of source MAC addresses that are dynamically currently reached to that of Maximum Configured MAC Limit.
Violation Shutdown Mode	After MAC limit has reached, action will shut down the ports participating in the VLAN.
Violation Trap Mode	After MAC limit has reached, a log message will be generated with violation MAC address details.

To configure The VLAN MAC Locking, use the following buttons to perform the tasks:

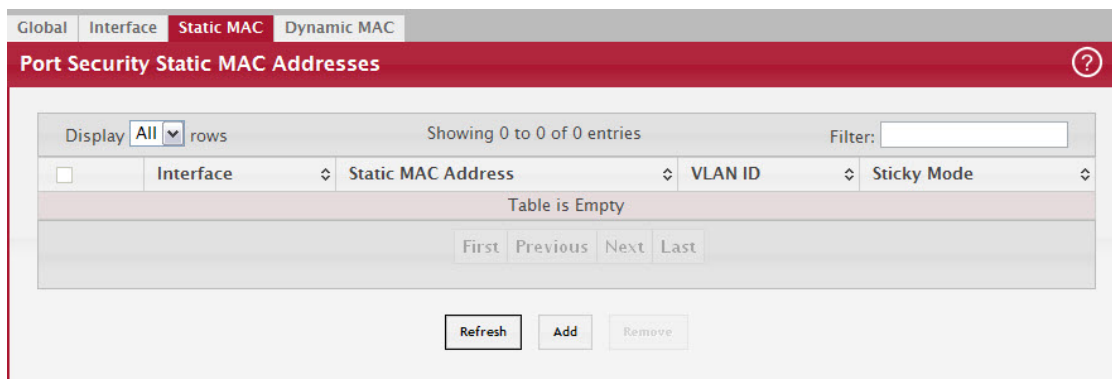
- Use Submit to enable or disable VLAN MAC Locking Admin Mode.
- Use Add to configure VLAN MAC Locking.
- Use Edit to modify configuration parameters of VLAN MAC Locking.
- Use Remove to remove configured VLANs.

### 5.27.4 Port Security Statically Configured MAC Addresses

Use the Port Security Statically Configured MAC Addresses page to view static MAC addresses configured on an interface. From this page, you can delete statically configured MAC addresses.

To access the Port Security Static page, click Switching > Port Security > Static MAC in the navigation menu.

**Figure 262: Port Security Statically Configured MAC Addresses**



**Table 248: Port Security Statically Configured MAC Address Fields**

Field	Description
Interface	Select the physical interface or the LAG on which to view the dynamically learned MAC addresses.
MAC Address	This column lists the static MAC addresses, if any, configured on the selected port.
VLAN ID	Displays the VLAN ID corresponding to the statically configured MAC address.
Delete a static MAC Address	Enter the address of the statically configured MAC address to delete. All MAC addresses that are available to be deleted appear in the MAC Address – VLAN ID table.
VLAN ID	Enter the VLAN ID that corresponds to the statically configured MAC address to delete.

After you enter the MAC address and VLAN ID of the statically configured MAC address to delete, click Submit to remove the MAC address from the port and apply the new settings to the system. The screen refreshes, and the MAC address no longer appears in the table on the page.

## 5.27.5 Port Security Dynamically Learned MAC Addresses

Use the Port Security Dynamically Learned MAC Addresses page to view a table with the dynamically learned MAC addresses on an interface. With dynamic locking, MAC addresses are learned on a first arrival basis. You specify how many addresses can be learned on the locked port.

To access the Port Security Dynamic page, click Switching > Port Security > Dynamic MAC in the navigation menu.

Figure 263: Port Security Dynamic MAC Address

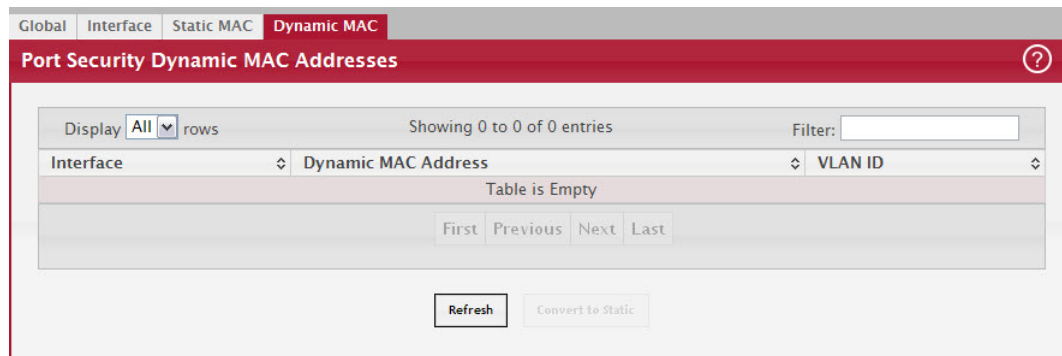


Table 249: Port Security Dynamic Fields

Field	Description
Interface	Select the physical interface or the LAG on which to view the dynamically learned MAC addresses.
MAC Address	This column lists the dynamically learned MAC addresses, if any, on the selected port.
VLAN ID	Displays the VLAN ID corresponding to the dynamically learned MAC address.

## 5.28 Managing LLDP

The IEEE 802.1AB defined standard, Link Layer Discovery Protocol (LLDP), allows stations residing on an 802 LAN to advertise major capabilities and physical descriptions. This information is viewed by a network manager to identify system topology and detect bad configurations on the LAN.

LLDP is a one-way protocol; there are no request/response sequences. Information is advertised by stations implementing the transmit function, and is received and processed by stations implementing the receive function. The transmit and receive functions can be enabled/disabled separately per port. By default, both transmit and receive are disabled on all ports. The application is responsible for starting each transmit and receive state machine appropriately, based on the configured status and operational state of the port.

FASTPATH allows LLDP to have multiple LLDP neighbors per interface. The number of such neighbors is limited by the memory constraints. A product-specific constant defines the maximum number of neighbors supported by the switch. There is no restriction on the number of neighbors supported on a per LLDP port. If all the remote entries on the switch are filled up, the new neighbors are ignored. In case of multiple VOIP devices on a single interface, the 802.1ab component sends the Voice VLAN configuration to all the VoIP devices.

### 5.28.1 Global Configuration

Use the LLDP Global Configuration page to specify LLDP parameters that are applied to the switch.

To display the LLDP Global Configuration page, click Switching > LLDP > Global in the navigation menu.

Figure 264: LLDP Global Configuration

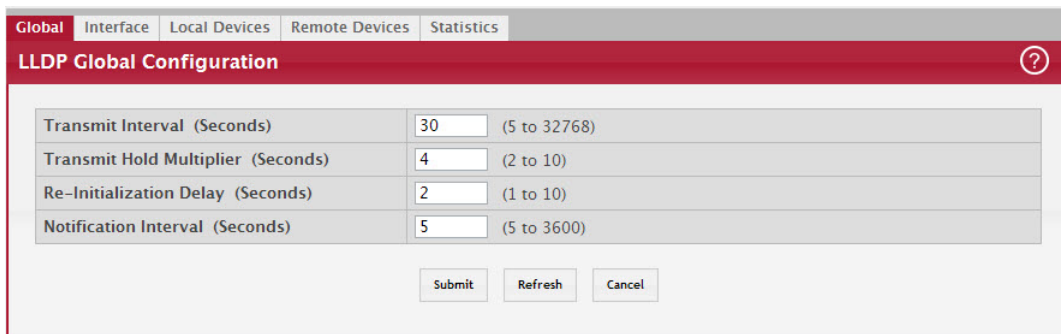


Table 250: LLDP Global Configuration Fields

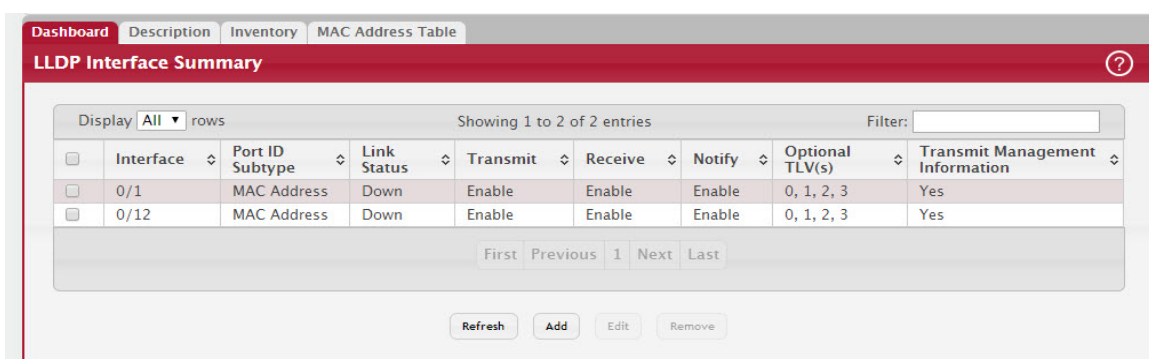
Field	Description
Transmit Interval	Specifies the interval at which LLDP frames are transmitted. The default is 30 seconds, and the valid range is 1-32768 seconds.
Transmit Hold Multiplier	Specifies multiplier on the transmit interval to assign to TTL. The default is 4, and the range is 2-10.
Re-Initialization Delay	Specifies the delay before a re-initialization. The default is 2 seconds, and the range is 1-10 seconds.
Notification Interval	Limits the transmission of notifications. The default is 5 seconds, and the range is 5-3600 seconds.

If you make any changes to the page, click Submit to apply the new settings to the system.

### 5.28.2 LLDP Interface Configuration

Use the LLDP Interface Configuration page to specify LLDP parameters that are applied to a specific interface. To display the LLDP Interface Configuration page, click Switching > LLDP > Interface in the navigation menu.

Figure 265: LLDP Interface Summary



**NOTICE**

When adding or editing LLDP settings on an interface, select the appropriate check box to enable a feature, or clear the check box to disable a feature:

Table 251: LLDP Interface Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that have at least one LLDP setting enabled appear in the table. In the Add LLDP Interface window, use this field to select the interface with the LLDP settings to configure. In the Edit LLDP Interface window, this field identifies the interface that is being configured.
Port ID Subtype	The LLDP Port ID subtype of the interface, which is either MAC Address or Interface Name.
Link Status	The link status of the interface, which is either Up or Down. An interface that is down does not forward traffic.
Transmit	The LLDP advertise (transmit) mode on the interface. If the transmit mode is enabled, the interface sends LLDP Data Units (LLDPDUs) that advertise the mandatory TLVs and any optional TLVs that are enabled.
Receive	The LLDP receive mode on the interface. If the receive mode is enabled, the device can receive LLDPDUs from other devices.
Notify	The LLDP remote data change notification status on the interface. If the notify mode is enabled, the interface sends SNMP notifications when a link partner device is added or removed.
Optional TLV(s)	Select each check box next to the type-length value (TLV) information to transmit. Choices include: <ul style="list-style-type: none"> <li>• System Name. To include system name TLV in LLDP frames. To configure the System Name, see <a href="#">Section 4.7.1: "System Description"</a>.</li> <li>• System Description. To include system description TLV in LLDP frames.</li> <li>• System Capabilities. To include system capability TLV in LLDP frames.</li> <li>• Port Description. To include port description TLV in LLDP frames. To configure the Port Description, see <a href="#">Section 4.13.2: "Port Description"</a>.</li> </ul>
Transmit Management Information	Select the check box to enable the transmission of management address instance. Clear the check box to disable management information transmission. The default is disabled.

Use the buttons to perform the following tasks:

- To configure LLDP settings on an interface that does not have any LLDP settings enabled, click Add.
- To change the LLDP settings for an interface in the table, select the entry to update and click Edit. If you clear (disable) all LLDP settings, the entry is removed from the table.
- To clear (disable) all LLDP settings from one or more interfaces, select each entry to clear and click Remove.

After you click Add or Edit, a window opens and allows you to configure the LLDP settings for an interface. The following information describes the additional fields that appear in the windows used for adding or editing per-interface LLDP settings.



Figure 266: LLDP Interface Add

In addition to some of the fields that Table 251, “LLDP Interface Summary Fields,” on page 298 describes, Table 252, “LLDP Interface Add Fields,” on page 299 shows the additional fields available on the Add LLDP Interface window.

Table 252: LLDP Interface Add Fields

Field	Description
System Name	Select this option to include the user-configured system name in the LLDPDU the interface transmits. The system name is configured on the System Description page and is the SNMP server name for the device.
System Description	Select this option to include a description of the device in the LLDPDU the interface transmits. The description includes information about the product model and platform.
System Capabilities	Select this option to advertise the primary function(s) of the device in the LLDPDU the interface transmits.
Port Description	Select this option to include the user-configured port description in the LLDPDU the interface transmits.

If you make any changes to the page, click Submit to apply the new settings to the system.

### 5.28.3 Local Devices

Use the LLDP Local Device page to view information about all interfaces on the device that are enabled to transmit LLDP information.

To display the LLDP Local Device Summary page, click Switching > LLDP > Local Devices in the navigation menu.

Figure 267: LLDP Local Devices

Table 253: LLDP Local Devices Columns

Field	Description
Interface	The interface associated with the rest of the LLDP - 802.1AB data in the row. When viewing the details for an interface, this field identifies the interface that is being viewed.
Port ID	The port identifier, which is the physical address associated with the interface.
Port Description	A description of the port. An administrator can configure this information on the Port Description page.

Click Refresh to update the information on the screen with the most current data.

After you click Details, a window opens and displays additional information about the data the interface transmits in its LLDPDUs. The following information describes the additional fields that appear in the LLDP Local Device Information window.

Table 254: LLDP Local Devices Details

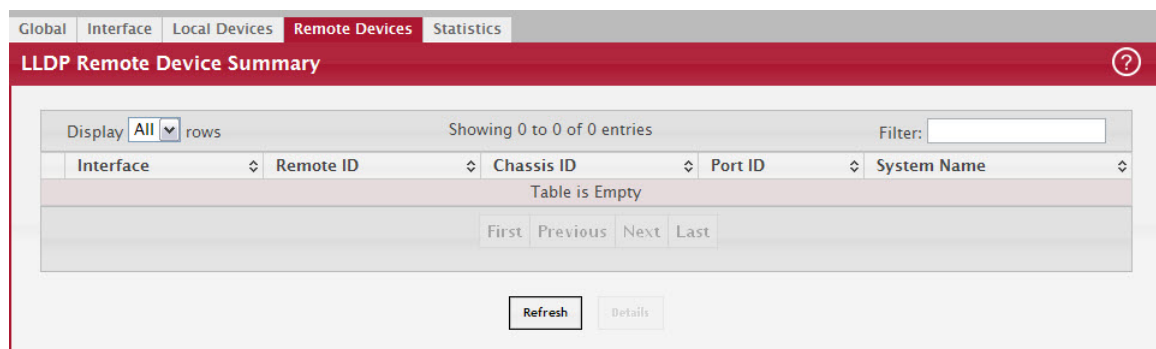
Field	Description
Chassis ID Subtype	The type of information used to identify the device in the Chassis ID field.
Chassis ID	The hardware platform identifier for the device.
Port ID Subtype	The type of information used to identify the interface in the Port ID field.
System Name	The user-configured system name for the device. The system name is configured on the System Description page and is the SNMP server name for the device.
System Description	The device description, which includes information about the product model and platform.
System Capabilities Supported	The primary function(s) the device supports.
System Capabilities Enabled	The primary function(s) the device supports that are enabled.
Management Address	The physical address associated with the management interface of the device.
Management Address Type	The protocol type or standard associated with the management address.

## 5.28.4 Remote Devices

Use the LLDP Remote Device Summary page to view information about all interfaces on the device that are enabled to transmit LLDP information.

To display the LLDP Remote Device Summary page, click Switching > LLDP > Remote Devices in the navigation menu.

Figure 268: LLDP Remote Device Summary



**Table 255: LLDP Remote Device Summary Columns**

Field	Description
Interface	The local interface that is enabled to receive LLDPDUs from remote devices.
Remote ID	The client identifier assigned to the remote system that sent the LLDPDU.
Chassis ID	The information the remote device sent as the Chassis ID TVL. This identifies the hardware platform for the remote system.
Port ID	The port on the remote system that transmitted the LLDP data.
System Name	The system name configured on the remote device.

Click Refresh to update the information on the screen with the most current data.

After you click Details, a window opens and displays additional information. If the interface has received LLDP data from a remote device, the window displays detailed information about the device. If the interface has not received any LLDPDUs from remote devices, the window displays a message indicating that no LLDP data has been received. The following information describes the additional fields that appear in the LLDP Remote Device Information window when LLDP data has been received on the selected interface.

**Table 256: LLDP Remote Device Summary Columns**

Field	Description
Chassis ID Subtype	The type of information used to identify the device in the Chassis ID field.
Port ID Subtype	The type of information used to identify the interface in the Port ID field.
System Description	The device description, which includes information about the product model and platform.
Port Description	The description of the port on the remote device that transmitted the LLDP data.
System Capabilities Supported	The primary function(s) the remote system supports. The possible capabilities include Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station.
System Capabilities Enabled	The primary function(s) of the remote system that are both supported and enabled. The possible capabilities include Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station.
Time To Live	The number of seconds the local device should consider the LLDP data it received from the remote system to be valid.

### 5.28.5 Statistics

Use the LLDP Statistics page to view the global and interface LLDP statistics.

To display the LLDP Statistics page, click Switching > LLDP > Statistics in the navigation menu.

Figure 269: LLDP Statistics

Table 257: LLDP Statistics Fields

Field	Description
<b>System-wide Statistics</b>	
Last Update	Displays the time when an entry was created, modified, or deleted in the tables associated with the remote systems.
Total Inserts	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been inserted into the tables associated with the remote systems.
Total Deletes	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from the tables associated with the remote systems.
Total Drops	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) could not be entered into tables associated with the remote systems because of insufficient resources.
Total Ageouts	Displays the number of times a complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from tables associated with the remote systems because the information timelines interval has expired.
<b>Port Statistics</b>	
Interface	Identifies the interfaces.
Transmit Total	Displays the total number of LLDP frames transmitted by the LLDP agent on the corresponding port.
Receive Total	Displays the total number of valid LLDP frames received by the LLDP agent on the corresponding port, while the LLDP agent is enabled.
Discards	Displays the number of LLDP TLVs discarded for any reason by the LLDP agent on the corresponding port.
Errors	Displays the number of invalid LLDP frames received by the LLDP agent on the corresponding port, while the LLDP agent is enabled.
Ageouts	Displays the number of age-outs that occurred on a given port. An age-out is the number of times the complete set of information advertised by a particular MAC Service Access Point (MSAP) has been deleted from tables associated with remote entries because the information timeliness interval had expired.
TLV Discards	Displays the number of LLDP TLVs (Type, Length, Value sets) discarded for any reason by the LLDP agent on the corresponding port.

Table 257: LLDP Statistics Fields (Continued)

Field	Description
TLV Unknowns	Displays the number of LLDP TLVs received on the local ports which were not recognized by the LLDP agent on the corresponding port.
TLV MED	Displays the total number of LLDP-MED TLVs received on the local ports.
TLV 802.1	Displays the total number of LLDP TLVs received on the local ports which are of type 802.1.
TLV 802.3	Displays the total number of LLDP TLVs received on the local ports which are of type 802.3.

- Click Refresh to update the page with the most current information.
- Click Clear to clear the LLDP statistics of all the interfaces.

## 5.28.6 LLDP-MED

The Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) is an enhancement to LLDP that features:

- Auto-discovery of LAN policies (such as VLAN, Layer 2 Priority and DiffServ settings), enabling plug and play networking.
- Device location discovery for creation of location databases.
- Extended and automated power management of Power over Ethernet endpoints.
- Inventory management, enabling network administrators to track their network devices and determine their characteristics (manufacturer, software and hardware versions, serial/asset number).

### 5.28.6.1 LLDP-MED Global Configuration

Use this page to set global parameters for LLDP-MED operation. To display this page, click Switching > LLDP-MED > Global in the navigation menu.

Figure 270: LLDP-MED Global Configuration

Table 258: LLDP Global Configuration Fields

Field	Description
Fast Start Repeat Count	Specifies the number of LLDP PDUs that will be transmitted when the protocol is enabled. The range is from (1 to 10). The default value is 3.
Device Class	Specifies local device's MED Classification. The following three represent the actual endpoints: <ul style="list-style-type: none"> <li>• Class I Generic [IP Communication Controller etc.]</li> <li>• Class II Media [Conference Bridge etc.]</li> <li>• Class III Communication [IP Telephone etc.]</li> </ul> The fourth device is Network Connectivity Device, which is typically a LAN switch/router, IEEE 802.1 bridge, IEEE 802.11 wireless access point, etc.

Click Submit to update the switch. The changes take effect but will not be retained across a power cycle unless a save is performed.

### 5.28.6.2 LLDP-MED Interface Configuration

Use this page to enable LLDP-MED mode on an interface and to configure its properties. To configure the settings for one or more interfaces, select each entry to modify and click Edit. The same LLDP-MED settings are applied to all selected interfaces.

To display this page, click Switching > LLDP-MED > Interface in the navigation menu.

Figure 271: LLDP-MED Interface Summary

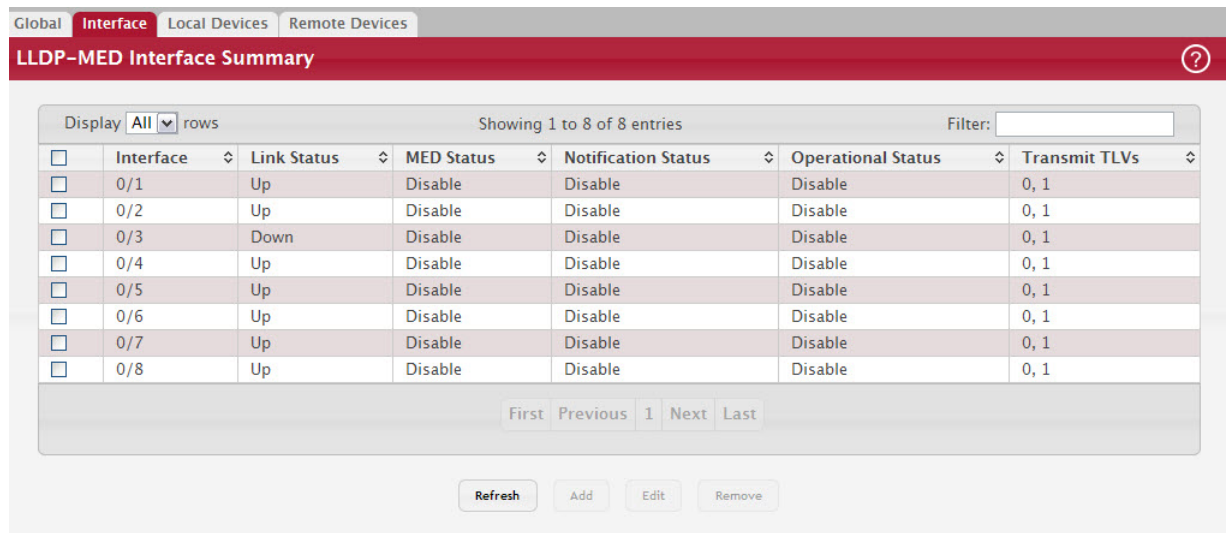


Table 259: LLDP-MED Interface Configuration Fields

Field	Description
Interface	Selects the port that you want to configure LLDP-MED-802.1AB on. You can select All to configure all interfaces on the DUT with the same properties. The Interface Configuration page will not be able to display the summary of 'All' interfaces. The summary of individual interfaces is visible from the Interface Configuration page. The Interface Configuration page for the 'All' option will always display the LLDP-MED mode and notification mode as 'disabled' and check boxes for 'Transmit TLVs' will always be unchecked.
Link Status	The link status of the interface, which is either Up or Down. An interface that is down does not forward traffic.
MED Status/LLDP-MED Mode	The administrative status of LLDP-MED on the interface. When LLDP-MED is enabled, the transmit and receive function of LLDP is effectively enabled on the interface.
Notification Status/Configuration Notification Mode	Indicates whether LLDP-MED topology change notifications are enabled or disabled on the interface.
Operational Status	Indicates whether the interface will transmit TLVs.
Transmit TLVs	The LLDP-MED TLV(s) that the interface transmits: <ul style="list-style-type: none"> <li>• MED Capabilities: 0</li> <li>• Network Policy: 1</li> </ul>

Click Submit to send the updated configuration to the switch. These changes take effect immediately but will not be retained across a power cycle unless a save is performed.

### 5.28.6.3 LLDP Local Device Information

This page displays information on LLDP-MED information advertised on the selected local interface. To display this page, click Switching > LLDP-MED > Local Devices in the navigation menu.

Figure 272: LLDP-MED Local Device Summary

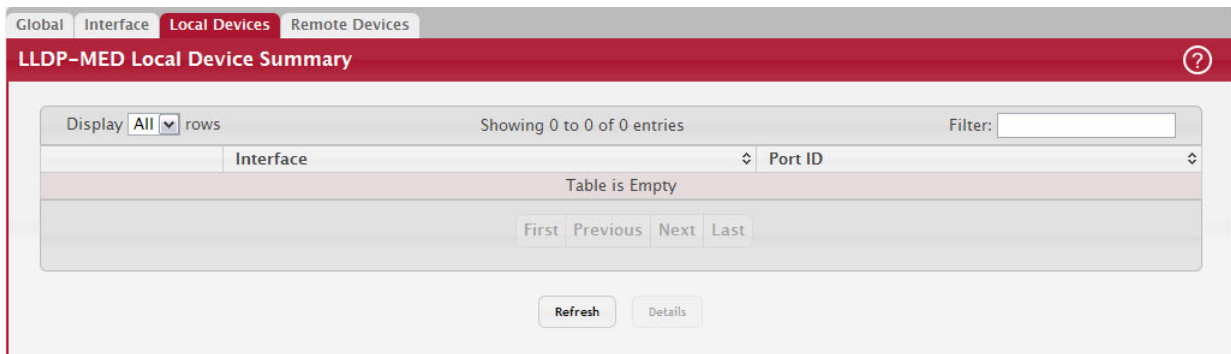


Table 260: LLDP-MED Local Device Information Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When viewing LLDP-MED details for an interface, this field identifies the interface that is being viewed.
Port ID	The MAC address of the interface. This is the MAC address that is advertised in LLDP-MED PDUs. After you click <i>Details</i> , a window opens and shows detailed information about the LLDP-MED information the selected interface transmits. The following information describes the additional fields that appear in the LLDP-MED Local Device Information window.
Network Policy Information The information in this table identifies the data transmitted in the Network Policy TLVs.	
Media Application Type	The media application type transmitted in the TLV. The application types are unknown, voicesignaling, guestvoice, guestvoicesignalling, softphonevoice, videoconferencing, streamingvideo, vidoesignalling. Each application type that is transmitted has the VLAN ID, priority, DSCP, tagged bit status and unknown bit status. A port may transmit one or many such application types. This information is displayed only when a network policy TLV has been transmitted.
VLAN ID	The VLAN ID associated with a particular policy type.
Priority	The user priority associated with a particular policy type.
DSCP	The DSCP value associated with a particular policy type.
Unknown Bit Status	The unknown bit associated with a particular policy type.
Tagged Bit Status	Identifies whether the network policy is defined for tagged or untagged VLANs.
Location Information	
Sub Type	The type of location information: <ul style="list-style-type: none"> <li>Coordinate Based – The location map coordinates (latitude, longitude and altitude) of the device.</li> <li>Civic Address – The civic or street address location of the device.</li> <li>ELIN – The Emergency Call Service (ECS) Emergency Location Identification Number (ELIN) of the device.</li> </ul>
Information	This column displays the information related to the coordinates, civic address, and ELIN for the device.

Click Refresh to update the page with the latest information from the router.

### 5.28.6.4 LLDP-MED Remote Device Information

This page displays information about the remote devices the local system has learned about through the LLDP-MED data units received on its interfaces. Information is available about remote devices only if an interface receives an LLDP-MED data unit from a device. To view additional information about a remote device, select the interface that received the LLDP-MED data and click Details. The information below is organized according to the order in which the fields appear in the LLDP-MED Remote Device Information window.

To display this page, click Switching > LLDP-MED > Remote Devices in the navigation menu.

Figure 273: LLDP Remote Device Summary

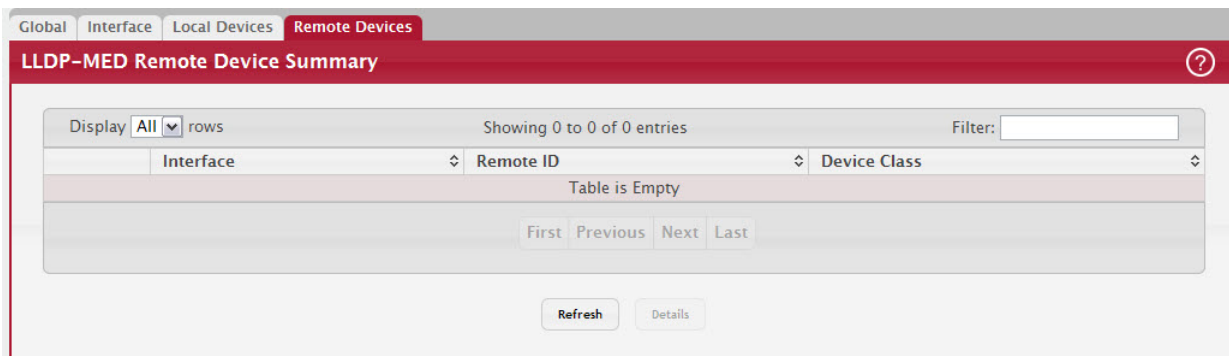


Table 261: LLDP-MED Remote Device Information Fields

Field	Description
Interface	The local interface that has received LLDP-MED data units from remote devices.
Remote ID	The client identifier assigned to the remote system that sent the LLDP-MED data unit.
Capability Information	
Supported Capabilities	The supported capabilities that were received in the MED TLV on this interface.
Enabled Capabilities	The supported capabilities on the remote device that are also enabled.
Device Class	The MED Classification advertised by the TLV from the remote device. The following three classifications represent the actual endpoints: <ul style="list-style-type: none"> <li>• Class I Generic (for example, IP Communication Controller)</li> <li>• Class II Media (for example, Conference Bridge)</li> <li>• Class III Communication (for example, IP Telephone)</li> </ul> The fourth device is Network Connectivity Device, which is typically a device such as a LAN switch or router, IEEE 802.1 bridge, or IEEE 802.11 wireless access point.
Network Policy Information	
This section describes the information in the network policy TLVs received in the LLDP-MED frames on this interface.	
Media Application Type	The media application type received in the TLV from the remote device. The application types are unknown, voicesignaling, guestvoice, guestvoicesignalling, softphonevoice, videoconferencing, streamingvideo, vidoesignalling. Each application type that is transmitted has the VLAN ID, priority, DSCP, tagged bit status and unknown bit status. The port on the remote device may transmit one or many such application types. This information is displayed only when a network policy TLV has been received.
VLAN ID	The VLAN ID associated with a particular policy type.
Priority	The user priority associated with a particular policy type.
DSCP	The DSCP value associated with a particular policy type.
Unknown Bit Status	The unknown bit associated with a particular policy type.
Tagged Bit Status	Identifies whether the network policy is defined for tagged or untagged VLANs.



**Table 261: LLDP-MED Remote Device Information Fields (Continued)**

Field	Description
<b>Inventory Information</b>	
This section describes the information in the inventory TLVs received in the LLDP-MED frames on this interface.	
Hardware Revision	The hardware version advertised by the remote device.
Firmware Revision	The firmware version advertised by the remote device.
Software Revision	The software version advertised by the remote device.
Serial Number	The serial number advertised by the remote device.
Manufacturer Name	The name of the system manufacturer advertised by the remote device.
Model Name	The name of the system model advertised by the remote device.
Asset ID	The system asset ID advertised by the remote device.
<b>Location Information</b>	
This section describes the information in the location TLVs received in the LLDP-MED frames on this interface.	
Sub Type	The type of location information advertised by the remote device.
Information	The text description of the location information included in the subtype.
Extended PoE	Indicates whether the remote device is advertised as a PoE device.
Device Type	If the remote device is a PoE device, this field identifies the PoE device type of the remote device connected to this port.

Click Refresh to update the page with the latest information from the router.

## 5.29 Loop Protection

L2 Loop Protection feature allows loop detection in downstream switches that do not run spanning tree. It can optionally disable the associated port on loop detection.

The Loop Protection feature is not intended for ports that serve as uplinks between spanning tree aware switches. Loop Protection feature is designed for unmanaged switches which drop spanning Tree BPDUs. This feature detects physical and logical loops between Ethernet ports on a device. The feature needs to be enabled globally before enabling it at the interface level for the system policy filter to be installed.

### 5.29.1 Loop Protection Configuration

Use this page to configure the Loop Protection feature. Loops on a network consume resources and can impact network performance. When loop protection is enabled on the switch and on one or more interfaces (ports and trunks), the interfaces send loop protection protocol data units (PDUs) to the multicast destination address 01:80:C2:00:00:08. When an interface receives a loop protection PDU, it compares the source MAC address with its own. If the MAC addresses match, a loop is detected and a configured action is taken, which may include shutting down the port for a specified period. An interface can also be configured to receive and take action in response to loop protection PDUs, but not to send out the PDUs itself.

To display this page, click Switching > Loop Protection > Configuration in the navigation menu.

Figure 274: Loop Protection Configuration

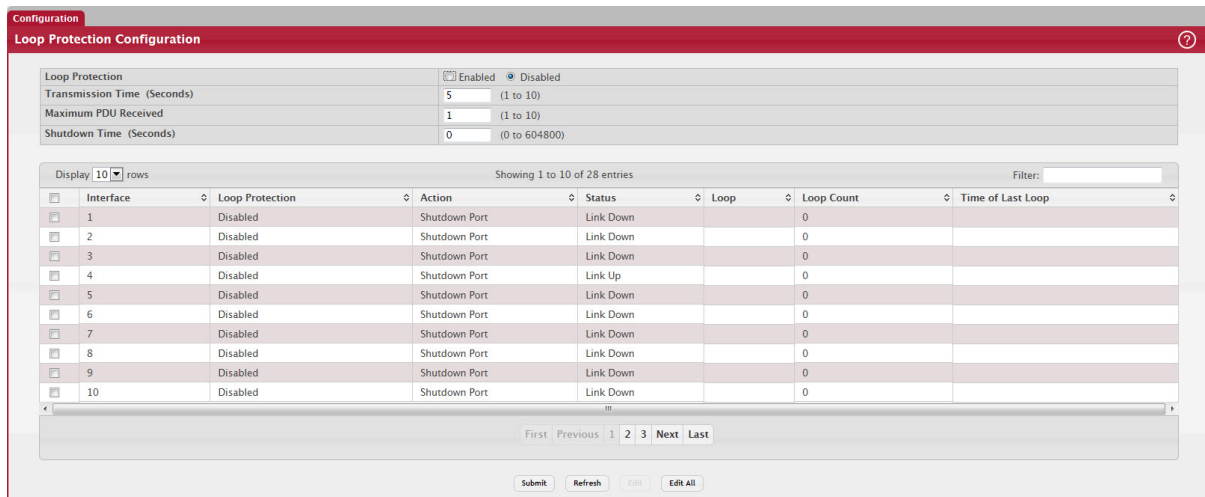


Table 262: Loop Protection Configuration Fields

Field	Description
Loop Protection	Enables or disables the loop protection feature globally on the switch. <b>Note:</b> The loop protection feature is not supported on dynamic trunks. The loop protection feature will be automatically disabled if it was previously enabled on a static trunk that is now configured as dynamic.
Transmission Time (Seconds)	The interval at which the switch sends loop protection PDUs on interfaces that are enabled to send them.
Maximum PDU Received	This configures the count of loop protection packets received by the switch after which the interface will be err-disabled.
Interface	The port or trunk ID.
Edit Loop Protection Port Configuration	Select an interface to and click Edit to edit the Loop Protection Port Configuration. Click Edit All to apply the same configuration to all interfaces.  <div data-bbox="587 1339 1268 1709" data-label="Image"> </div>
Action	The action to be taken when a loop is detected on the port: <ul style="list-style-type: none"> <li>Shutdown Port: Shut down the port for the configured Transmission Time.</li> <li>Shutdown Port and Log: Shut down the port for the configured Transmission Time and send a message to the system log.</li> <li>Log Only: Send a message to the system log but do not shut down the port.</li> </ul>
Status	The current status of the interface. Link Up indicates the interface is operating normally. Link Down indicates that the port has been shut down due to the detection of a loop.

Table 262: Loop Protection Configuration Fields

Field	Description
Loop	Indicates whether a loop is currently detected on the interface. If blank, then no loop is detected.
Loop Count	The number of times a loop has occurred on the interface.
Time of Last Loop	The date and time the most recent loop was detected.

Click Submit to updated the switch. The changes take effect but will not be retained across a power cycle unless a save is performed.

## 5.30 IEEE 802.1ag Connectivity Fault Management (CFM)

With the emergence of Ethernet as a Metropolitan and Wide-Area Networking technology, different operators often work together to provide end-to-end services to enterprise customers. This has driven the need of a new set of OAM (Operations, Administration, and Maintenance) Protocols.

IEEE 802.1ag, also known as Service-Level Connectivity Fault Management (CFM), is the OAM protocol provision for end-to-end service-layer instances in carrier networks. CFM provides mechanisms to support the administrator in performing connectivity checks, fault detection, fault verification and isolation, and fault notification per service in the network domain of interest. Unlike Ethernet OAM (IEEE 802.3ah), where the faults are detected and notified on a single point-to-point IEEE Std. 802.3 LAN, Dot1ag addresses fault diagnosis at the service layer across networks comprising multiple LANs, including LANs other than 802.3 media.

You can configure the switch to use IEEE 802.1ag (Dot1ag) CFM functionality. Dot1ag enables customers to perform the following in a network domain of interest:

- Connectivity checks
- Fault detection
- Fault verification and isolation
- Fault notification per service

---

The Dot1ag feature is available in the optional Metro package:

### NOTICE

### 5.30.1 Dot1ag Global Configuration

Use the Dot1ag Configuration page to set the administrative mode of IEEE 802.1ag CFM and to configure the archive hold time.

To display this page, click Switching > Dot1ag > Configuration in the navigation menu.

Figure 275: Dot1ag Global Configuration

Table 263: Dot1ag Global Configuration

Field	Definition
CFM Admin Mode	Enables or Disables the CFM feature on the switch.
Archive Hold Time	The time interval in seconds after which inactive Remote Maintenance End Points (RMEPs) are removed.

- If you make any configuration changes, click Submit to apply the new settings to the switch.
- Click Cancel to cancel the change.
- Click Refresh to update the screen with most recent data.

### 5.30.2 Dot1ag Maintenance Domain (MD) Configuration

Use this page to view and configure Dot1ag maintenance domains for the device.

To display the page, click Switching > Ddot1ag > MD Configuration.

Figure 276: Dot1ag MD Configuration

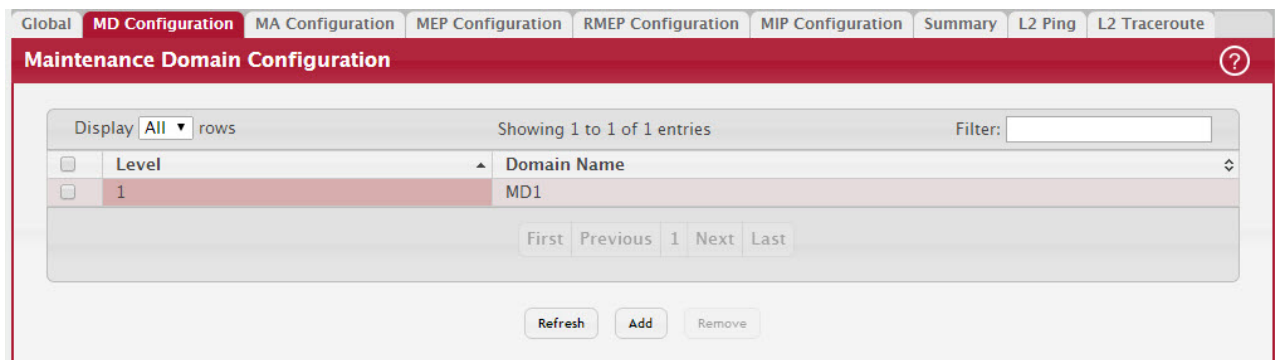


Table 264: MD Configuration

Field	Definition
Level	The Level at which the Maintenance Domain is to be created.
Domain Name	The maintenance domain name configured for the level. The domain name may contain only alphanumeric characters including hyphens, underscores, and quotation marks (-, _, ').

Use the buttons to perform the following:

- To add a maintenance domain name, click Add. The Maintenance Domain Add window displays. Select the MD level, and specify a name for the MD.

Figure 277: Add Dot1ag MD

- To remove one or more configured maintenance domain name(s), select each entry to delete and click Remove. You must confirm the action before the entry is deleted.
- Click Cancel to abandon the changes.

### 5.30.3 Dot1ag Maintenance Association (MA) Configuration

Use this page to view and configure Dot1ag maintenance associations for the device.

To display the page, click Switching > Ddot1ag > MA Configuration.

Figure 278: Dot1ag MA Configuration

Table 265: MA Configuration

Field	Definition
Domain Name—Level	The maintenance domain name configured for the level, followed by the level at which the maintenance domain was created.
Primary VLAN ID	The primary VLAN ID associated with the MA service.
MA Name	The maintenance association service name. The service name may only contain alphanumeric characters, including -, _, ' '.
CCM Interval	The time between CCM frames transmission used by all MEPs in the given Maintenance Association.

Use the buttons to perform the following:

- To add a maintenance domain association, click Add. The Maintenance Association Add window displays.

Figure 279: Add Dot1ag MA

- To change the settings for an existing maintenance association, select each entry to configure, click Edit, and update the fields.
- To remove one or more configured maintenance domain name(s), select each entry to delete and click Remove. You must confirm the action before the entry is deleted.
- Click Cancel to abandon the changes.

### 5.30.4 Dot1ag Maintenance Association End-Point (MEP) Configuration

Use this page to view and configure Dot1ag maintenance association end-points for the device. To display the page, click Switching > Ddot1ag > MEP Configuration.

Figure 280: Dot1ag MEP Configuration

Table 266: MEP Configuration

Field	Definition
Domain Name—Level	The maintenance domain name configured for the level, followed by the Level at which the maintenance domain was created.
Primary VLAN ID	The primary VLAN ID associated with the MA service.
MEP ID	An integer ID for the MEP. <b>Note:</b> The MEP ID must be an integer that is unique (cannot be the same as the identifier for other MEPs in the same MA,
Interface	The physical port or port-channel interface to which the MEP is attached.
Direction	The direction (UP/DOWN) in which the MEP faces on the bridge port.
MEP Active	A boolean value indicating the administrative state of the MEP. True indicates that the MEP is configured to function normally, and False that it is configured to cease functioning.
CCI Enabled	If Set to true, the MEP will generate CCM Messages.

Use the buttons to perform the following:

- To add a MEP, click Add. The Maintenance Association End-Point Add window displays.

**Figure 281: Add Dot1ag MEP**

Domain Name - Level	MD1 - 1
Primary VLAN ID	100
MEP ID	(1 to 8191)
Interface	1/0/1
Direction	<input checked="" type="radio"/> Down <input type="radio"/> Up
MEP Active	<input type="radio"/> True <input checked="" type="radio"/> False
CCI Enabled	<input type="radio"/> True <input checked="" type="radio"/> False

- To change the settings for an existing maintenance association, select the MEP to configure, click Edit, and update the fields.
- To remove one or more configured MEPs, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.
- To view additional information about a configured MEP, select the entry, and click Details.

**Figure 282: Dot1ag MEP Details**

Domain Name - Level	MD1 - 1
Primary VLAN ID	100
MEP ID	10
Interface	1/0/1
Direction	Down
MEP Active	True
CCI Enabled	True
Operational Status	False
MAC Address	00:10:18:7F:FC:F1
Defects	

- Click Cancel to abandon the changes.

### 5.30.5 Dot1ag Remote Maintenance Association End-Point (RMEP) Configuration

Use this page to view and configure Dot1ag RMEPs for the device.

To display the page, click Switching > Dot1ag > RMEP Configuration.

Figure 283: Dot1ag RMEP Configuration

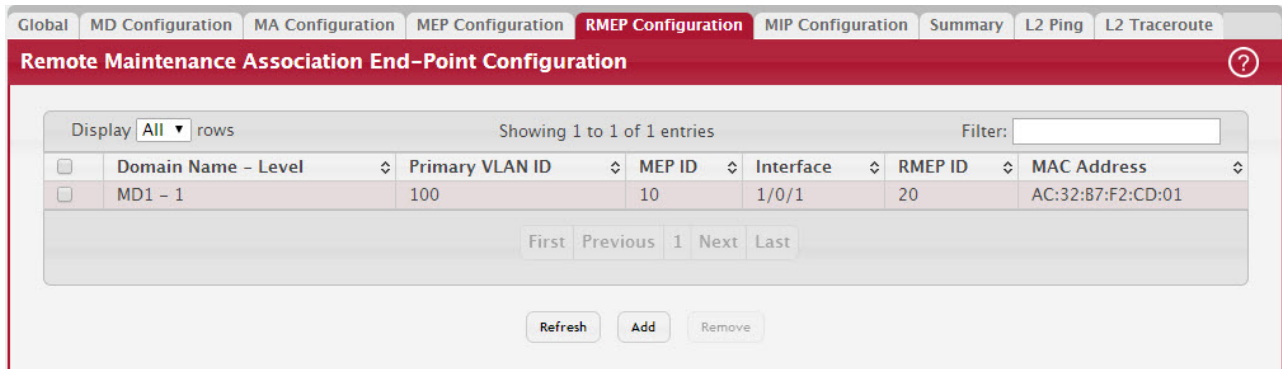


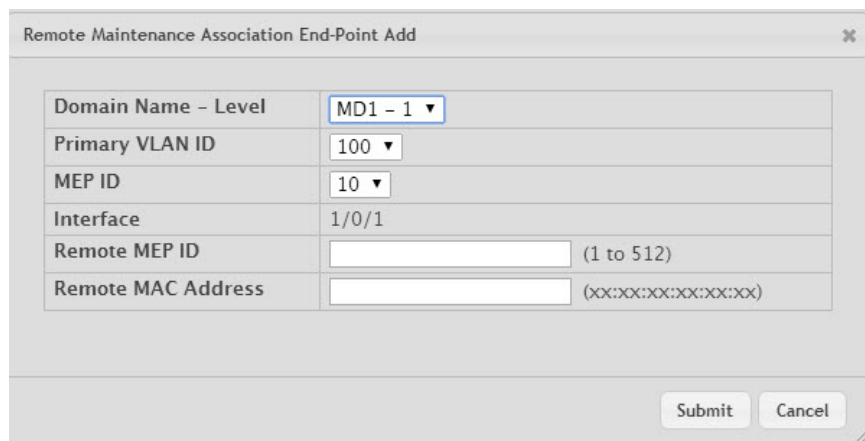
Table 267: RMEP Configuration

Field	Definition
Domain Name—Level	The maintenance domain name configured for the level, followed by the level at which the maintenance domain was created.
Primary VLAN ID	The primary VLAN ID associated with the MA service.
MEP ID	An integer ID for the MEP, which is unique within the MA.
Interface	The physical port or port-channel interface to which the MEP is attached.
RMEP ID	The Maintenance Association End Point Identifier of a remote MEP.
MAC Address	MAC Address of the Remote MEP.

Use the buttons to perform the following:

- To add an RMEP, click Add. The Maintenance Association End-Point Add window displays.

Figure 284: Add Dot1ag RMEP



- To remove one or more configured RMEPs, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.
- Click Refresh to update the screen with the current information.



### 5.30.6 Dot1ag Maintenance Intermediate Point (MIP) Configuration

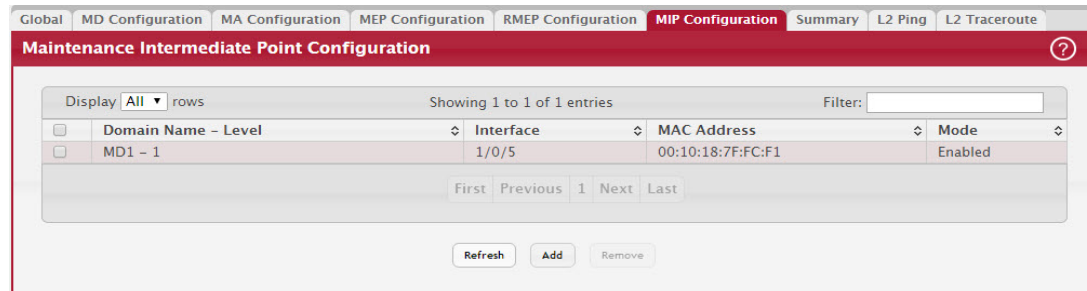
Use this page to view and configure Dot1ag maintenance intermediate points (MIPs) for the device.

Only enabled MIP devices are displayed in the configuration table.

**NOTICE**

To display the page, click Switching > Dot1ag > MIP Configuration.

**Figure 285: Dot1ag MIP Configuration**



**Table 268: MIP Configuration**

Field	Definition
Domain Name—Level	The maintenance domain name configured for the level, followed by the level at which the maintenance domain was created.
Interface	The physical port or a port-channel interface to which the MIP is attached.
MAC Address	The MAC Address of the MIP.
Mode	Indicates whether a MIP is configured on an interface for a given maintenance domain.

Use the buttons to perform the following:

- To add a MIP, click Add. The Maintenance Intermediate Point Add window displays.

**Figure 286: Add Dot1ag MIP**



- To remove one or more configured MIPs, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.
- Click Refresh to update the screen with the current information.
- Click Cancel to abandon any in-progress changes.

### 5.30.7 Dot1ag Remote Maintenance Association End-Point Summary

Use this page to view a summary of Dot1ag RMEP end-points for the device.

To display the page, click Switching > Dot1ag > Summary.

Figure 287: Dot1ag RMEP Summary

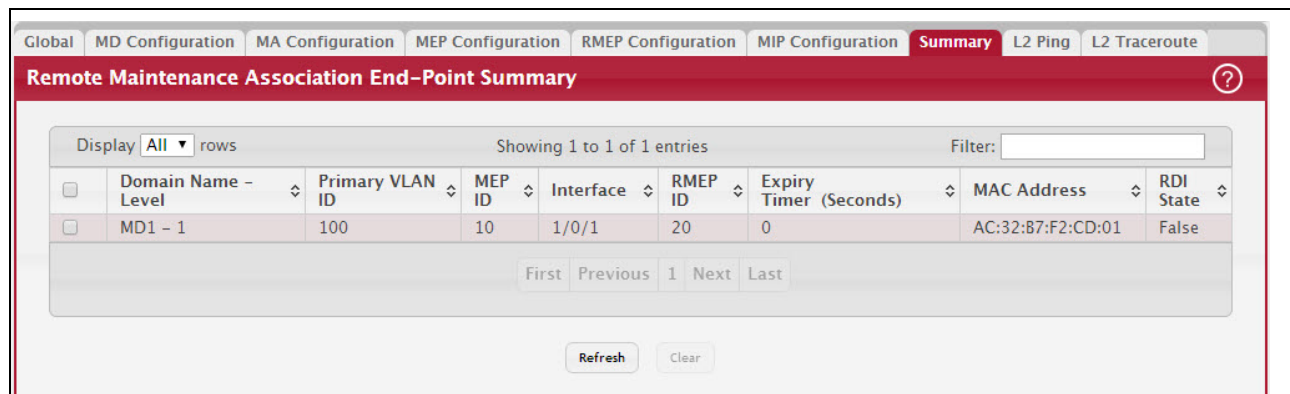


Table 269: RMEP Configuration

Field	Definition
Domain Name—Level	The maintenance domain name configured for the level, followed by the level at which the maintenance domain was created.
Primary VLAN ID	The primary VLAN ID associated with the MA service.
MEP ID	An integer ID for the MEP, which is unique within the MA.
Interface	The physical port or port-channel interface to which the MEP is attached.
RMEP ID	The Maintenance Association End Point Identifier of a remote MEP.
Expiry Timer	Time allowed for the last received Continuity Check Message (CCM) entry to expire, on a given RMEP.
MAC Address	MAC Address of the Remote MEP.
RDI State	The state of the Remote Defect Indication (RDI) bit in the last received CCM. Possible value are True and False. If the value is True, the RDI bit is set and indicates that a defect condition has been encountered.

Use the buttons to perform the following:

- Click Refresh to update the screen with the current information.
- To clear one or more configured/learned remote maintenance association end-point(s), select each entry to delete and click Clear. You must confirm the action before the entry is deleted.

### 5.30.8 Dot1ag L2 Ping

Use this page to initiate a Dot1ag L2 ping. To display the page, click Switching > Dot1ag > L2 Ping.

Figure 288: Dot1ag L2 Ping

The screenshot shows the 'L2 Ping' configuration page. At the top, there are navigation tabs: Global, MD Configuration, MA Configuration, MEP Configuration, RMEP Configuration, MIP Configuration, Summary, L2 Ping (selected), and L2 Traceroute. Below the tabs is a red header with 'L2 Ping' and a help icon. The main area contains a form with the following fields:

- Domain Name - Level: MD1 - 1
- Primary VLAN ID: 100
- MEP ID: 10
- Interface: 1/0/1
- Target Identifier:  MAC Address  MEP ID
- Target MAC Address:  (xx:xx:xx:xx:xx:xx)
- Target MEP ID:  (1 to 8191)
- Count: 1 (1 to 255)
- Results:

A 'Submit' button is located at the bottom center of the form.

Table 270: Dot1ag L2 Ping

Field	Definition
Domain Name—Level	The maintenance domain name configured for the level, followed by the level at which the maintenance domain was created.
Primary VLAN ID	The primary VLAN ID associated with the MA service.
MEP ID	An integer ID for the MEP, which is unique within the MA.
Interface	The physical port or port-channel interface to which the MEP is attached.
Target Identifier	Specifies whether to use the MEP ID or MAC address of the remote MEP in the Link Trace Message.
Target MAC Address	The target MAC address field of another MEP in the same MA to which the LBM is to be transmitted. This MAC address should be unicast. This field is visible only if the MAC address is chosen as the Target Identifier
Target MEP ID	The maintenance association end-point ID of the MEP in the same MA to which the LBM is to be sent. This field is visible only if the MEP ID is chosen as the Target Identifier.
Count	The number of loopback messages to be transmitted.
Results	The response from the remote MEP for the transmitted loopback messages. The response will be available and displayed 5 seconds after transmission.

Click Submit to transmit the loopback messages to the switch.

### 5.30.9 Dot1ag L2 Traceroute

Use this page to configure and transmit a Dot1ag L2 traceroute. To display the page, click Switching > Dot1ag > L2 Traceroute.

Figure 289: Dot1ag L2 Traceroute

Table 271: Dot1ag L2 Traceroute Fields

Field	Definition
Domain Name—Level	The maintenance domain name configured for the level, followed by the level at which the maintenance domain was created.
Primary VLAN ID	The primary VLAN ID associated with the MA service.
MEP ID	An integer ID for the MEP, which is unique within the MA.
Interface	The physical port or port-channel interface to which the MEP is attached.
Target Identifier	Specifies whether to use the MEP ID or MAC address of the remote MEP in the loopback message.
Target MAC Address	The target MAC address field of another MEP in the same MA to which the LBM is to be transmitted. This MAC address should be unicast. This field is visible only if the MAC address is chosen as the Target Identifier.
Target MEP ID	The maintenance association end-point ID of the MEP in the same MA to which the LBM is to be sent. This field is visible only if the MEP ID is chosen as the Target Identifier.
TTL	The number of hops remaining to the LTM. Each LinkTrace responder that handles the LTM decrements this count by 1. The default value, if not specified, is 64. If the LTM TTL is 0 or 1, the LTM is not forwarded to the next hop, and if 0, no LTR is generated.
Results	The results of the Traceroute are available on the L2 Traceroute Cache page.

Click Submit to transmit the LinkTrace Messages to the switch. You are redirected to the L2 Traceroute Cache page.

### 5.30.10 Dot1ag L2 Traceroute Cache

Use this page to view the dot1ag L2 traceroute cache response. To display the page, click Switching > Dot1ag > L2 Traceroute Cache.

Figure 290: Dot1ag L2 Traceroute Cache

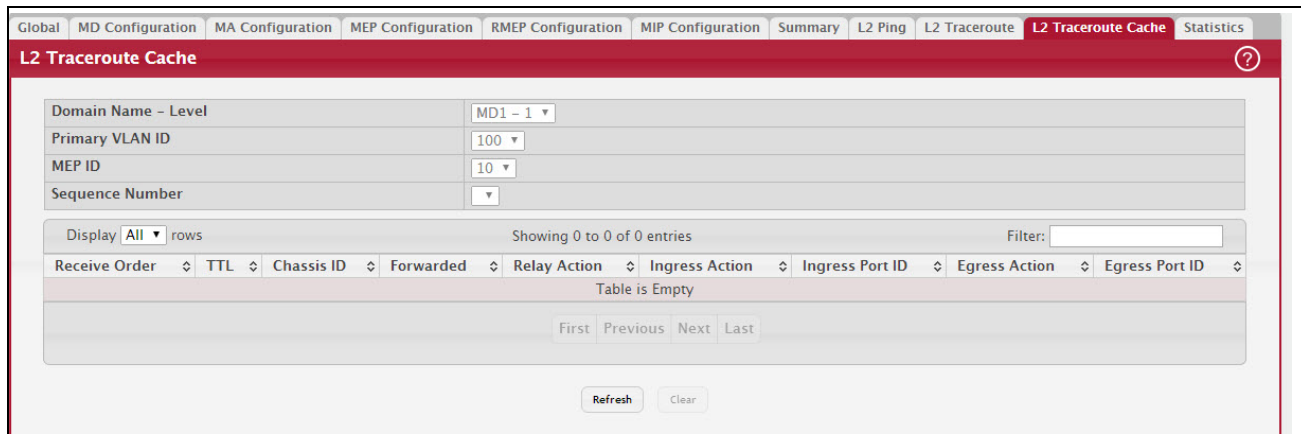


Table 272: Dot1ag L2 Traceroute Cache Fields

Field	Description
Domain Name - Level	The maintenance domain name configured for the level, followed by the level at which the maintenance domain was created.
Primary VLAN ID	The primary VLAN ID associated with the MA service.
MEP ID	An integer ID for the MEP, which is unique within the MA.
Sequence Number	A sequence number that identifies the LTM transaction.
Traceroute Cache Summary	
Receive Order	An index that is used to distinguish multiple LinkTrace responses that have the same LinkTrace Transaction Identifier field value.
TTL	The TTL field value for a returned LinkTrace response.
Chassis ID	The ID of the bridge on which the responding MP resides.
Forwarded	Indicates whether a LinkTrace message was forwarded by the responding MP, as returned in the 'FwdYes' flag of the flags field. Possible values are Forwarded and Not Forwarded.
Relay Action	The Relay Action field value for a returned LinkTrace response. Possible values are: relayHIT, relayFDB, or relayMPDB.
Ingress Action	Indicates how the data frame targeted by the LinkTrace message would be received on the receiving MP. The possible values are: ingOK, ingDown, ingBlocked, or ingVid.
Ingress Port ID	ID of the port on which the LinkTrace message has arrived.
Egress Action	Indicates how the data frame targeted by the LinkTrace message would be passed through egress bridge port. Possible values are: egrOK, egrDown, egrBlocked, or egrVid.
Egress Port ID	ID of the port to which the LinkTrace message was forwarded.

Use the buttons to perform the following:

- Click Refresh to update the screen with the current information.
- To clear one or more configured/learned remote maintenance association end-point(s), select each entry to delete and click Clear. You must confirm the action before the entry is deleted.

### 5.30.11 Dot1ag Statistics

Use this page to view dot1ag statistics. To display the page, click Switching > Dot1ag > Statistics.

Figure 291: Dot1ag Statistics

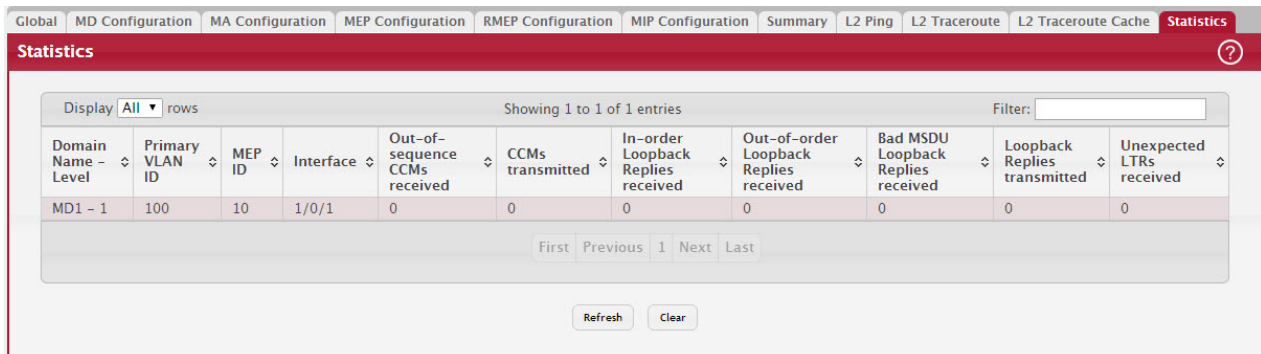


Table 273: Dot1ag Statistics Fields

Field	Definition
Domain Name—Level	The maintenance domain name configured for the level, followed by the level at which the maintenance domain was created.
Primary VLAN ID	The primary VLAN ID associated with the MA service.
MEP ID	An integer ID for the MEP, which is unique within the MA.
Interface	The physical port or port-channel interface to which the MEP is attached.
Out-of-sequence CCMs Received	Number of out-of-sequence CCMs received from all remote MEPs.
CCMs Transmitted	Number of CCMs transmitted.
In-order Loopback Replies Received	Number of valid in-order loopback replies received.
Out-of-order Loopback Replies Received	Number of valid out-of-order loopback replies received.
Bad MSDU Loopback Replies Received	Number of LBRs received with a mac_service_data_unit that did not match the mac_service_data_unit of the corresponding LBM. <b>Note:</b> The OpCode is not included for the match.
Loopback Replies Transmitted	Number of loopback replies transmitted.
Unexpected LTRs Received	Number of unexpected LTRs received.

Click Refresh to update the page with the most current information.

### 5.31 IEEE 802.3ah Ethernet in the First Mile

FASTPATH supports the Ethernet in the First Mile (EFM) protocol as part of the Operations, Administration, and Maintenance (OAM) set of tools and features that help manage and maintain the network and to troubleshoot network issues. EFM-OAM is specified by the IEEE 802.3ah specification.

#### 5.31.1 Dot3ah Configuration

Use the Dot3ah Configuration page to configure the Dot3ah (EFM OAM) settings on one or more interfaces. Dot3ah can be enabled on ports that are connected to service provider networks (WAN port).

OAM provides the mechanisms for monitoring link operation such as remote fault indication and remote loopback control. It helps the network administrators to monitor, test and troubleshoot the OAM-enabled links on the network. The hardware identifies all incoming OAM PDUs on ports where Dot3ah is enabled.

To display this page, click Switching > Dot3ah > Configuration in the navigation menu.

Figure 292: Dot3ah Configuration

Interface	Ethernet OAM	Dot3ah Mode	Max PDU Rate	Min PDU Rate	Org Specific TLV	Peer Timeout
1/0/1	Disabled	Passive	1	1	Enabled	10
1/0/2	Disabled	Passive	1	1	Enabled	10
1/0/3	Disabled	Passive	1	1	Enabled	10
1/0/4	Disabled	Passive	1	1	Enabled	10
1/0/5	Disabled	Passive	1	1	Enabled	10
1/0/6	Disabled	Passive	1	1	Enabled	10
1/0/7	Disabled	Passive	1	1	Enabled	10
1/0/8	Disabled	Passive	1	1	Enabled	10
1/0/9	Disabled	Passive	1	1	Enabled	10
1/0/10	Disabled	Passive	1	1	Enabled	10

Table 274: Dot3ah Configuration

Field	Description
Interface	The port associated with the rest of the data in the row.
Ethernet OAM	The status of Ethernet OAM protocol on the interface.
Dot3ah Mode	The EFM-OAM mode on the interface. An active-mode device can exert more control on its peer than a passive-mode device. For example, an active-mode OAM entity can put a passive-mode OAM entity into loop-back mode, but not vice versa
Max PDU Rate	Indicates the maximum transmission rate (pdu_timer) in seconds when one OAM PDU is sent per second.
Min PDU Rate	Indicates the minimum transmission rate (pdu_timer) in seconds when one OAM PDU is sent per second.
Org Specific TLV	Indicates whether the interface support interpreting Organization Specific Information TLVs.
Peer Timeout	Indicates the link lost timer value in seconds. If any OAM PDUs are not received from the remote DTE, the local client executing FAULT state of Discovery state machine.

- To configure one or more interface, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.
- To apply the same settings to all interfaces, click Edit All.
- If you make any configuration changes in the Edit Dot3ah Configuration window, click Submit to apply the new settings to the interface(s), or click Cancel to abandon the changes and close the window.
- Click Refresh to update the screen with most recent data.

### 5.31.2 Dot3ah Link Monitor Configuration

Use the Dot3ah Link Monitor Configuration page to configure an interface for link monitoring. The link monitoring mechanism is provided to support event notifications that include information regarding the transmission errors as detected at the MAC sub-layer in OAM PDUs in link events.

To display this page, click Switching > Dot3ah > Link Monitor Configuration in the navigation menu.

Figure 293: Dot3ah Link Monitor Configuration

Interface	Ethernet OAM	Link Monitor Capability	Link Monitor Mode	Link Monitor Frame Error	Link Monitor Frame Period Error	Link Monitor Frame Seconds Error	Link Monitor Last TLV Send	Link Monitor Last TLV Received
<input type="checkbox"/> 1/0/1	Disabled	Disabled	Disabled	High 1   Low 1   Window 1	High 1   Low 1   Window 10000	High 1   Low 1   Window 60	Frame - 0 days 0:0:0 Frame Period - 0 days 0:0:0 Frame Sec - 0 days 0:0:0	Frame - 0 days 0:0:0 Frame Period - 0 days 0:0:0 Frame Sec - 0 days 0:0:0
<input checked="" type="checkbox"/> 1/0/2	Enabled	Enabled	Disabled	High 1   Low 1   Window 1	High 1   Low 1   Window 10000	High 1   Low 1   Window 60	Frame - 0 days 0:0:0 Frame Period - 0 days 0:0:0 Frame Sec - 0 days 0:0:0	Frame - 0 days 0:0:0 Frame Period - 0 days 0:0:0 Frame Sec - 0 days 0:0:0
<input type="checkbox"/> 1/0/3	Disabled	Disabled	Disabled	High 1   Low 1   Window 1	High 1   Low 1   Window 10000	High 1   Low 1   Window 60	Frame - 0 days 0:0:0 Frame Period - 0 days 0:0:0 Frame Sec - 0 days 0:0:0	Frame - 0 days 0:0:0 Frame Period - 0 days 0:0:0 Frame Sec - 0 days 0:0:0
<input type="checkbox"/> 1/0/4	Disabled	Disabled	Disabled	High 1   Low 1   Window 1	High 1   Low 1   Window 10000	High 1   Low 1   Window 60	Frame - 0 days 0:0:0 Frame Period - 0 days 0:0:0 Frame Sec - 0 days 0:0:0	Frame - 0 days 0:0:0 Frame Period - 0 days 0:0:0 Frame Sec - 0 days 0:0:0
<input type="checkbox"/> 1/0/5	Disabled	Disabled	Disabled	High 1   Low 1   Window 1	High 1   Low 1   Window 10000	High 1   Low 1   Window 60	Frame - 0 days 0:0:0 Frame Period - 0 days 0:0:0 Frame Sec - 0 days 0:0:0	Frame - 0 days 0:0:0 Frame Period - 0 days 0:0:0 Frame Sec - 0 days 0:0:0

Table 275: Dot3ah Link Monitor Configuration

Field	Description
Interface	The port associated with the rest of the data in the row.
Ethernet OAM	The status of Ethernet OAM protocol on the interface.
Link Monitor Capability	Indicates whether link monitoring is supported.
Link Monitor Mode	This indicates the status of Ethernet OAM (Dot3ah) link monitoring.
Link Monitor Frame Error	The high and low error threshold values in number of symbols and the window used to calculate and generate a trap if the errors are crossing the thresholds within the window time frame.
Link Monitor Frame Period Error	The high and low thresholds for error frames that trigger an error-frame link event and the window used to calculate and generate a trap if the errors are crossing the thresholds within the window time frame.
Link Monitor Frame Seconds Error	The high and low thresholds for the error-frame seconds that triggers an error-frame seconds link event and the window used to calculate and generate a trap if the errors are crossing the thresholds within the window time frame.
Link Monitor Last TLV Send	Timestamp when the last Frame Error Event TLV, Frame Period Error Event, and Frame Second Summary Error Event TLV were transmitted on this interface.
Link Monitor Last TLV Received	Timestamp when the last Frame Error Event TLV, Frame Period Error Event, and Frame Second Summary Error Event TLV were received on this interface.

- To configure one or more interface, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.
- To apply the same settings to all interfaces, click Edit All.
- If you make any configuration changes in the Edit Link Monitor Configuration window, click Submit to apply the new settings to the interface(s), or click Cancel to abandon the changes and close the window.
- Click Refresh to update the screen with most recent data.



### 5.31.3 Dot3ah Remote Loopback Configuration

Use the Dot3ah Remote Loopback Configuration page to configure an interface for link monitoring. Fault localization and link performance testing can be done by Remote Loopback Mechanism.

To display this page, click Switching > Dot3ah > Remote Loopback Configuration in the navigation menu.

Figure 294: Dot1ag Remote Loopback Configuration

The screenshot shows the 'Dot3ah Remote Loopback Configuration' page. At the top, there are tabs for 'Configuration', 'Link Monitor Configuration', and 'Remote Loopback Configuration'. Below the tabs is a red header with the title 'Dot3ah Remote Loopback Configuration' and a help icon. The main content area contains a table with the following columns: 'Interface', 'Ethernet OAM', 'Remote Loopback Capability', 'Remote Loopback Status', and 'Remote Loopback Timeout'. The table displays 10 entries for interfaces 1/0/1 through 1/0/10. Each row has a checkbox on the left. Below the table is a pagination control with buttons for 'First', 'Previous', '1', '2', '3', '4', '5', 'Next', and 'Last'. A 'Display' dropdown is set to '10 rows' and a 'Filter' input field is present.

Interface	Ethernet OAM	Remote Loopback Capability	Remote Loopback Status	Remote Loopback Timeout
1/0/1	Disabled	Enabled	Disabled	50
1/0/2	Enabled	Enabled	Disabled	50
1/0/3	Disabled	Enabled	Disabled	50
1/0/4	Disabled	Enabled	Disabled	50
1/0/5	Disabled	Enabled	Disabled	50
1/0/6	Disabled	Enabled	Disabled	50
1/0/7	Disabled	Enabled	Disabled	50
1/0/8	Disabled	Enabled	Disabled	50
1/0/9	Disabled	Enabled	Disabled	50
1/0/10	Disabled	Enabled	Disabled	50

Table 276: Dot3ah Remote Loopback Configuration

Field	Description
Interface	The port associated with the rest of the data in the row.
Ethernet OAM	The status of Ethernet OAM protocol on the interface.
Remote Loopback Capability	Indicates whether the local entity can respond to Remote-Loopback OAM PDUs.
Remote Loopback Status	Indicates the status of the Remote-Loopback on the Ethernet OAM interface.
Remote Loopback Timeout	Indicates the Remote-Loopback timeout value.

- To configure one or more interface, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.
- To apply the same settings to all interfaces, click Edit All.
- If you make any configuration changes in the Edit Link Monitor Configuration window, click Submit to apply the new settings to the interface(s), or click Cancel to abandon the changes and close the window.
- Click Refresh to update the screen with most recent data.

## 5.32 Data Center Features

FASTPATH includes the following data center features that can be configured only by using the CLI.

- Data Center Bridging Exchange protocol (DCBX)
- FIP Snooping Bridge
- Congestion Notification (IEEE 802.1Qau)
- Enhanced Transmission Selection (ETS), defined in IEEE 802.1Qaz
- OpenFlow

This section provides an overview of these data center features. For information about the CLI commands you use to configure the data center features, refer to the FASTPATH CLI Command Reference Guide (FP8XXENT-SWUM2xx-R).

### 5.32.1 Data Center Bridging Exchange Protocol

The Data Center Bridging Exchange protocol (DCBX) is defined in the 802.1Qaz standard. DCBX is used by DCB devices to exchange configuration information with directly connected peers. DCBX uses type-length-value (TLV) information elements over LLDP to exchange information, so LLDP must be enabled on the port to enable the information exchange.

The main objective of DCBX is to perform the following operations:

- Discovery of DCB capability in a peer: DCBX is used to learn about the capabilities of the peer device. It is a means to determine if the peer device supports a particular feature such as PFC.
- DCB feature misconfiguration detection: DCBX can be used to detect misconfiguration of a feature between the peers on a link. Misconfiguration detection is feature-specific because some features may allow asymmetric configuration.
- Peer configuration of DCB features: DCBX can be used by a device to perform configuration of DCB features in its peer device if the peer device is willing to accept configuration.

DCBX is expected to be deployed in Fibre Channel over Ethernet (FCoE) topologies in support of lossless operation for FCoE traffic. In these scenarios, all network elements are DCBX enabled. In other words, DCBX is enabled end-to-end.

### 5.32.2 FIP Snooping

The FCoE Initialization Protocol (FIP) is used to perform the functions of FC\_BB\_E device discovery, initialization, and maintenance. FIP uses a separate EtherType from FCoE to distinguish discovery, initialization, and maintenance traffic from other FCoE traffic. FIP frames are standard Ethernet size (1518 Byte 802.1q frame), whereas FCoE frames are a maximum of 2240 bytes.

FIP snooping is a frame inspection method used by FIP Snooping Bridges to monitor FIP frames and apply policies based upon the L2 header information in those frames.

FIP snooping allows for:

- Auto-configuration of Ethernet ACLs based on information in the Ethernet headers of FIP frames.
- Emulation of FC point-to-point links within the DCB Ethernet network.
- Enhanced FCoE security/robustness by preventing FCoE MAC spoofing.
- The role of FIP snooping-enabled ports on the switch falls under one of the following types:
  - **Perimeter or Edge port (connected directly to a Fibre Channel end node or ENode).**
  - **Fibre Channel forwarder (FCF) facing port (that receives traffic from FCFs targeted to the ENodes).**

---

#### **NOTICE**

The FIP Snooping Bridge feature supports the configuration of the perimeter port role and FCF-facing port roles and is intended for use only at the edge of the switched network.

---

The default port role in an FCoE-enabled VLAN is as a perimeter port. FCF-facing ports are configured by the user.

### 5.32.3 Congestion Notification (Qau)

Congestion Notification (CN) is defined in the 802.1Qau standard. A consequence of link level pausing (i.e. 802.1Qbb) is congestion spreading—the domino effect of buffer congestion propagating upstream causing secondary bottlenecks. A layer two congestion control algorithm allows a primary bottleneck to directly reduce the rates of those sources whose packets pass through it, thereby preventing secondary bottlenecks.

Congestion notification is broken up into two algorithms: CP and RP. CP, Switch or Congestion Point Dynamics is the mechanism in which a switch buffer samples incoming packets and generates a feedback message addressed to the source of the sampled packets with the extent of the congestion. RP, Rate Limiter or Reaction Point Dynamics is the mechanism by which a Rate Limiter (RL) decreases its sending rate based on feedback and increases its rate voluntarily to recover lost bandwidth and probe for available bandwidth.

### 5.32.4 Enhanced Transmission Selection

Enhanced Transmission Selection (ETS) enables the sharing and redistribution of network bandwidth between various protocols. To support ETS, FASTPATH switches accept the ETS traffic class group and bandwidth information Application Priority TLV from auto-upstream devices and propagate them to auto-downstream devices. In addition, if iSCSI CoS is enabled, an additional entry in the Application Priority TLV is added. FASTPATH switches support the reception and propagation of ETS information in the automatic configuration port roles. They do not use the ETS information to configure traffic class groups or bandwidth allocations.

### 5.32.5 OpenFlow

In a classical router or switch, the packet forwarding (data path) and the high level routing decisions (control path) occur on the same device. An OpenFlow Switch separates these two functions. The data path portion still resides on the switch, while high-level routing decisions are moved to a separate controller, typically a standard server. The OpenFlow Switch and Controller communicate via the OpenFlow protocol, which defines messages, such as packet-received, send-packet-out, modify-forwarding-table, and get-stats. The data path of an OpenFlow Switch presents a clean flow table abstraction—each flow table entry contains a set of packet fields to match, and an action (such as send-out-port, modify-field, or drop). When an OpenFlow Switch receives a packet it has never seen before, for which it has no matching flow entries, it sends this packet to the controller. The controller then makes a decision on how to handle this packet. It can drop the packet, or it can add a flow entry directing the switch on how to forward similar packets in the future. OpenFlow is supported on both standalone switches and switch stacks.

## 5.33 802.1AS

IEEE 802.1AS is an Audio Video Bridging (AVB) feature available on some FASTPATH platforms. AVB helps ensure that the synchronization requirements are met for time-sensitive applications, such as audio and video, across bridged and virtual bridged local area networks consisting of LAN media where the transmission delays are fixed and symmetrical. This includes the maintenance of synchronized time during normal operation and following addition, removal, or failure of network components and network reconfiguration.

The IEEE 802.1AS standard specifies the protocol and procedures used to ensure that the QoS requirements are guaranteed for time-sensitive applications, such as audio and video. The IEEE 1588 Precision Time Protocol (PTP) forms the basis of the IEEE 802.1AS standard. PTP specifies a precise clock synchronization protocol that relies on time-stamped packets. The PTP protocol is applicable to distributed systems consisting of one or more nodes communicating over some set of communication media. The distribution of synchronous time information is performed in a hierarchical manner with a grandmaster clock at the root of the hierarchy. The grandmaster provides a common and precise time reference for one or more directly-attached slave devices by periodically exchanging timing information. In other words, all slave devices synchronize their clocks with the grandmaster clock. The slave devices can, in-turn, act as master devices for further hierarchical layers of slave devices.

### 5.33.1 802.1AS Configuration

Use this page to configure the global 802.1AS settings on the device. IEEE 802.1AS is an Audio Video Bridging (AVB) feature. AVB helps ensure that the synchronization requirements are met for time-sensitive applications, such as audio and video, across bridged and virtual bridged local area networks consisting of LAN media where the transmission delays are fixed and symmetrical. This includes the maintenance of synchronized time during normal operation and following addition, removal, or failure of network components and network reconfiguration.

To display the 802.1AS Configuration page, click Switching > Timing and Synchronization (802.1AS) > 802.1AS Configuration.

Figure 295: 802.1AS Global Configuration

802.1AS Global Configuration	
Administrative Mode	<input checked="" type="radio"/> Enable <input type="radio"/> Disable
Grandmaster Present	True
Grandmaster Change Count	0
Grandmaster Last Change Timestamp	0
Best Clock Identity	00:02:BC:FF:FE:4D:B9:E6
Best Clock Priority1	246
Best Clock Priority2	248
Steps to Best Clock	0
Local Clock Identity	00:02:BC:FF:FE:4D:B9:E6
Local Clock Priority1	<input type="text" value="246"/> (0 to 255)
Local Clock Priority2	<input type="text" value="248"/> (0 to 255)

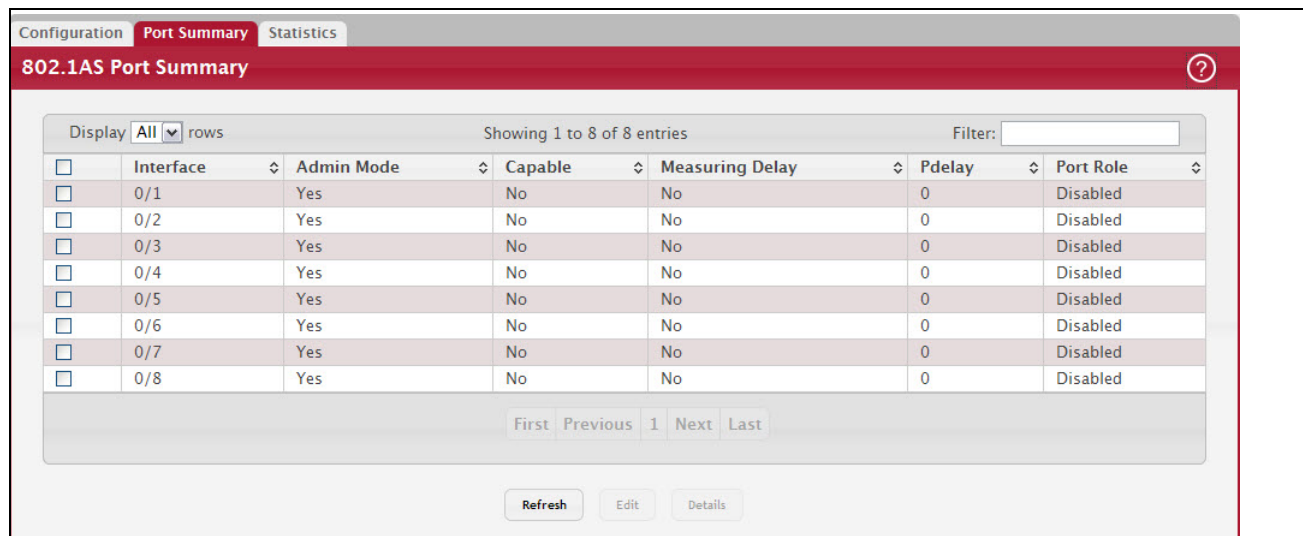
Table 277: 802.1AS Configuration Fields

Field	Description
Administrative Mode	The administrative mode of the 802.1AS feature on the device. When enabled, 802.1AS can calculate the time delay between devices on a given link and maintain an accurate view of a network clock.
Grandmaster Present	Identifies whether a grandmaster clock is present on the network segment. The grandmaster clock is the root device in the timing reference hierarchy.
Grandmaster Change Count	The number of times the grandmaster has changed from one device to another.
Grandmaster Last Change Timestamp	The system time when the most recent grandmaster change occurred.
Best Clock Identity	The MAC address of the clock device selected as the best master clock based on the best master clock (BMC) algorithm, which uses master clock selection criteria such as the clock priority, class, and accuracy.
Best Clock Priority1	The configured Priority 1 value on the clock device selected as the best master clock. A lower Priority 1 value has precedence over a higher Priority 1 value.
Best Clock Priority2	The configured Priority 2 value on the clock device selected as the best master clock. A lower Priority 2 value has precedence over a higher Priority 2 value.
Steps to Best Clock	The number of links in the path from the best master clock to this time-aware bridge. If this time-aware bridge is the best, the value is zero.
Local Clock Identity	The MAC address of the local time aware bridge (this device).
Local Clock Priority1	The Priority 1 value for the device. Configuring a lower value for the Priority 1 value increases the chances that this device will be selected as the best master clock. The priority values are the only best clock selection criteria that can be configured by an administrator. When the priority values for clock devices are the same, other properties, such as clock quality and clock accuracy, are used to determine which device is the best master clock.
Local Clock Priority2	The Priority 2 value for the device.

### 5.33.2 802.1AS Port Summary

Use the 802.1AS Port Settings page to configure and view per-port 802.1AS settings. To display the 802.1AS Port Settings page, click Switching > Timing and Synchronization (802.1AS) > Port Summary.

Figure 296: 802.1AS Port Summary



To configure one or more ports or LAGS, select the check box next to each port or LAG to configure and click Edit. You can select multiple ports or LAGs to apply the same settings to the selected interfaces.

Table 278: 802.1AS Port Settings Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When editing 802.1AS settings or viewing 802.1AS details for an interface, this field identifies the interface(s) being edited or viewed.
Admin Mode	Indicates whether 802.1AS is administratively enabled or disabled on the interface.
Capable	Indicates whether the interface is capable of communicating with a peer device using the 802.1AS protocol.
Measuring Delay	Indicates whether the interface is receiving PDELAY response messages from a peer device at the other end of the link.
Pdelay	The mean propagation delay on the interface.
Port Role	Indicates the 802.1AS role of the interface. The possible roles are as follows: <ul style="list-style-type: none"> <li>• Disabled</li> <li>• Master</li> <li>• Slave</li> <li>• Passive</li> </ul>
Pdelay Threshold	The propagation delay threshold on the interface. The threshold determines whether the port is capable of participating in the 802.1AS protocol. If the propagation delay on the interface is above the threshold you configure, the interface is not considered capable of participating in the 802.1AS protocol. The peer delay must be less than the threshold value configured on the interface.
Pdelay Lost Responses Allowed	The maximum number of PDELAY_REQ messages that can be sent without receiving a valid response before the interface is considered to be failing to exchange peer delay messages with its neighbor.
Initial Announce Interval	The logarithm to the base 2 of the mean-time interval between successive ANNOUNCE messages sent on this interface.
Initial Pdelay Interval	The logarithm to the base 2 of the mean-time interval between successive PDELAY_REQ messages sent on this interface.
Initial Sync Interval	The logarithm to the base 2 of the mean-time interval between successive SYNC messages sent on this interface.

**Table 278: 802.1AS Port Settings Fields (Continued)**

Field	Description
Announce Receipt Timeout	The number of ANNOUNCE intervals that must pass without receipt of an ANNOUNCE PDU before considering that the master clock is no longer transmitting.
Sync Receipt Timeout	The number of SYNC intervals that must pass without receipt of SYNC information before considering that the master clock is no longer transmitting.
Neighbor Rate Ratio	This table shows information about the initial and current intervals for sending timing messages. When an interface exchanges timing messages with a connected device, the connected device might change the interval values to correct time skew. The Initial value for each interval is the value specified in the Edit 802.1AS Port Configuration window. The Current value is interval on the interface after it has exchanged timing messages with its peer. <ul style="list-style-type: none"> <li>• Current Announce Interval–The current value of the ANNOUNCE interval.</li> <li>• Current Pdelay Interval–The current value of the PDELAY interval.</li> <li>• Current Sync Interval–The current value of the SYNC interval.</li> </ul>

### 5.33.3 802.1AS Statistics

Use this page to view information regarding the 802.1AS messages transmitted and received by each interface. To view additional 802.1AS statistics for an interface, select the interface with the information to view and click Details. The information below is organized according to the order in which the fields appear in the Details of 802.1AS Port Statistics window.

To display the 802.1AS Port Statistics page, click Switching > Timing and Synchronization (802.1AS) > Statistics.

**Figure 297: 802.1AS Statistics**

Interface	Pdelay Req Msgs Tx	Pdelay Req Msgs Rx	Announce Req Msgs Tx	Announce Req Msgs Rx	Sync Req Msgs Tx	Sync Req Msgs Rx
0/1	9762	0	0	0	0	0
0/2	9762	0	0	0	0	0
0/3	0	0	0	0	0	0
0/4	9762	0	0	0	0	0
0/5	9762	0	0	0	0	0
0/6	9762	0	0	0	0	0
0/7	9762	0	0	0	0	0
0/8	9762	0	0	0	0	0

The following table describes the information the 802.1AS Statistics page displays.

**Table 279: 802.1AS Port Statistics Fields**

Field	Description
Interface	The interface associated with the rest of the data in the row.
Propagation Delay	The statistics in this section provide information about the messages transmitted or received by the interface that are used to determine the propagation delay between the device and another clock device.
Pdelay Req Messages Transmitted	The number of PDELAY_REQ packets transmitted without error.

Table 279: 802.1AS Port Statistics Fields (Continued)

Field	Description
Pdelay Req Messages Received	The number of PDELAY_REQ packets received without error.
Pdelay Resp Messages Transmitted	The number of PDELAY_RESP packets transmitted without error.
Pdelay Resp Messages Received	The number of PDELAY_RESP packets received without error.
Pdelay Resp Followup Messages Transmitted	The number of PDELAY_RESP_FOLLOWUP packets transmitted without error.
Pdelay Resp Followup Messages Received	The number of PDELAY_RESP_FOLLOWUP packets received without error.
Pdelay Receipt Timeouts	The number of PDELAY_REQ messages that were transmitted that did not receive a PDELAY_RESP message within the allowed time.
Pdelay Receipt Discards	The number of PDELAY packets discarded.
Pdelay Allowed Lost Responses	The number of PDELAY packets that were sent to the link partner that did not receive a response and were considered to be lost.
Time Synchronization	The statistics in this section provide information about the messages transmitted or received by the interface that are used to synchronize the time between the device and another clock device.
Sync Req Messages Transmitted	The number of SYNC packets transmitted without error.
Sync Req Messages Received	The number of SYNC packets received without error.
Followup Messages Transmitted	The number of FOLLOWUP packets transmitted without error.
Followup Messages Received	The number of FOLLOWUP packets received without error.
Sync Receipt Timeouts	The number of SYNC messages that were transmitted that did not receive a response message within the allowed time.
Sync Messages Discarded	The number of SYNC packets that were discarded.
Best Master Clock Algorithm	The statistics in this section provide information about the ANNOUNCE messages transmitted or received by the interface. ANNOUNCE messages are used to select the best master clock on the network segment and to build the clock hierarchy.
Announce Req Messages Transmitted	The number of ANNOUNCE packets transmitted without error.
Announce Req Messages Received	The number of ANNOUNCE packets received without error.
Announce Receipt Timeouts	The number of ANNOUNCE packets that were transmitted that did not receive a response message within the allowed time.
Announce Messages Discarded	The number of ANNOUNCE packets that were discarded.
Invalid 802.1AS Messages Received	The number of messages that were received but included an error, such as a bad header.
Signaling	The statistics in this section provide information about the signaling messages transmitted or received by the interface. Signaling messages are used for non-time critical communication between the interface and other clock devices.
Messages Transmitted	The number of SIGNALING packets transmitted without error.
Messages Received	The number of SIGNALING packets received without error.

To reload the page, click Refresh. To reset the statistics for all interfaces, click Clear.

## 5.34 Multiple Registration Protocol Configuration

Like 802.1AS, Multiple Registration Protocol (MRP) is an AVB feature that is available on some FASTPATH platforms. MRP is a base registration protocol that enables devices running an MRP application to register attributes to other devices in a network. MRP provides an application to register attributes, such as the bandwidth requirement for a given AV stream and MAC address information. It is used by various applications to propagate the registration. FASTPATH switches support the following MRP applications:

- Multiple MAC Registration Protocol (MMRP)
- Multiple Stream Reservation Protocol (MSRP)
- Multiple VLAN Registration Protocol (MVRP)

MMRP allows for the propagation MAC address information in the network, and allows for the registration and deregistration of both individual MAC address information and group MAC address membership. End stations may request to join or leave a multicast group, or to register an individual MAC address with a specific VLAN. MAC address entries can be dynamically registered and deregistered if MMRP is administratively enabled on the switch.

MSRP reserves necessary resources in the network to facilitate time sensitive traffic to flow end to end. In a typical network, there are multiple Talkers (those who transmit streams) and multiple Listeners (those who receive streams from one or many Talkers). Each flow has specific bandwidth, frame rate, and time sync requirements. With the use of MSRP these resources are guaranteed through all intermediate devices that are between any talker and listener.

MVRP registers VLANs in the network, enabling automatic VLAN configuration on the switch. In a typical network, VLAN tagging is common. Many nodes require ingress traffic to be tagged with specific VLAN ID, and other nodes require egress traffic to be transmitted with a specific VLAN ID. With the use of MVRP on both ingress and egress, no manual VLAN configuration is required to pass tagged traffic through the network.

---

### **NOTICE**

MRP framework must be available and enabled in all intermediate devices to ensure that the propagation of the attributes occurs throughout the network.

---

With MRP, network attributes are declared, registered, withdrawn, and removed completely dynamically without any user intervention. This dynamic nature is especially useful in networks where:

- Network attributes are likely to change frequently, requiring reconfiguration of the intermediate devices.
- Recipients of these attributes frequently increase or decrease in number.
- Each of these changes without a dynamic self-adjusting framework would require constant attention from the network administrator.

### 5.34.1 MRP Configuration

Use the MRP Configuration page to configure global MRP settings for the switch. To access the basic MRP Configuration page, click Switching > MRP > Configuration.

---

### **NOTICE**

The fields available on the MRP Configuration page vary based on the platform and its supported features.

---



Figure 298: MRP Configuration

Table 280: MRP Configuration Fields

Field	Description
MVRP	<p>Multiple VLAN Registration Protocol (MVRP) registers VLANs in the network, enabling automatic VLAN configuration on the device. In a typical network, VLAN tagging is common. Many nodes require ingress traffic to be tagged with a specific VLAN ID, and other nodes require egress traffic to be transmitted with a specific VLAN ID. With the use of MVRP on both ingress and egress, no manual VLAN configuration is required to pass tagged traffic through the network.</p> <ul style="list-style-type: none"> <li>• Admin Mode–The administrative mode of MVRP on the device.</li> <li>• Periodic State Machine–Select this option to enable the MRP Periodic State Machine for MVRP on the device. When enabled, the state machine can help limit the effect of topology changes and reduce the number of protocol data units (PDUs) transmitted between devices.</li> </ul>
MMRP	<p>Multiple MAC Registration Protocol (MMRP) allows the propagation of MAC address information in the network, and allows for the registration and deregistration of both individual MAC address information and group MAC address membership. End stations may request to join or leave a multicast group, or to register an individual MAC address with a specific VLAN. MAC address entries can be dynamically registered and deregistered if MMRP is administratively enabled on the device.</p> <ul style="list-style-type: none"> <li>• Admin Mode–The administrative mode of MMRP on the device.</li> <li>• Periodic State Machine–Select this option to enable the MRP Periodic State Machine for MMRP on the device. When enabled, the state machine can help limit the effect of topology changes and reduce the number of protocol data units (PDUs) transmitted between devices.</li> </ul>
MSRP	<p>Multiple Stream Reservation Protocol (MSRP) reserves necessary resources in the network to facilitate the end-to-end flow of time sensitive traffic. In a typical network, there are multiple Talkers (those who transmit streams) and multiple Listeners (those who receive streams from one or many Talkers). Each flow has specific bandwidth, frame rate, and time sync requirements. With the use of MSRP these resources are guaranteed through all intermediate devices that are between any Talker and Listener.</p> <ul style="list-style-type: none"> <li>• Admin Mode–The administrative mode of MSRP on the device.</li> <li>• Periodic State Machine–Select this option to enable the MRP Periodic State Machine for MSRP on the device. When enabled, the state machine can help limit the effect of topology changes and reduce the number of protocol data units (PDUs) transmitted between devices.</li> </ul>
Talker Pruning	<p>Select this option to enable MSRP Talker pruning. The MSRP Talker is the source of an AV stream. When enabled, Talker pruning stops MSRP declarations sent by the talker unless a Listener registers and requests them.</p>

Table 280: MRP Configuration Fields (Continued)

Field	Description
Boundary Propagation	Select this option to enable MSRP boundary propagation on the device.
Max Fan In Ports	The maximum number of ports where MSRP registrations are allowed.

### 5.34.2 MRP Interface Configuration

Use the MRP Interface Configuration page to view and configure the per-interface Multiple Registration Protocol (MRP) settings. To change the current settings for one or more interfaces, select each interface to modify and click Edit. The same MRP settings are applied to all selected interfaces.

To access the MRP Interface Configuration page, click Switching > MRP > Interface in the navigation menu. In the following image, the MMRP mode on ports g4 and g5 is being enabled.

Figure 299: MRP Port Configuration

The screenshot shows the 'MRP Interface Configuration' page with a table of 14 entries. The table has columns for Interface, MVRP Mode, MMRP Mode, MSRP Mode, MSRP SR Class PVID, Join Timer, Leave Timer, and Leave All Timer. The first 10 entries are visible, showing interfaces 0/1 through 0/8 and 1/1 through 1/2. The MMRP Mode column is highlighted in red for interfaces 0/4 and 0/5, indicating they are being configured.

Interface	MVRP Mode	MMRP Mode	MSRP Mode	MSRP SR Class PVID	Join Timer	Leave Timer	Leave All Timer
<input type="checkbox"/> 0/1	Enabled	Disabled	Enabled	2	20	300	2000
<input type="checkbox"/> 0/2	Enabled	Disabled	Enabled	2	20	300	2000
<input type="checkbox"/> 0/3	Enabled	Disabled	Enabled	2	20	300	2000
<input type="checkbox"/> 0/4	Enabled	Disabled	Enabled	2	20	300	2000
<input type="checkbox"/> 0/5	Enabled	Disabled	Enabled	2	20	300	2000
<input type="checkbox"/> 0/6	Enabled	Disabled	Enabled	2	20	300	2000
<input type="checkbox"/> 0/7	Enabled	Disabled	Enabled	2	20	300	2000
<input type="checkbox"/> 0/8	Enabled	Disabled	Enabled	2	20	300	2000
<input type="checkbox"/> 1/1	Enabled	Disabled			20	300	2000
<input type="checkbox"/> 1/2	Enabled	Disabled			20	300	2000

To configure one or more ports or LAGs, select the check box next to each port or LAG to configure. You can select multiple ports to apply the same settings to the selected interfaces.

Table 281: MRP Port Configuration Fields

Field	Description
Interface	Identifies the interface associated with the rest of the information in the row.
MVRP Mode	The administrative mode of Multiple VLAN Registration Protocol (MVRP) on the interface. MVRP registers VLANs in the network, enabling automatic VLAN configuration on the device.
MMRP Mode	The administrative mode of Multiple MAC Registration Protocol (MMRP) on the interface. MMRP allows the propagation of MAC address information in the network and allows for the registration and deregistration of both individual MAC address information and group MAC address membership.
MSRP Mode	The administrative mode of Multiple Stream Reservation Protocol (MSRP) on the interface. MSRP reserves necessary resources in the network to facilitate the end-to-end flow of time sensitive traffic.
MSRP SR Class PVID	The default VLAN ID to be used for MSRP stream traffic.
Join Timer	The amount of time to wait for JoinIn messages from other MRV participants after the interface sends a Join message. If the amount of time specified in this field passes before the interface receives a JoinIn message, the interface resends the Join message.

Table 281: MRP Port Configuration Fields (Continued)

Field	Description
Leave Timer	The amount of time to wait before the interface deregisters attributes from other MRV participants. If the interface receives Join messages from other participants before the Leave timer expires, the attributes are not deregistered.
Leave All Timer	The amount of time to wait, after the interface starts the MRP registration process, before the participants refresh and re-register their attributes.

### 5.34.3 Qav Mapping

Use the 802.1Qav MSRP Mapping page to configure and view the IEEE 802.1Qav-to-802.1p priority mappings on the device. The IEEE 802.1Qav standard supports time-sensitive traffic streams by placing all device traffic, including legacy asynchronous Ethernet traffic, through queuing and forwarding. When a Talker declares a stream, it identifies whether the stream is class A or class B and specifies the stream's bandwidth requirements. Class A traffic has a higher transmission priority than class B traffic. In an EAV network, Audio/Video (AV) traffic is given priority by reserving bandwidth for a given stream ID.

On the Qav Parameters page, you can view and configure selected bandwidth allocations for Class A and Class B traffic. To display the Qav Statistics page, click Switching > MRP > Qav Mapping.

Figure 300: Qav Mapping

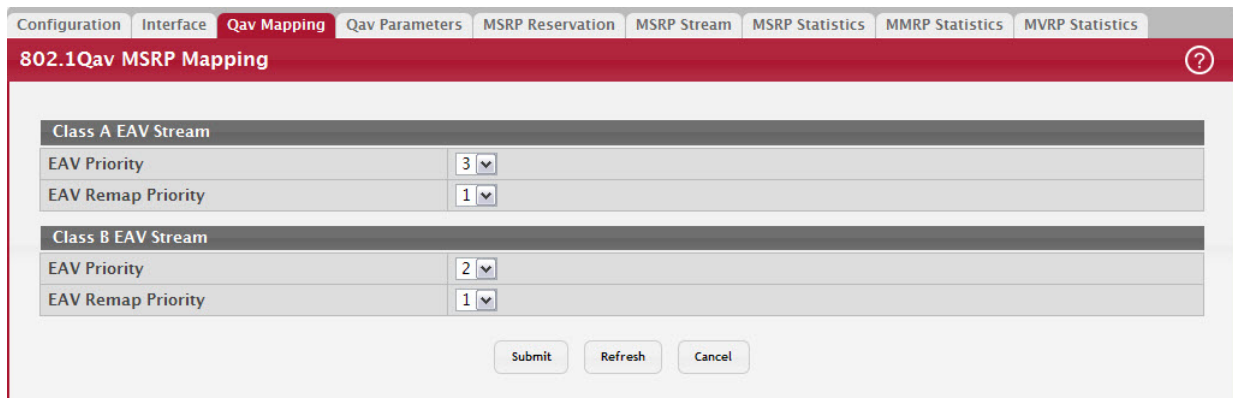


Table 282: Qav Mapping Fields

Field	Description
Class A EAV Stream	This section allows the configuration of the 802.1p priority value for class A Ethernet Audio/Video (EAV) traffic and non-AV traffic that has the same 802.1p priority as the class A traffic stream. <ul style="list-style-type: none"> <li>EAV Priority–The 802.1p priority value to assign to class A EAV traffic. A higher value indicates a higher priority for the traffic.</li> <li>EAV Priority–The 802.1p value to assign to non-EAV traffic that has the same 802.1p priority as class A EAV streams when the non-EAV traffic enters an EAV cloud.</li> </ul>
Class B EAV Stream	This section allows the configuration of the 802.1p priority value for class B EAV traffic and non-AV traffic that has the same 802.1p priority as the class B traffic stream. <ul style="list-style-type: none"> <li>EAV Priority–The 802.1p priority value to assign to class B EAV traffic. A higher value indicates a higher priority for the traffic.</li> <li>EAV Priority–The 802.1p value to assign to non-EAV traffic that has the same 802.1p priority as class B EAV streams when the non-EAV traffic enters an EAV cloud.</li> </ul>

### 5.34.4 Qav Parameters

Use the Qav Parameters page to configure and view the per-port IEEE 802.1Qav settings. The IEEE 802.1Qav standard supports time-sensitive traffic streams by pacing all switch traffic, including legacy asynchronous Ethernet traffic,

through queuing and forwarding. When a Talker declares a stream, it identifies whether the stream is Class A or Class B and specifies the stream's bandwidth requirements. Class A traffic has a higher transmission priority than Class B traffic. On the Qav Parameters page, you can view and configure selected bandwidth allocations for Class A and Class B traffic. To display the Qav Statistics page, click Switching > MRP > Qav Parameters.

Figure 301: Qav Parameters

Interface	Class A Delta Bandwidth	Class A Bandwidth Allocated	Class A Bandwidth Remaining	Class B Delta Bandwidth	Class B Bandwidth Allocated	Class B Bandwidth Remaining	Total Bandwidth Allocated	Total Bandwidth Remaining
0/1	75	0	9351000	0	0	9351000	0	9351000
0/2	75	0	9351000	0	0	9351000	0	9351000
0/3	75	0	0	0	0	0	0	0
0/4	75	0	9351000	0	0	9351000	0	9351000
0/5	75	0	9351000	0	0	9351000	0	9351000
0/6	75	0	9351000	0	0	9351000	0	9351000
0/7	75	0	9351000	0	0	9351000	0	9351000
0/8	75	0	9351000	0	0	9351000	0	9351000

To configure Qav parameters, select the check box next to the port to configure and edit the fields as desired. You can select multiple ports to apply the same settings to the selected interfaces.

Table 283: Qav Parameter Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring the settings for one or more interfaces, this field identifies the interface(s) being configured.
Class A MSRP delta bandwidth	The additional bandwidth that can be allocated to a port for class A traffic. The value is represented as a percentage of port transmit rate that can be reserved for class A traffic. Class A traffic has a higher priority than class B traffic.
Class A Bandwidth Allocated	Shows the current rate of the class A traffic on interface (in Bps).
Class A Bandwidth Remaining	Shows the maximum rate of the class A traffic available on interface (in Bps).
Class B MSRP delta bandwidth	The additional bandwidth that can be allocated to a port for class B traffic. The value is represented as a percentage of port transmit rate that can be reserved for class B traffic. Class B traffic has a lower priority than class A traffic.
Class B Bandwidth Allocated	Shows the current rate of the class B traffic on interface (in Bps).
Class B Bandwidth Remaining	Shows the maximum rate of the class B traffic available on interface (in Bps).
Total Bandwidth Allocated	Sum of the allocated Class A and Class B traffic rates on interface (in Bps).
Total Bandwidth Remaining	This value is 75% of the interface speed minus total allocated bandwidth (in Bps/sec).

### 5.34.5 MSRP Reservation Parameters

Use the MSRP Reservation Parameters page to view information about the talker, listener, and intermediate device sta-

tus for the devices involved in each MSRP stream flowing through the switch.

To display the MSRP Reservation Parameters page, click the Switching tab, then click MRP > Advanced > MSRP Reservation.

Figure 302: MSRP Reservation Parameters

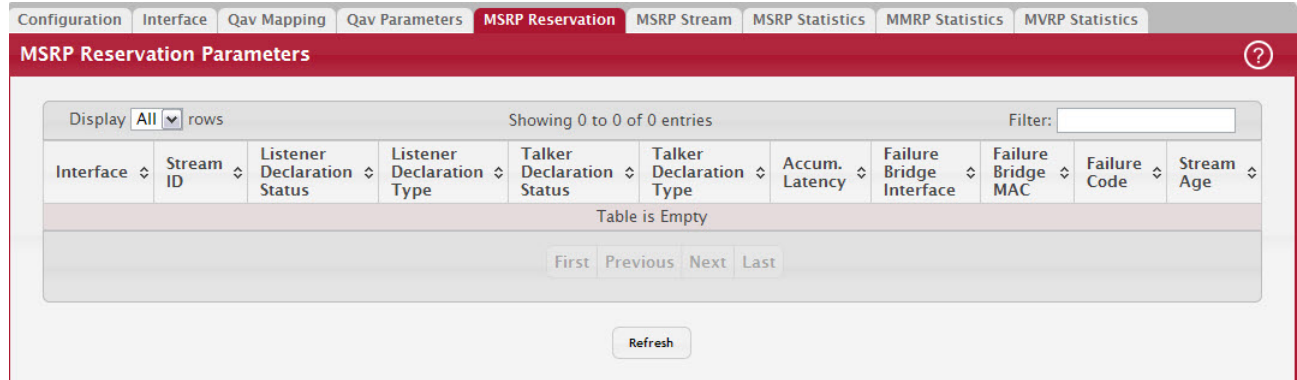


Table 284: MSRP Reservation Parameters Fields

Field	Description
Interface	Identifies the interface associated with the rest of the information in the row.
Stream ID	A 16-bit unsigned integer value, Unique ID, used to distinguish among multiple streams sourced by the same system.
Listener Declaration Status	Identifies the MSRP declaration status of the listener attribute.
Listener Declaration Type	Identifies the MSRP declaration type of the listener attribute.
Talker Declaration Status	Identifies the MSRP declaration status of the talker attribute.
Talker Declaration Type	Identifies the MSRP declaration type of the talker attribute.
Accumulated Latency	Identifies how much latency, in nanoseconds, the stream has suffered in its path from the Talker to a given Listener. It starts as a 0 in a Talker Advertise Declaration at the Talker, and its value is increased by one for each bridge as the Talker Advertise Declaration propagates through the network.
Failure Bridge Interface	Identifies the interface on the Bridge where the failure occurred.
Failure Bridge MAC	Identifies the MAC address of the switch where the failure occurred.

Table 284: MSRP Reservation Parameters Fields (Continued)

Field	Description
Failure Code	Shows the number that represents the reason for the failure. The switch supports the following codes: <ul style="list-style-type: none"> <li>• 1–Insufficient bandwidth</li> <li>• 3–Insufficient bandwidth for the traffic class</li> <li>• 5–Stream destination_address is already in use</li> <li>• 7–Reported latency has changed</li> <li>• 8–Egress port is not Audio/Video Bridging (AVB) capable</li> <li>• 9–Use a different destination_address (i.e. MAC DA hash table full)</li> <li>• 12–Cannot store destination_address (i.e., Bridge is out of MAC DA resources)</li> <li>• 13–Requested priority is not an SR Class priority</li> <li>• 14–MaxFrameSize is too large for media</li> <li>• 15–msrpMaxFanInPorts limit has been reached</li> <li>• 16–Changes in FirstValue for a registered StreamID</li> <li>• 17–VLAN is blocked on this egress port (Registration Forbidden)</li> </ul>
Stream Age	The time, in seconds, since the stream destination address was added to the Dynamic Reservations Entries table. A value of zero indicates the destination address has not been added to the table.

### 5.34.6 MSRP Streams Information

Use the MSRP Stream Information page to view information about MSRP streams flowing through each interface. To display the MSRP Stream Information page, click Switching > MRP > MSRP Stream.

Figure 303: MSRP Streams Information

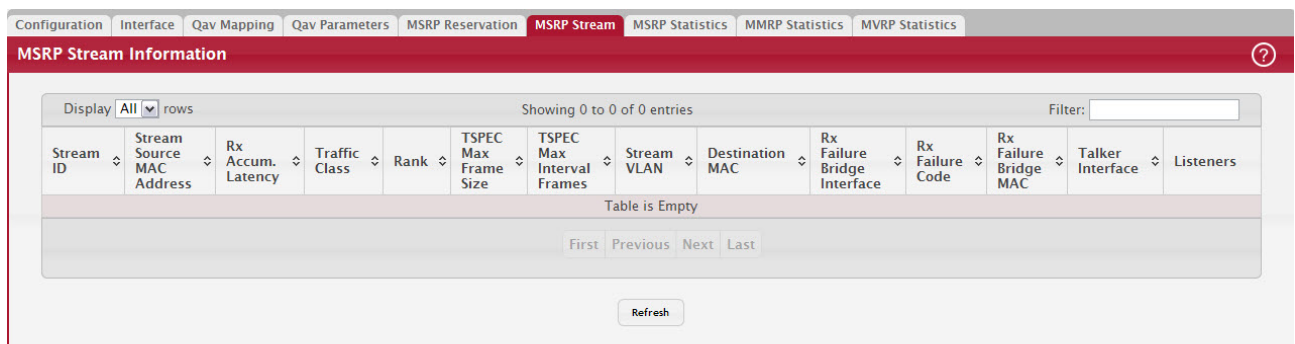


Table 285: MSRP Streams Information Fields

Field	Description
Stream ID	A 16-bit unsigned integer value, Unique ID, used to distinguish among multiple streams sourced by the same system.
Stream Source MAC Address	Identifies the MAC address of the traffic stream's source.
Received Accumulated Latency	The 32-bit unsigned Accumulated Latency component is used to determine the worst-case latency that a Stream can suffer in its path from the Talker to a given Listener. It starts as a 0 in a Talker Advertise Declaration at the Talker, and its value is increased by each Bridge as the Talker Advertise Declaration propagates through the network.
Traffic Class	Identifies whether the stream is Class A or Class B. Class A traffic has a higher priority than Class B traffic.

Table 285: MSRP Streams Information Fields (Continued)

Field	Description
Rank	The 5-bit unsigned Rank component is used by systems to decide which streams can and cannot be served, when the MSRP registrations exceed the capacity of a Port to carry the corresponding data streams. If a Bridge becomes oversubscribed (e.g. network reconfiguration, 802.11 bandwidth reduction) the Rank will also be used to help determine which Stream or Streams can be dropped. A lower numeric value is more important than a higher numeric value.
TSPEC Max Frame Size	The 32-bit unsigned Bandwidth component is used to allocate resources and adjust queue selection parameters to supply the quality of service requested by an MSRP Talker Declaration. It represents the maximum rate, in units of 1024 octets per second, at which frames in the Stream referenced by the Talker Declaration may be transmitted.
TSPEC Max Interval Frames	The 32-bit unsigned Frame Rate component is used to allocate resources and adjust queue selection parameters to supply the quality of service requested by an MSRP Talker Declaration. It represents the maximum number of frames that the Talker may transmit in one second.
Stream VLAN	Identifies the VLAN ID of the traffic stream.
Destination MAC	Identifies the MAC address of the traffic stream's destination.
Rx Failure Bridge Interface	Identifies the interface on the Bridge where the failure occurred.
Received Failure Code	Identifies the code value of the failure. For more information about the failure codes, see the Failure Code field description in Table 284, "MSRP Reservation Parameters Fields," on page 335.
Rx Failure Bridge MAC	Identifies the MAC address of the switch where the failure occurred.
Talker Interface	Identifies the interface on which the Talker is present.
Listeners	Identifies the interface on which Listeners are present.

### 5.34.7 MSRP Statistics

The MSRP Statistics page displays information about the MSRP frames transmitted and received by the switch and by each interface. To access the MSRP Statistics page, click Switching > MRP > MSRP Statistics.

Figure 304: MSRP Statistics

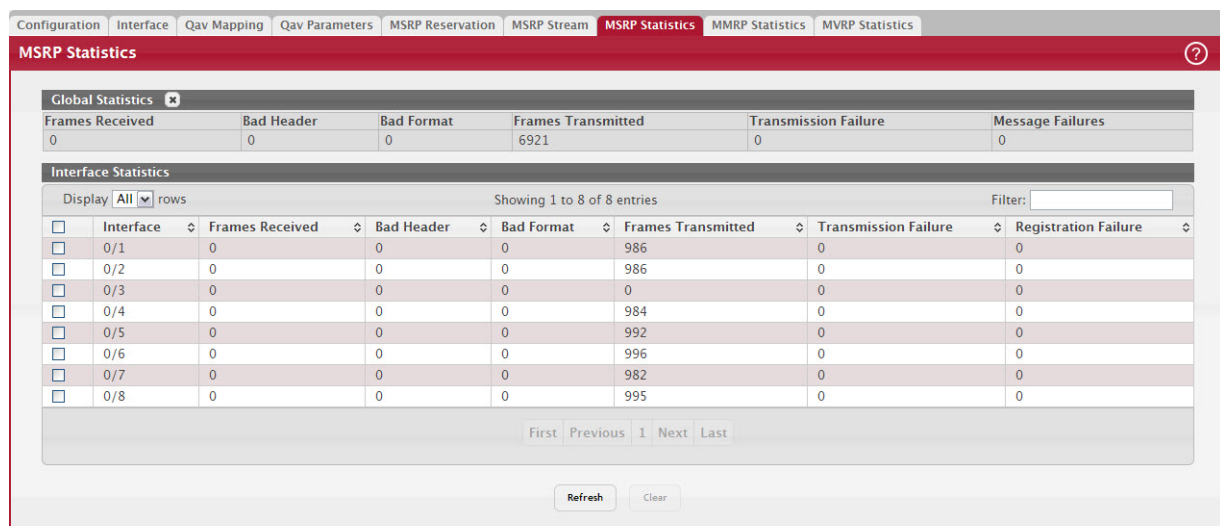




Table 286: MSRP Statistics

Field	Description
Interface	In the Interface Statistics table, this field identifies the interface associated with the rest of the data in the row.
Frames Received	Shows number of MSRP frames that have been received on the switch.
Bad Header	Shows number of MSRP frames with bad headers that have been received on the switch.
Bad Format	Shows number of MSRP frames with bad PDUs body formats that have been received on the switch.
Frames Transmitted	Shows number of MSRP frames which that have been transmitted on the switch.
Transmission Failures	Shows number of MSRP frames the switch failed to transmit.
Message Failures	Shows the number of messages that failed to be added to the queue.
Registration Failures	Shows the number of MSRP frames that failed to register on a device or particular interface.

To reload the page, click Refresh. To clear the statistics for one or more ports, select the check box next to the interface or interfaces, and click Clear. To clear the statistics for all interfaces, select the check box in the heading row, and click Clear.

### 5.34.8 MMRP Statistics

The MMRP Statistics page displays information regarding the MMRP frames transmitted and received by the switch and by each interface. To access the MMRP Statistics page, click the Switching tab, then click MRP > Advanced > MMRP Statistics.

Figure 305: MMRP Statistics

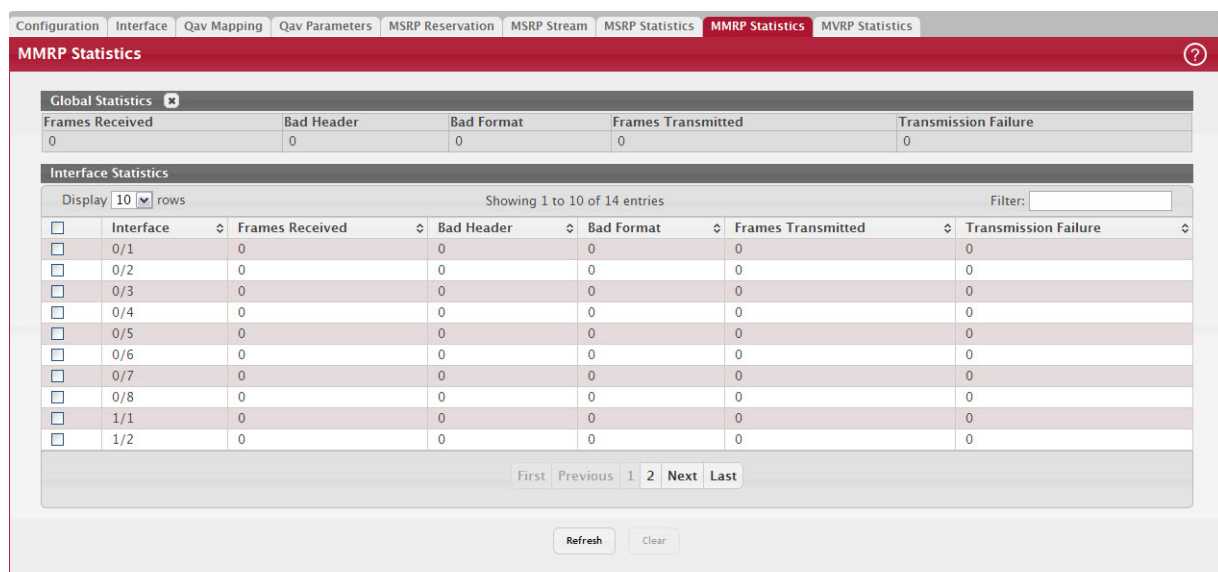


Table 287: MMRP Statistics Fields

Field	Description
Interface	In the Interface Statistics table, this field identifies the interface associated with the rest of the data in the row.
Frames Received	Shows the number of MMRP frames which were received on the switch.
Bad Header	Shows number of MMRP frames with bad headers which were received on the switch.



**Table 287: MMRP Statistics Fields (Continued)**

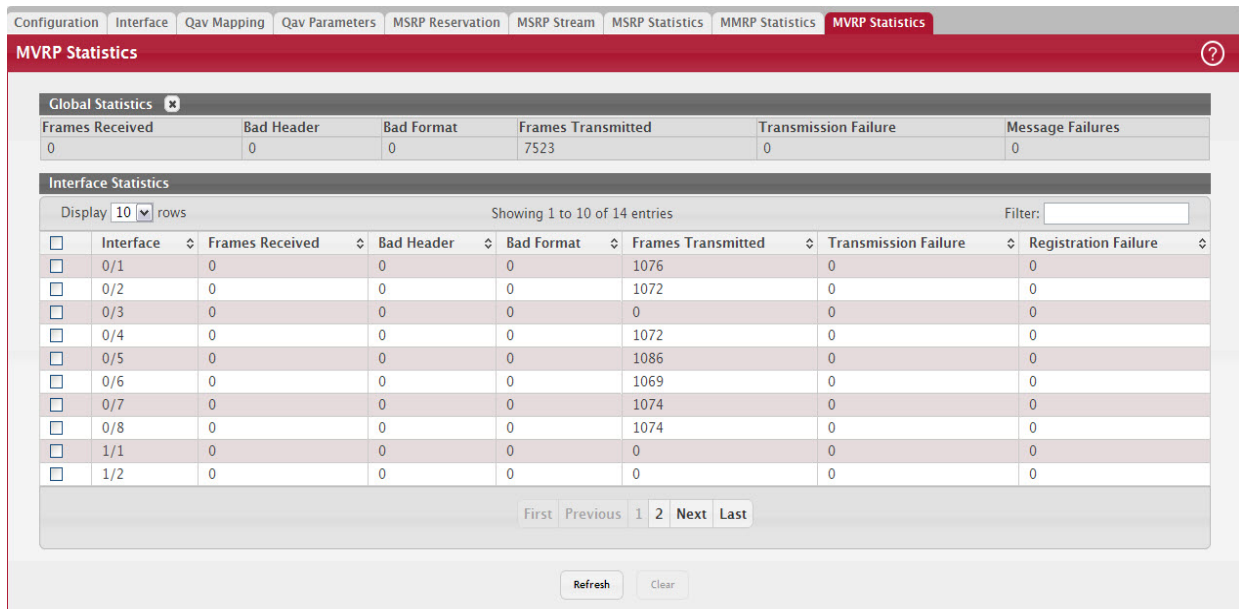
Field	Description
Bad Format	Shows number of MMRP frames with bad PDUs body formats which were received on the switch.
Frames Transmitted	Shows number of MMRP frames which were transmitted on the switch.
Transmission Failures	Shows number of MMRP frames that the switch failed to transmit.

To reload the page, click Refresh. To clear the statistics for one or more ports, select the check box next to the interface or interfaces, and click Clear. To clear the statistics for all interfaces, select the check box in the heading row, and click Clear.

### 5.34.9 MVRP Statistics

The MVRP Statistics page displays information about the MVRP frames transmitted and received by the switch and by each interface. To access the MVRP Statistics page, click Switching > MRP > MVRP Statistics.

**Figure 306: MVRP Statistics**



**Table 288: MVRP Statistics**

Field	Description
Interface	In the Interface Statistics table, this field identifies the interface associated with the rest of the data in the row.
Frames Received	Shows number of MVRP frames that have been received on the switch.
Bad Header	Shows number of MVRP frames with bad headers that have been received on the switch.
Bad Format	Shows number of MVRP frames with bad PDUs body formats that have been received on the switch.
Frames Transmitted	Shows number of MVRP frames which that have been transmitted on the switch.
Transmission Failures	Shows number of MVRP frames the switch failed to transmit.
Message Failures	Shows the number of messages that failed to be added to the queue.
Registration Failures	Shows the number of MVRP frames that failed to register on a device or particular interface.

To reload the page, click Refresh. To clear the statistics for one or more ports, select the check box next to the interface or interfaces, and click Clear. To clear the statistics for all interfaces, select the check box in the heading row, and click Clear Counters.

## 5.35 IP Device Tracking

The IPv4 Device Tracking (IPDT) feature enables the network administrator to see which IPv4 addresses are attached to which physical ports or LAGs. This information is available for non-routing-enabled switches as well as VLAN routing interfaces on routing-enabled switches.

The DHCP Snooping feature (see [Section 5.15: "Configuring DHCP Snooping"](#)) also provides mapping from the host IP address to a physical port on an L2 switch, for the IP address acquired using DHCP. However, the DHCP Snooping feature cannot track the statically-configured hosts, nor can it detect the movement of the hosts in a VLAN. The IPDT feature snoops the ARP packets exchanged in a VLAN and populates the tracking table with the IP address, MAC address, VLAN, and interface for each host.

### 5.35.1 Device Tracking Global Configuration

Use the Device Tracking Global Configuration page to view and configure the global settings for IPDT. To access the Device Tracking Global Configuration page, click Switching > Device Tracking > Global.

Figure 307: Device Tracking Global Configuration

Field	Description
Admin Mode	The administrative mode of the IPDT feature on the device. Disabling the administrative mode clears all the entries in the IPDT table.
Probe Generation	The ARP probe generation mode for the entries in the IPDT table. For each device entry in the IPDT table, an ARP probe is sent periodically to check the reachability of the device. If there are no ARP responses received for the configured number of retransmit ARP probes, the device entry is marked inactive.
Probe Count	The number of probes sent to the device, without any response from the device, before the device is declared as inactive in the IPDT table.
Probe Interval (Seconds)	The number of seconds that IPDT should wait before sending an ARP probe to the device entries in the IPDT table.

Table 289: Device Tracking Global Configuration

Field	Description
Admin Mode	The administrative mode of the IPDT feature on the device. Disabling the administrative mode clears all the entries in the IPDT table.
Probe Generation	The ARP probe generation mode for the entries in the IPDT table. For each device entry in the IPDT table, an ARP probe is sent periodically to check the reachability of the device. If there are no ARP responses received for the configured number of retransmit ARP probes, the device entry is marked inactive.
Probe Count	The number of probes sent to the device, without any response from the device, before the device is declared as inactive in the IPDT table.
Probe Interval (Seconds)	The number of seconds that IPDT should wait before sending an ARP probe to the device entries in the IPDT table.

**Table 289: (Continued) Device Tracking Global Configuration**

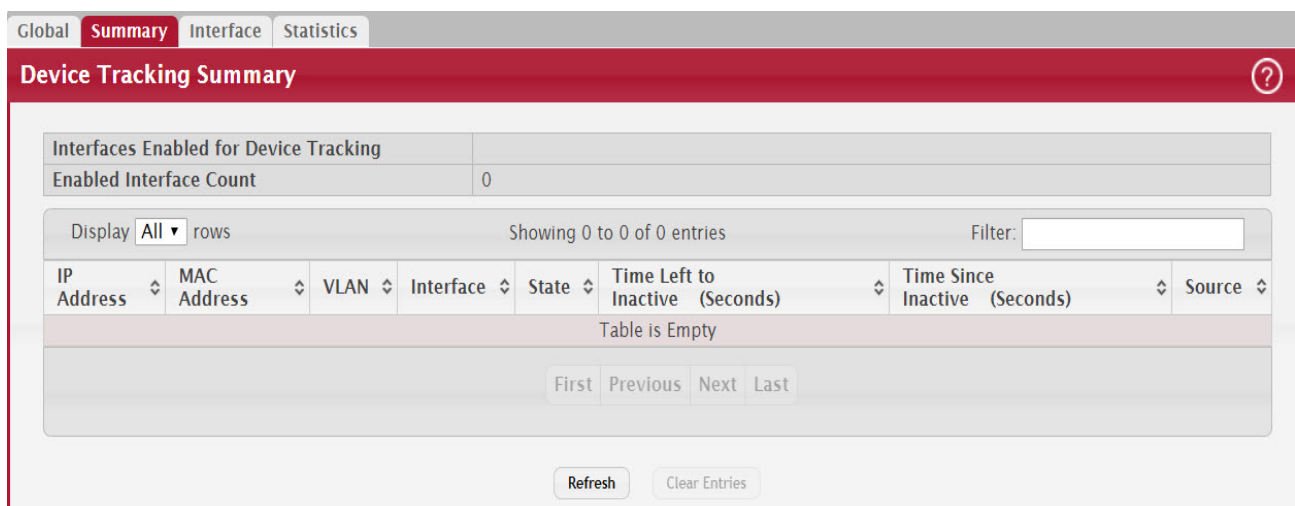
Field	Description
Probe Delay (Seconds)	The number of seconds to delay sending the first ARP probe to the IPDT table entries, when the interface associated with the device entry moves from the non-forwarding state to the forwarding state.
Host IP Address / Mask	The source IP address and mask in the ARP probes generated by IPDT and sent to the device entries in the IPDT table on non-routing interfaces.

Click Submit to send the updated configuration to the switch. Click Refresh to update the page with the most current information. Click Cancel to cancel the changes.

### 5.35.2 Device Tracking Summary

Use the Device Tracking Summary page to display information about the device entries in the IPDT table that are learned on the IPDT-enabled interfaces. To access the Device Tracking Summary page, click Switching > Device Tracking > Summary.

**Figure 308: Device Tracking Summary**



**Table 290: Device Tracking Summary**

Field	Description
Interfaces Enabled for Device Tracking	The list of interfaces that are enabled for IPDT.
Enabled Interface Count	The total number of IPDT-enabled interfaces.
IP Address	The learned IP address of the device.
MAC Address	The MAC address associated with the learned IP address of the device.
VLAN	The VLAN ID associated with the interface on which the device is learned.
Interface	The interface on which the device is learned.
State	The state of the learned device in the IPDT table, which can be one of the following: <ul style="list-style-type: none"> <li>Active–The device is reachable and responding to the ARP probes sent.</li> <li>Inactive–The device is not reachable.</li> </ul>
Time Left to Inactive (Seconds)	The number of seconds remaining before an active device in the IPDT table is marked inactive.

Table 290: Device Tracking Summary (Continued)

Field	Description
Time Since Inactive (Seconds)	The number of seconds elapsed since the inactive device in the IPDT table was last reachable.
Source	The source of the learned device in the IPDT table, which can be one of the following: <ul style="list-style-type: none"> <li>• ARP–ARP snooping detected new devices.</li> <li>• DHCP–DHCP snooping detected DHCP client devices.</li> </ul>

The administrative mode of the IPDT feature must be enabled on the device to clear the IPDT table entries. To remove all entries from the IPDT table, click the Clear Entries button. You must confirm the action before the entries are deleted. The table is repopulated as new devices are learned by the IPDT feature. Click Refresh to update the page with the most current information.

### 5.35.3 Device Tracking Interface Configuration

Use the Device Tracking Interface Configuration page to configure the IPDT settings on specific interfaces. To configure the settings for one or more interfaces, select each entry to modify and click Edit. The same IPDT settings are applied to all selected interfaces. To access the Device Tracking Interface Configuration page, click Switching > Device Tracking > Interface.

Figure 309: Device Tracking Interface Configuration

The screenshot shows the 'Device Tracking Interface Configuration' page. At the top, there is a red header with the title. Below the header, there is a control bar with 'Display 10 rows', 'Showing 1 to 10 of 78 entries', and a 'Filter:' input field. The main content is a table with two columns: 'Interface' and 'Maximum Learned Entries'. Each row in the table has a checkbox on the left. The 'Interface' column lists interfaces from 0/1 to 0/10. The 'Maximum Learned Entries' column shows 'No Limit' for all interfaces. At the bottom of the table, there is a pagination control with buttons for 'First', 'Previous', '1', '2', '3', '4', '5', 'Next', and 'Last'. Below the table, there are two buttons: 'Refresh' and 'Edit'.

Interface	Maximum Learned Entries
<input type="checkbox"/> 0/1	No Limit
<input type="checkbox"/> 0/2	No Limit
<input type="checkbox"/> 0/3	No Limit
<input type="checkbox"/> 0/4	No Limit
<input type="checkbox"/> 0/5	No Limit
<input type="checkbox"/> 0/6	No Limit
<input type="checkbox"/> 0/7	No Limit
<input type="checkbox"/> 0/8	No Limit
<input type="checkbox"/> 0/9	No Limit
<input type="checkbox"/> 0/10	No Limit

Table 291: Device Tracking Interface Configuration

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring IPDT settings, this field identifies the interfaces that are being configured.
Maximum Learned Entries	The maximum number of learned entries that can be added to the IPDT table per interface. If the maximum limit configured on an interface is zero, IPDT is effectively disabled on that interface. If the current number of entries learned on an interface is already more than the maximum limit configured on the interface, all the entries associated with the interface are deleted from the IPDT table, and ARP probes are sent again to the devices previously learned on that interface. By default, there is no limit on the number of entries that can be learned per interface.

Click Refresh to update the page with the most current information.

### 5.35.4 Device Tracking Statistics

The Device Tracking Statistics page displays information about the number and type of the learned entries in the IPDT table. To access the Device Tracking Statistics page, click Switching > Device Tracking > Statistics.

Figure 310: Device Tracking Statistics

Field	Description
ARP Entries	0
DHCP Entries	0
Active Entries	0
Inactive Entries	0
Total Entries	0

Refresh

Table 292: Device Tracking Statistics

Field	Description
ARP Entries	The number of device entries learned by ARP snooping.
DHCP Entries	The number of client devices learned by DHCP snooping.
Active Entries	The number of device entries currently in an active state.
Inactive Entries	The number of device entries currently in an inactive state.
Total Entries	The total number of device entries in the IPDT table.

Click Refresh to update the page with the most current information.

## 6/ Configuring Routing

FASTPATH supports IP routing. Use the links in the Routing navigation menu folder to manage routing on the system.

When a packet enters the switch, the destination MAC address is checked to see if it matches any of the configured routing interfaces. If it does, then the silicon searches the host table for a matching destination IP address. If an entry is found, then the packet is routed to the host. If there is not a matching entry, then the switch performs a longest prefix match on the destination IP address. If an entry is found, then the packet is routed to the next hop. If there is no match, then the packet is routed to the next hop specified in the default route. If there is no default route configured, then the packet is passed to the 6200 series software to be handled appropriately.

The routing table can have entries added either statically by the administrator or dynamically via a routing protocol. The host table can have entries added either statically by the administrator or dynamically via ARP.

### NOTICE

FASTPATH supports the Border Gateway Protocol (BGP). BGP is available as a separate module and might not be available on all platforms. The BGP features can be configured only by using the CLI. No web-based administrative pages are available for BGP configuration.

## 6.1 Configuring ARP

The ARP protocol associates a Layer 2 MAC address with a Layer 3 IPv4 address. FASTPATH software features both dynamic and manual ARP configuration. With manual ARP configuration, you can statically add entries into the ARP table.

ARP is a necessary part of the Internet protocol (IP) and is used to translate an IP address to a media (MAC) address, defined by a local area network (LAN) such as Ethernet. A station needing to send an IP packet must learn the MAC address of the IP destination, or of the next hop router, if the destination is not on the same subnet. This is achieved by broadcasting an ARP request packet, to which the intended recipient responds by unicasting an ARP reply containing its MAC address. Once learned, the MAC address is used in the destination address field of the Layer 2 header prepended to the IP packet.

The ARP cache is a table maintained locally in each station on a network. ARP cache entries are learned by examining the source information in the ARP packet payload fields, regardless of whether it is an ARP request or response. Thus, when an ARP request is broadcast to all stations on a LAN segment or virtual LAN (VLAN), every recipient has the opportunity to store the sender's IP and MAC address in their respective ARP cache. The ARP response, being unicast, is normally seen only by the requester, who stores the sender information in its ARP cache. Newer information always replaces existing content in the ARP cache.

The number of supported ARP entries is platform-dependent.

Devices can be moved in a network, which means the IP address that was at one time associated with a certain MAC address is now found using a different MAC, or may have disappeared from the network altogether (i.e., it has been reconfigured, disconnected, or powered off). This leads to stale information in the ARP cache unless entries are updated in reaction to new information seen on the network, periodically refreshed to determine if an address still exists, or removed from the cache if the entry has not been identified as a sender of an ARP packet during the course of an age-out interval, usually specified via configuration.

### 6.1.1 ARP Create

Use the ARP Create page to add an entry to the Address Resolution Protocol table.

To display the page, click Routing > ARP Table > Summary in the navigation menu.

Figure 311: ARP Table

IP Address	MAC Address	Interface	Type	Age
Table is Empty				

The ARP Table displays at the bottom of the page, and contains the following fields:

Use the buttons to perform the following tasks:

- To add a static ARP entry, click Add. The Add Static ARP Entry dialog box opens. Specify the new entry information in the available fields.
- To delete one or more ARP entries, select each entry to delete and click Remove. Note that ARP entries designated as Local cannot be removed.

**Table 293: ARP Create Fields**

Field	Description
IP Address	The IP address of a network host on a subnet attached to one of the device's routing interfaces. When adding a static ARP entry, specify the IP address for the entry after you click Add.
MAC Address	The unicast MAC address (hardware address) associated with the network host. When adding a static ARP entry, specify the MAC address to associate with the IP address in the entry.
Interface	The routing interface associated with the ARP entry. The network host is associated with the device through this interface.
Type	The ARP entry type: <ul style="list-style-type: none"> <li>• Dynamic – An ARP entry that has been learned by the router</li> <li>• Gateway – A dynamic ARP entry that has the IP address of a routing interface</li> <li>• Local – An ARP entry associated with the MAC address of a routing interface on the device</li> <li>• Static – An ARP entry configured by the user</li> </ul>
Age	The age of the entry since it was last learned or refreshed. This value is specified for Dynamic or Gateway entries only (it is left blank for all other entry types).

After you enter an IP address and the associated MAC address, click Submit to apply the changes to the system and create the entry in the ARP table.

### 6.1.2 ARP Table Configuration

Use this page to change the configuration parameters for the Address Resolution Protocol Table. You can also use this screen to display the contents of the table.

To display the page, click Routing > ARP Table> Configuration in the navigation menu.

**Figure 312: ARP Table Configuration**

The screenshot shows the 'ARP Table Configuration' page with a navigation bar at the top containing 'Summary', 'Configuration', and 'Statistics'. The main content area has a red header with the title 'ARP Table Configuration' and a help icon. Below the header is a table of configuration parameters:

Age Time (Seconds)	1200	(15 to 21600)
Response Time (Seconds)	1	(1 to 10)
Retries	4	(0 to 10)
Cache Size	6144	(384 to 6144)
Dynamic Renew	<input type="checkbox"/>	

At the bottom of the configuration area are three buttons: 'Submit', 'Refresh', and 'Cancel'.

**Table 294: ARP Table Configuration Fields**

Field	Description
Age Time	The amount of time, in seconds, that a dynamic ARP entry remains in the ARP table before aging out.
Response Time	The amount of time, in seconds, that the device waits for an ARP response to an ARP request that it sends.
Retries	The maximum number of times an ARP request will be retried after an ARP response is not received. The number includes the initial ARP request.
Cache Size	The maximum number of entries allowed in the ARP table. This number includes all static and dynamic ARP entries.
Dynamic Renew	When selected, this option allows the ARP component to automatically attempt to renew dynamic ARP entries when they age out.

If you make any changes to the page, click Submit to apply the changes to the system.

## 6.2 Configuring Global IP Settings

The Routing > IP folder contains links to web pages that configure and display IP routing data.

### 6.2.1 Configuration

Use the Configuration page to configure global routing settings on the device. Routing provides a means of transmitting IP packets between subnets on the network. Routing configuration is necessary only if the device is used as a Layer 3 device that routes packets between subnets. If the device is used as a Layer 2 device that handles switching only, it typically connects to an external Layer 3 device that handles the routing functions; therefore, routing configuration is not required on the Layer 2 device.

To display the page, click Routing > IP > Configuration in the navigation menu.



Figure 313: Configuration

Routing IP Configuration	
Routing Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
ICMP Echo Replies	<input checked="" type="checkbox"/>
ICMP Redirects	<input checked="" type="checkbox"/>
ICMP Rate Limit Interval	1000 (0 to 2147483647)
ICMP Rate Limit Burst Size	100 (1 to 200)
Static Route Preference	1 (1 to 255)
Local Route Preference	0
Maximum Next Hops	4
Maximum Routes	16352
Global Default Gateway	<input type="text"/> <input type="button" value="edit"/> <input type="button" value="power"/>
ECMP Resilient Hashing	Enable <input type="button" value="edit"/>
ECMP Dynamic Hashing	Disabled <input type="button" value="edit"/>
Operational ECMP Dynamic Hashing	Disabled

Table 295: Configuration Fields

Field	Description
Routing Mode	The administrative mode of routing on the device. The options are as follows: <ul style="list-style-type: none"> <li>Enable – The device can act as a Layer 3 device by routing packets between interfaces configured for IP routing.</li> <li>Disable – The device acts as a Layer 2 bridge and switches traffic between interfaces. The device does not perform any internetwork routing.</li> </ul>
ICMP Echo Replies	Select the <code>Enable</code> or <code>Disable</code> check box. If you select <code>Enable</code> , then only the router can send ECHO replies. By default, ICMP Echo Replies are sent for echo requests.
ICMP Redirects	Select this option to allow the device to send ICMP Redirect messages to hosts. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.
ICMP Rate Limit Interval	To control the ICMP error packets, you can specify the number of ICMP error packets that are allowed per burst interval. By default, the rate limit is 100 packets per second, i.e. the burst interval is 1000 milliseconds. To disable ICMP rate limiting, set this field to zero. The valid rate interval range is 0 to 2147483647 milliseconds.
ICMP Rate Limit Burst Size	To control the ICMP error packets, you can specify the number of ICMP error packets that are allowed per burst interval. By default, the burst size is 100 packets. When the burst interval is zero, then configuring this field is not a valid option. The valid burst size range is 1 to 200.

Table 295: Configuration Fields (Continued)

Field	Description
Static Route Preference	The default distance (preference) for static routes. Lower route-distance values are preferred when determining the best route. The value configured for Static Route Preference is used when using the CLI to configure a static route and no preference is specified. Changing the Static Route Preference does not update the preference of existing static routes.
Local Route Preference	The default distance (preference) for local routes.
Maximum Next Hops	The maximum number of hops supported by the switch. This is a read-only value.
Maximum Routes	The maximum number of routes (routing table size) supported by the switch. This is a read-only value.
Global Default Gateway	The IP address of the default gateway for the device. If the destination IP address in a packet does not match any routes in the routing table, the packet is sent to the default gateway. The gateway specified in this field is more preferred than a default gateway learned from a DHCP server. Use the icons associated with this field to perform the following tasks: <ul style="list-style-type: none"> <li>To configure the default gateway, click the Edit icon and specify the IP address of the default gateway in the available field.</li> <li>To reset the IP address of the default gateway to the factory default value, click the Reset icon associated with this field.</li> </ul>
ECMP Resilient Hashing	Use this field to configure ECMP Resilient Hashing mode. Resilient hashing is enabled by default. Click the Edit icon next to the field to open a modal page to enable or disable resilient hashing. See <a href="#">Section 6.2.1.1: "Dynamic Load Balancing"</a> .
ECMP Dynamic Hashing	Use this field to view and configure the ECMP Dynamic Hashing mode.
Operational ECMP Dynamic Hashing	Use this field to view the Operational ECMP Dynamic Hashing mode.

If you make any changes to the page, click Submit to apply the changes to the system.

### 6.2.1.1 Dynamic Load Balancing

Dynamic Load Balancing (DLB) is a load balancing feature that works across LAG, HiGig trunk (stack links), and ECMP. The Dynamic Load Balancing mechanism improves upon a hash-based load balancing scheme by performing the following:

- Consider the state/loading of aggregate members when assigning a new flow.
- Accounting for existing flow assignment when changing loading across members.
- Identifying instances where active flows can be moved to another aggregate member while avoiding re-ordering.

DLB only works for known unicast traffic. It should not be used for multicast, broadcast, unknown unicast, and mirrored packets.

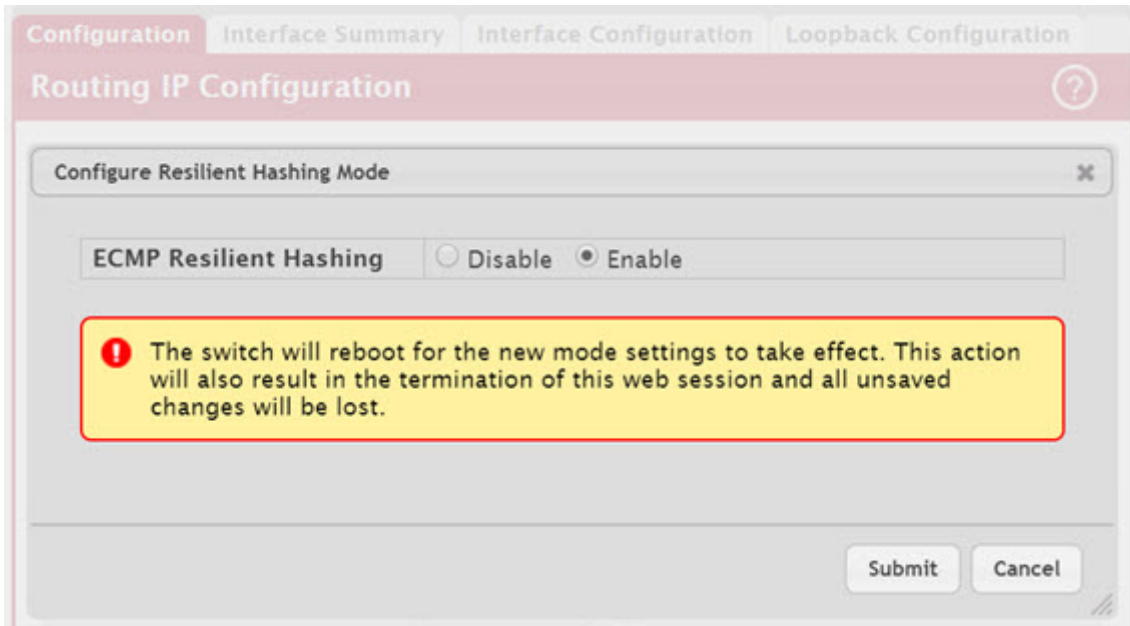
The level of DLB support varies from ASIC to ASIC. The list of supported ASICs and their level of DLB support in FASTPATH is shown in the following table.

Table 296: ASIC Dynamic Load Balancing Support

ASIC	Applications
BCM56970	LAG or HiGig trunk AND ECMP
BCM56870	LAG or HiGig trunk AND ECMP
BCM56771	LAG or HiGig trunk AND ECMP
BCM56850	Only on HiGig trunk
BCM56640	LAG, HiGig trunk and ECMP
BCM56547, BCM56340	LAG, HiGig trunk

Use the following modal page to to enable or disable resilient hashing.

Figure 314: Configure Resilient Hashing Mode



Resilient hashing is enabled by default. If the mode is changed, the switch reboots for the new mode settings to take effect. This action also results in the termination of the web session. When resilient hashing is enabled, dynamic hashing cannot be configured.

### 6.2.2 Interface Summary

This page shows summary information about the routing configuration for all interfaces. To view additional routing configuration information for an interface, select the interface with the settings to view and click Details.

To display the page, click Routing > IP > Interface Summary in the navigation menu.

Figure 315: Interface Summary

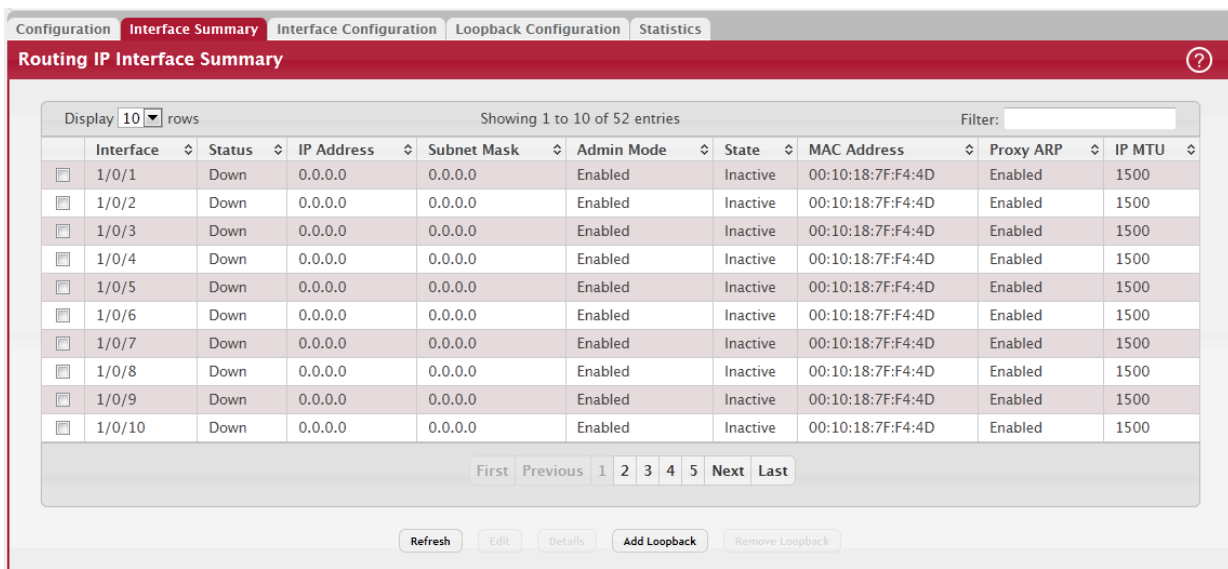


Table 297: Interface Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When viewing details about the routing settings for an interface, this field identifies the interface being viewed.
Status	Indicates whether the interface is capable of routing IP packets (Up) or cannot route packets (Down). For the status to be Up, the routing mode and administrative mode for the interface must be enabled. Additionally, the interface must have an IP address and be physically up (active link).
IP Address	The IP address of the interface.
Subnet Mask	The IP subnet mask for the interface (also known as the network mask or netmask). It defines the portion of the interface's IP address that is used to identify the attached network.
Admin Mode	The administrative mode of the interface, which is either Enabled or Disabled.
State	The state of the interface, which is either Active or Inactive. An interface is considered active if the link is up, and the interface is in a forwarding state.
MAC Address	The burned-in physical address of the interface. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.
Proxy ARP	Indicates whether proxy ARP is enabled or disabled on the interface. When proxy ARP is enabled, the interface can respond to an ARP request for a host other than itself. An interface can act as an ARP proxy if it is aware of the destination and can route packets to the intended host, which is on a different subnet than the host that sent the ARP request.
IP MTU	The largest IP packet size the interface can transmit, in bytes. The IP Maximum Transmission Unit (MTU) is the maximum frame size minus the length of the Layer 2 header.

After you click Details, the Details window opens and displays detailed routing information for the selected interface. The following information describes the fields in this window that are not displayed on the summary page.

Table 298: Interface Summary Details Fields

Field	Description
Routing Mode	Indicates whether routing is administratively enabled or disabled on the interface.
Link Speed Data Rate	The physical link data rate of the interface.
IP Address Configuration Method	The source of the IP address, which is one of the following: <ul style="list-style-type: none"> <li>None – The interface does not have an IP address.</li> <li>Manual – The IP address has been statically configured by an administrator.</li> <li>DHCP – The IP address has been learned dynamically through DHCP. If the method is DHCP but the interface does not have an IP address, the interface is unable to acquire an address from a network DHCP server.</li> </ul>
Bandwidth	The configured bandwidth on this interface. This setting communicates the speed of the interface to higher-level protocols.
Encapsulation Type	The link layer encapsulation type for packets transmitted from the interface, which can be either Ethernet or SNAP.
Forward Net Directed Broadcasts	Indicates how the interface handles network-directed broadcast packets. A network-directed broadcast is a broadcast directed to a specific subnet. The possible values are as follows: <ul style="list-style-type: none"> <li>Enabled – Network directed broadcasts are forwarded.</li> <li>Disabled – Network directed broadcasts are dropped.</li> </ul>

Table 298: Interface Summary Details Fields (Continued)

Field	Description
Local Proxy ARP	Indicates whether local proxy ARP is enabled or disabled on the interface. When local proxy ARP is enabled, the interface can respond to an ARP request for a host other than itself. Unlike proxy ARP, local proxy ARP allows the interface to respond to ARP requests for a host that is on the same subnet as the host that sent the ARP request. This feature is useful when a host is not permitted to reply to an ARP request from another host in the same subnet, for example when using the protected ports feature.
Destination Unreachables	Indicates whether the interface is allowed to send ICMP Destination Unreachable message to a host if the intended destination cannot be reached for some reason. If the status of this field is Disabled, this interface will not send ICMP Destination Unreachable messages to inform the host about the error in reaching the intended destination.
ICMP Redirects	Indicates whether the interface is allowed to send ICMP Redirect messages. The device sends an ICMP Redirect message on an interface only if ICMP Redirects are enabled both globally and on the interface. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.

### 6.2.3 Interface Configuration

Use the Interface Configuration page to configure the IP routing settings for each interface.

To display the page, click Routing > IP > Interface Configuration in the navigation menu.

Figure 316: Interface Configuration

The screenshot shows the 'Routing IP Interface Configuration' page. At the top, there are navigation tabs: Configuration, Interface Summary, Interface Configuration (active), Loopback Configuration, and Statistics. Below the tabs is a red header with the title 'Routing IP Interface Configuration' and a help icon. The main configuration area is a table with the following fields:

- Type:  VLAN  Interface
- VLAN: VLAN 1
- Interface: 1/0/1
- Status: Down
- Routing Mode:  Disable  Enable
- Admin Mode:  Disable  Enable
- State: (empty)
- Link Speed Data Rate: (empty)
- IP Address Configuration Method:  None  Manual  DHCP
- DHCP Client Identifier:
- IP Address: (input field) (x.x.x.x)
- Subnet Mask: (input field) (x.x.x.x)
- MAC Address: (empty)
- IP MTU: 1500 (68 to 1500) ⚙
- Bandwidth: 10000 (1 to 10000000)
- Encapsulation Type:  Ethernet  SNAP
- Forward Net Directed Broadcasts:
- Proxy ARP:
- Local Proxy ARP:
- Destination Unreachables:
- ICMP Redirects:

Below the main configuration area is a table for 'Secondary IP Address' and 'Secondary Subnet Mask'. The table is empty, with the text 'Table is Empty' displayed. At the bottom of the page are three buttons: Submit, Refresh, and Cancel.

Table 299: Interface Configuration Fields

Field	Description
Interface	The menu contains all interfaces that can be configured for routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page.
Status	Indicates whether the interface is currently capable of routing IP packets (Up) or cannot route packets (Down). For the status to be Up, the routing mode and administrative mode for the interface must be enabled. Additionally, the interface must have an IP address and be physically up (active link).
Routing Mode	The administrative mode of IP routing on the interface.
Admin Mode	The administrative mode of the interface. If an interface is administratively disabled, it cannot forward traffic.
State	The state of the interface, which is either Active or Inactive. An interface is considered active if the link is up, and the interface is in a forwarding state.
Link Speed Data Rate	The physical link data rate of the interface.
IP Address Configuration Method	The method to use for configuring an IP address on the interface, which can be one of the following: <ul style="list-style-type: none"> <li>• None – No address is to be configured.</li> <li>• Manual – The address is to be statically configured. When this option is selected you can specify the IP address and subnet mask in the available fields.</li> <li>• DHCP – The interface will attempt to acquire an IP address from a network DHCP server.</li> </ul>
IP Address	The IP address of the interface. This field can be configured only when the selected IP Address Configuration Method is Manual. If the method is DHCP, the interface attempts to lease an IP address from a DHCP server on the network, and the IP address appears in this field (read-only) after it is acquired. If this field is blank, the IP Address Configuration Method might be None, or the method might be DHCP and the interface is unable to lease an address.
Subnet Mask	The IP subnet mask for the interface (also known as the network mask or netmask). This field can be configured only when the selected IP Address Configuration Method is Manual.
MAC Address	The burned-in physical address of the interface. The format is six two-digit hexadecimal numbers separated by colons, for example 00:06:29:32:81:40.
IP MTU	The largest IP packet size the interface can transmit, in bytes. The IP Maximum Transmission Unit (MTU) is the maximum frame size minus the length of the Layer 2 header.
Bandwidth	The configured bandwidth on this interface. This setting communicates the speed of the interface to higher-level protocols.
Encapsulation Type	The link layer encapsulation type for packets transmitted from the interface, which can be either Ethernet or SNAP.
Forward Net Directed Broadcasts	Determines how the interface handles network-directed broadcast packets. A network-directed broadcast is a broadcast directed to a specific subnet. If this option is selected, network directed broadcasts are forwarded. If this option is clear, network directed broadcasts are dropped.
Proxy ARP	When this option is selected, proxy ARP is enabled, and the interface can respond to an ARP request for a host other than itself. An interface can act as an ARP proxy if it is aware of the destination and can route packets to the intended host, which is on a different subnet than the host that sent the ARP request.
Local Proxy ARP	When this option is selected, local proxy ARP is enabled, and the interface can respond to an ARP request for a host other than itself. Unlike proxy ARP, local proxy ARP allows the interface to respond to ARP requests for a host that is on the same subnet as the host that sent the ARP request. This feature is useful when a host is not permitted to reply to an ARP request from another host in the same subnet, for example when using the protected ports feature.

Table 299: Interface Configuration Fields (Continued)

Field	Description
Destination Unreachables	When this option is selected, the interface is allowed to send ICMP Destination Unreachable message to a host if the intended destination cannot be reached for some reason. If this option is clear, the interface will not send ICMP Destination Unreachable messages to inform the host about the error in reaching the intended destination.
ICMP Redirects	When this option is selected, the interface is allowed to send ICMP Redirect messages. The device sends an ICMP Redirect message on an interface only if ICMP Redirects are enabled both globally and on the interface. An ICMP Redirect message notifies a host when a better route to a particular destination is available on the network segment.
Secondary IP Address	To add a secondary IP address on the interface, click the + (plus) symbol in the header row and enter the address in the appropriate field in the Secondary IP Address Configuration window. You can add one or more secondary IP addresses to an interface only if the interface already has a primary IP address. To remove a configured secondary IP address, click the – (minus) symbol associated with the entry to remove. To remove all configured secondary IP addresses, click the – (minus) symbol in the header row.
Secondary Subnet Mask	The subnet mask associated with the secondary IP address. You configure this field in the Secondary IP Address Configuration window.

## 6.2.4 IP Loopback Configuration

Use this page to configure the IP routing settings for each loopback interface.

To display the IP Loopback Configuration page, click Routing > IP > Loopback Configuration in the navigation menu.

Figure 317: IP Loopback Configuration

Table 300: IP Loopback Configuration Fields

Field	Description
Interface	The menu contains all loopback interfaces that can be configured for routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page.
IP Address	The IP address of the loopback interface.
Subnet Mask	The IP subnet mask for the interface (also known as the network mask or netmask).
Secondary IP Address	To add a secondary IP address on the interface, click the + (plus) symbol in the header row and enter the address in the appropriate field in the Secondary IP Address Configuration window. You can add one or more secondary IP addresses to an interface only if the interface already has a primary IP address. To remove a configured secondary IP address, click the – (minus) symbol associated with the entry to remove. To remove all configured secondary IP addresses, click the – (minus) symbol in the header row.
Secondary Subnet Mask	The subnet mask associated with the secondary IP address. You configure this field in the Secondary IP Address Configuration window.

Click Refresh to update the information on the screen.

## 6.2.5 IP Statistics

The statistics reported on the IP Statistics page are as specified in RFC 1213.

To display the page, click Routing > IP > Statistics in the navigation menu. A partial page is shown.

Figure 318: "IP Statistics," on page 354 does not show all of the fields on the page.

### NOTICE

Figure 318: IP Statistics

Field	Value
IpInReceives	9976
IpInHdrErrors	0
IpAddrErrors	25
IpFwdDatagrams	0
IpInUnknownProtos	0
IpInDiscards	0
IpInDelivers	9951
IpOutRequests	5585
IpOutDiscards	0
IpOutNoRoutes	0
IpReasmTimeout	0
IpReasmReqds	0
IpReasmOKs	0
IpReasmFails	0
IpFragOKs	0

Table 301: IP Statistics Fields

Field	Description
IpInReceives	The total number of input datagrams received from interfaces, including those received in error.
IpInHdrErrors	The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.
IpInAddrErrors	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (e.g., 0.0.0.0) and addresses of unsupported Classes (e.g., Class E). For entities which are not IP Gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
IpForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP Gateways, this counter includes only those packets which were Source-Routed via this entity, and the Source-Route option processing was successful.
IpInUnknownProtos	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
IpInDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
IpInDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
IpOutRequests	The total number of IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission. Note that this counter does not include any datagrams counted in ipForwDatagrams.



Table 301: IP Statistics Fields (Continued)

Field	Description
IpOutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (e.g., for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
IpOutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this 'no-route' criterion. Note that this includes any datagrams which a host cannot route because all of its default gateways are down.
IpReasmTimeout	The maximum number of seconds which received fragments are held while they are awaiting reassembly at this entity.
IpReasmReqds	The number of IP fragments received which needed to be reassembled at this entity.
IpReasmOKs	The number of IP datagrams successfully re-assembled.
IpReasmFails	The number of failures detected by the IP re-assembly algorithm (for whatever reason: timed out, errors, etc.). Note that this is not necessarily a count of discarded IP fragments since some algorithms can lose track of the number of fragments by combining them as they are received.
IpFragOKs	The number of IP datagrams that have been successfully fragmented at this entity.
IpFragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, e.g., because their Don't Fragment flag was set.
IpFragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
IpRoutingDiscards	The number of routing entries which were chosen to be discarded even though they are valid. One possible reason for discarding such an entry could be to free-up buffer space for other routing entries.
IcmpInMsgs	The total number of ICMP messages which the entity received. Note that this counter includes all those counted by icmpInErrors.
IcmpInErrors	The number of ICMP messages which the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).
IcmpInDestUnreachs	The number of ICMP Destination Unreachable messages received.
IcmpInTimeExcds	The number of ICMP Time Exceeded messages received.
IcmpInParmProbs	The number of ICMP Parameter Problem messages received.
IcmpInSrcQuenchs	The number of ICMP Source Quench messages received.
IcmpInRedirects	The number of ICMP Redirect messages received.
IcmpInEchos	The number of ICMP Echo (request) messages received.
IcmpInEchoReps	The number of ICMP Echo Reply messages received.
IcmpInTimestamps	The number of ICMP Timestamp (request) messages received.
IcmpInTimestampReps	The number of ICMP Timestamp Reply messages received.
IcmpInAddrMasks	The number of ICMP Address Mask Request messages received.
IcmpInAddrMaskReps	The number of ICMP Address Mask Reply messages received.
IcmpOutMsgs	The total number of ICMP messages which this entity attempted to send. Note that this counter includes all those counted by icmpOutErrors.
IcmpOutErrors	The number of ICMP messages which this entity did not send due to problems discovered within ICMP such as a lack of buffers. This value should not include errors discovered outside the ICMP layer such as the inability of IP to route the resultant datagram. In some implementations there may be no types of error which contribute to this counter's value.
IcmpOutDestUnreachs	The number of ICMP Destination Unreachable messages sent.
IcmpOutTimeExcds	The number of ICMP Time Exceeded messages sent.
IcmpOutParmProbs	The number of ICMP Parameter Problem messages sent.

Table 301: IP Statistics Fields (Continued)

Field	Description
IcmpOutSrcQuenchs	The number of ICMP Source Quench messages sent.
IcmpOutRedirects	The number of ICMP Redirect messages sent. For a host, this object is always zero, since hosts do not send redirects.
IcmpOutEchos	The number of ICMP Echo (request) messages sent.
IcmpOutEchoReps	The number of ICMP Echo Reply messages sent.
IcmpOutTimestamps	The number of ICMP Timestamp (request) messages.
IcmpOutTimestampReps	The number of ICMP Timestamp Reply messages sent.
IcmpOutAddrMasks	The number of ICMP Address Mask Request messages sent.
IcmpOutAddrMaskReps	The number of ICMP Address Mask Reply messages sent.

## 6.3 Router

The Routing > Router menu contains links to web pages that configure and display route tables.

### 6.3.1 Route Table

The route table manager collects routes from multiple sources: static routes, RIP routes, OSPF routes, BGP routes, local routes. The route table manager may learn multiple routes to the same destination from multiple sources. The route table lists all routes. The best routes table displays only the most preferred route to each destination.

To display the page, click Routing > Router > Route Table in the navigation menu.

Figure 319: Route Table

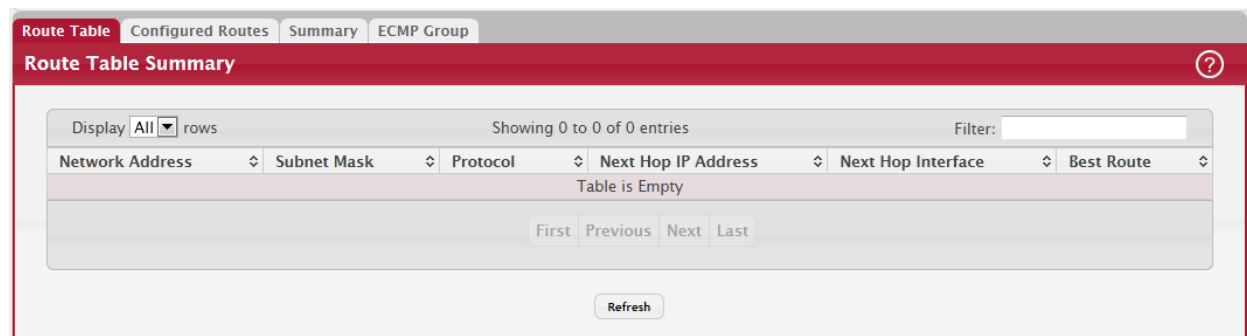


Table 302: Route Table Fields

Field	Description
Network Address	The IP route prefix for the destination.
Subnet Mask	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.

**Table 302: Route Table Fields (Continued)**

Field	Description
Protocol	This field tells which protocol created the specified route. The possibilities are one of the following: <ul style="list-style-type: none"> <li>Local</li> <li>Static</li> <li>Default</li> <li>OSPF Intra</li> <li>OSPF Inter</li> <li>OSPF Type-1</li> <li>OSPF Type-2</li> <li>RIP</li> </ul>
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly-attached network.
Next Hop Interface	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null.
Best Route	Indicates whether the route is the preferred route to the network. If the field is blank, a better route to the same network exists in the routing table.

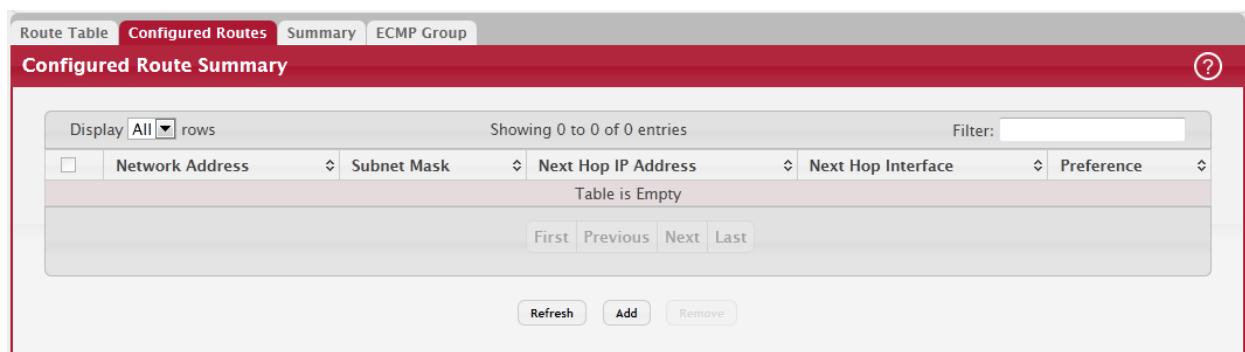
Click Refresh to update the information on the screen.

### 6.3.2 Configured Routes

Use the Configured Routes page to create and display static routes.

To display the page, click Routing > Router > Configured Routes in the navigation menu.

**Figure 320: Configured Routes**



Use the buttons to perform the following tasks:

- To configure a route, click Add and specify the desired settings in the available fields.
- To remove a configured route, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

**Table 303: Configured Routes Fields**

Field	Description
Network Address	The IP route prefix for the destination.
Subnet Mask	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.
Next Hop IP	The next hop router address to use when forwarding traffic to the destination.

Table 303: Configured Routes Fields (Continued)

Field	Description
Next Hop Unit/Slot Port	The outgoing interface to use when forwarding traffic to the destination. For static reject routes it would be Null0.
Preference	The preferences configured for the added routes.

### 6.3.2.1 Adding a Static Route

1. Open the Configured Routes page.
2. Click Add.

The Router Route Entry Configuration page displays:

The screenshot shows a dialog box titled "Add Route" with a close button (X) in the top right corner. Inside the dialog, there is a "Route Type" section with three radio buttons: "Default", "Static" (which is selected), and "Static Reject". Below this are five input fields: "Network Address" with a placeholder "(x.x.x.x)", "Subnet Mask" with a placeholder "(x.x.x.x)", "Next Hop IP Address" with a placeholder "(x.x.x.x)", and "Preference" with a value of "1" and a range "(1 to 255)". At the bottom right of the dialog are "Submit" and "Cancel" buttons.

3. Next to Route Type, select Default route, Static or Static Reject from the menu.

**Default:** Enter the default gateway address in the Next Hop IP Address field.

**Static:** Enter values for Network Address, Subnet Mask, Next Hop IP Address, and Preference.

**Static Reject:** Packets to these destinations will be dropped.

#### **NOTICE**

The route type you select determines the fields available on the page. Some of the fields that Table 304, "Route Entry Create Fields," on page 358 describes are not available when configuring certain types of routes.

Table 304: Route Entry Create Fields

Field	Description
Network Address	Specify the IP route prefix for the destination from the drop-down menu. To create a route, a valid routing interface must exist and the next hop IP Address must be on the same network as the routing interface. Routing interfaces are created on the IP Interface Configuration page. Valid next hop IP Addresses can be viewed on the Route Table page.
Subnet Mask	Also referred to as the subnet/network mask, this indicates the portion of the IP interface address that identifies the attached network.
Protocol	This field tells which protocol created the specified route. Possible values are: <ul style="list-style-type: none"> <li>• Local</li> <li>• Static</li> <li>• Default</li> <li>• OSPF Intra</li> <li>• OSPF Inter</li> <li>• OSPF Type-1</li> <li>• OSPF Type-2</li> <li>• RIP</li> </ul>

Table 304: Route Entry Create Fields (Continued)

Field	Description
Next Hop Slot/Port	The outgoing router interface to use when forwarding traffic to the destination.
Next Hop IP Address	The outgoing router IP address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IP address of the local interface for a directly attached network. When creating a route, the next hop IP must be on the same network as the routing interface. Valid next hop IP Addresses can be seen on the 'Route Table' page.
Metric	Administrative cost of the path to the destination. If no value is entered, default is 1. The range is 0 to 255. This field is present only when creating a static route.
Preference	Specifies a preference value for the configured next hop.
Route Type	Specifies whether the route is to be a Default route or a Static route.

4. Click Submit. The new route is added, and you are returned to the Configured Routes page.

### 6.3.2.2 Deleting a Route

Click Delete to remove a configured route.

### 6.3.3 Summary

The Summary page displays summary information about the entries in the IP routing table.

To display the page, click Routing > Router > Summary in the navigation menu.

Figure 321: Summary

Route Types	
Connected Routes	0
Static Routes	0
RIP Routes	0
BGP Routes	0
External	0
Internal	0
Local	0
OSPF Routes	0
Intra Area Routes	0
Inter Area Routes	0
External Type-1 Routes	0
External Type-2 Routes	0
Reject Routes	0
Total Routes	0

Route Table Counters	
Best Routes (High)	0 (0)
Alternate Routes	0
Route Adds	0
Route Modifies	0
Route Deletes	0
Unresolved Route Adds	0
Invalid Route Adds	0
Failed Route Adds	0
Reserved Locals	0
Unique Next Hops (High)	0 (0)
Next Hop Groups (High)	0 (0)
ECMP Groups (High)	0 (0)
ECMP Routes	0
Truncated ECMP Routes	0
ECMP Retries	0

Table 305: Summary Fields

Field	Description
Connected Routes	The total number of connected routes in the IP routing table.
Static Routes	The total number of static routes in the IP routing table.
RIP Routes	The total number of routes installed by the RIP protocol.
BGP Routes	The total number of routes installed by the BGP protocol.
External	The total number of external routes installed by the BGP protocol.
Internal	The total number of internal routes installed by the BGP protocol.
Local	The total number of local routes installed by the BGP protocol.

Table 305: Summary Fields (Continued)

Field	Description
OSPF Routes	The total number of routes installed by the OSPF protocol.
Intra Area Routes	The total number of intra-area routes installed by the OSPF protocol.
Inter Area Routes	The total number of inter-area routes installed by the OSPF protocol.
External Type-1 Routes	The total number of external type-1 routes installed by the OSPF protocol.
External Type-2 Routes	The total number of external type-2 routes installed by the OSPF protocol.
Reject Routes	The total number of reject routes installed by all protocols.
Total Routes	The total number of routes in the routing table.
Best Routes (High)	The number of best routes currently in the routing table. This number only counts the best route to each destination.
Alternate Routes	The number of alternate routes currently in the routing table. An alternate route is a route that was not selected as the best route to its destination.
Route Adds	The number of routes that have been added to the routing table.
Route Modifies	The number of routes that have been changed after they were initially added to the routing table.
Route Deletes	The number of routes that have been deleted from the routing table.
Unresolved Route Adds	The number of route adds that failed because none of the route's next hops were on a local subnet. Note that static routes can fail to be added to the routing table at startup because the routing interfaces are not yet up. This counter gets incremented in this case. The static routes are added to the routing table when the routing interfaces come up.
Invalid Route Adds	The number of routes that failed to be added to the routing table because the route was invalid. A log message is written for each of these failures.
Failed Route Adds	The number of routes that failed to be added to the routing table because of a resource limitation in the routing table.
Reserved Locals	The number of routing table entries reserved for a local subnet on a routing interface that is down. Space for local routes is always reserved so that local routes can be installed when a routing interface bounces.
Unique Next Hops (High)	The number of distinct next hops used among all routes currently in the routing table. These include local interfaces for local routes and neighbors for indirect routes.
Next Hop Groups (High)	The current number of next hop groups in use by one or more routes. Each next hop group includes one or more next hops.
ECMP Groups (High)	The number of next hop groups with multiple next hops.
ECMP Routes	The number of routes with multiple next hops currently in the routing table.
Truncated ECMP Routes	The number of ECMP routes that are currently installed in the forwarding table with just one next hop. The forwarding table may limit the number of ECMP routes or the number of ECMP groups. When an ECMP route cannot be installed because such a limit is reached, the route is installed with a single next hop.
ECMP Retries	The number of ECMP routes that have been installed in the forwarding table after initially being installed with a single next hop.
Routes with n Next Hops	The current number of routes with each number of next hops.
Clear Counters	This button resets to zero IPv4 routing table counters reported in this page. This only resets event counters. Counters that report the current state of the routing table, such as the number of routes of each type, are not reset.

Click Refresh to update the information on the screen.

## 6.3.4 ECMP Group

The ECMP Groups Summary page displays all current ECMP groups in the IPv4 routing table. An ECMP group is a set of two or more next hops used in one or more routes.

To display the page, click Routing > Router > ECMP Group in the navigation menu.

Figure 322: ECMP Groups Summary

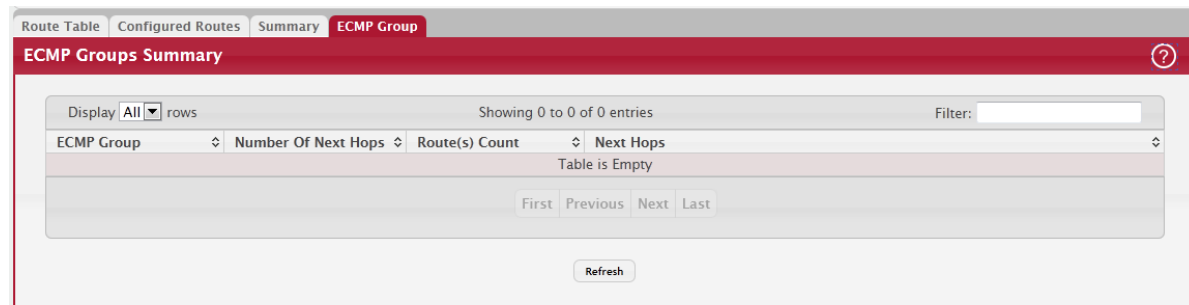


Table 306: ECMP Groups Summary Fields

Field	Description
ECMP Group	The ECMP group number associated with the rest of the data in the row. The device assigns an arbitrary ECMP group number from 1 to n to identify the ECMP group.
Number Of Next Hops	The number of next hops in the group.
Route(s) Count	The number of routes that use the set of next hops.
Next Hops	The IPv4 address and outgoing interface of each next hop in the group.

Click Refresh to update the information on the screen.

## 6.4 Configuring IPv6 Settings

The Routing > IPv6 folder contains links to web pages that configure and display IP routing data.

### 6.4.1 IPv6 Global Configuration

Use this page to configure global IPv6 routing settings on the device. IPv6 routing provides a means of transmitting IPv6 packets between subnets on the network. IPv6 routing configuration is necessary only if the device is used as a Layer 3 device that routes IPv6 packets between subnets. IPv6 is the next generation of the Internet Protocol. With 128-bit addresses, versus 32-bit addresses for IPv4, IPv6 solves the address depletion issues seen with IPv4 and removes the requirement for Network Address Translation (NAT), which is used in IPv4 networks to reduce the number of globally unique IP addresses required for a given network.

To display the IPv6 Global Configuration page, click Routing > IPv6 > Configuration in the navigation menu.



Figure 323: Configuration

Field	Value	Range
IPv6 Unicast Routing Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable	
IPv6 Neighbors Dynamic Renew	<input type="checkbox"/>	
IPv6 Hop Limit	64	(1 to 255)
IPv6 Unresolved Packets Rate Limit (pps)	1024	(50 to 1024)
NUD Maximum Unicast Solicits	3	(3 to 10)
NUD Maximum Multicast Solicits	3	(3 to 255)
NUD Back-off Multiple	1	(1 to 5)
ICMPv6 Rate Limit Error Interval (Msecs)	1000	(0 to 2147483647)
ICMPv6 Rate Limit Burst Size	100	(1 to 200)
Static Route Preference	1	(1 to 255)
Local Route Preference	0	

Buttons: Submit, Refresh, Cancel

Table 307: Configuration Fields

Field	Description
IPv6 Unicast Routing Mode	The administrative mode of IPv6 routing on the device. The options are as follows: <ul style="list-style-type: none"> <li>Enable – The device can act as a Layer 3 device by routing IPv6 packets between interfaces configured for IPv6 routing.</li> <li>Disable – The device does not support IPv6 routing.</li> </ul>
IPv6 Neighbors Dynamic Renew	Select this option to enable dynamic renewal mode for the periodic Neighbor Unreachability Detection (NUD) run on the existing IPv6 neighbor entries in the IPv6 neighbor cache. If NUD attempts to communicate with IPv6 neighbors and no response is received after the maximum number of solicits is reached, its entry is removed from the cache.
IPv6 Hop Limit	The unicast hop count used in IPv6 packets originated by the device. This value is also included in router advertisements.
IPv6 Unresolved Packets Rate Limit	The rate in packets-per-second for the number of IPv6 data packets trapped to the CPU when the packet fails to be forwarded in the hardware due to the unresolved hardware address of the destined IPv6 node.
NUD Maximum Multicast Solicits	The maximum number of multicast neighbor solicitations sent during NUD when a neighbor is in the UNREACHABLE state.
NUD Back-off Multiple	The exponential backoff multiplier to be used in the calculation of the next timeout value for neighbor solicitation transmission during NUD following the exponential backoff algorithm.
ICMPv6 Rate Limit Error Interval	The maximum burst interval for ICMPv6 error messages transmitted by the device. The rate limit for ICMPv6 error messages is configured as a token bucket. The ICMPv6 Rate Limit Error Interval specifies how often the token bucket is initialized with tokens of the size configured in the ICMPv6 Rate Limit Burst Size field.
ICMPv6 Rate Limit Burst Size	The number of ICMPv6 error messages that can be sent during the burst interval configured in the ICMPv6 Rate Limit Error Interval field.
Static Route Preference	The default distance (preference) for static IPv6 routes. Lower route-distance values are preferred when determining the best route. The value configured for Static Route Preference is used when using the CLI to configure a static route and no preference is specified. Changing the Static Route Preference does not update the preference of existing static routes.
Local Route Preference	The default distance (preference) for local IPv6 routes.

If you make any changes to the page, click Submit to apply the changes to the system.

## 6.4.2 IPv6 Interface Summary

This page shows summary information about the IPv6 routing configuration for all interfaces.

To display the IPv6 Interface Summary page, click Routing > IPv6 > Interface Summary in the navigation menu.

Figure 324: IPv6 Interface Summary

Interface	Operational Status	IPv6 Mode	Routing Mode	Admin Mode	IPv6 Prefix	Prefix Length	State
1/0/1	Disabled	Enabled	Disabled	Enabled			
1/0/2	Disabled	Disabled	Disabled	Enabled			
1/0/3	Disabled	Disabled	Disabled	Enabled			
1/0/4	Disabled	Disabled	Disabled	Enabled			
1/0/5	Disabled	Disabled	Disabled	Enabled			
1/0/6	Disabled	Disabled	Disabled	Enabled			
1/0/7	Disabled	Disabled	Disabled	Enabled			
1/0/8	Disabled	Disabled	Disabled	Enabled			
loopback0	Disabled	Disabled	Enabled	Disabled	fe80::202:bcff:fe4d:b9e6	128	Inactive

Use the buttons to perform the following tasks:

- To edit any interface, select the interface and click Edit. You are redirected to the IPv6 Interface Configuration or IPv6 Loopback Configuration page for the selected interface.
- To view additional routing configuration information for an interface, select the interface with the settings to view and click Details.
- To add the next available loopback interface, click Add Loopback. You are redirected to the IPv6 Loopback Configuration page.

Table 308: IPv6 Interface Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When viewing details about the routing settings for an interface, this field identifies the interface being viewed.
Operational Status	The IPv6 routing operational mode on the interface. The operational mode is enabled under the conditions in the following list; otherwise, the mode is disabled: <ul style="list-style-type: none"> <li>The IPv6 mode is enabled on the interface.</li> <li>The routing mode is enabled on the interface.</li> <li>The administrative mode is enabled on the interface.</li> <li>The link is up.</li> </ul>
IPv6 Mode	The IPv6 mode on the interface. When IPv6 mode is enabled, the interface is capable of IPv6 operation without a global address. In this case, an EUI-64 based link-local address is used.
Routing Mode	The IPv6 mode on the interface. When IPv6 mode is enabled, the interface is capable of IPv6 operation without a global address. In this case, an EUI-64 based link-local address is used.
Admin Mode	The administrative mode on the interface.
IPv6 Prefix	The IPv6 routing prefix dynamically or manually configured on the interface.
Prefix Length	The number of bits used for the IPv6 prefix.
State	The state of the IPv6 address. The state is TENT if routing is disabled or Duplicate Address Detection (DAD) fails. The state is Active if the interface is active and DAD is successful. Otherwise, the state is Inactive.

Table 308: IPv6 Interface Summary Fields (Continued)

Field	Description
	After you click Details, the Details window opens and displays detailed IPv6 routing information for the selected interface. The following information describes the fields in this window that are not displayed on the summary page.
Link Local Prefix	The autoconfigured link-local address which is: <ul style="list-style-type: none"> <li>Allocated from part of the IPv6 unicast address space</li> <li>Not visible off the local link</li> <li>Not globally unique</li> </ul>
Link Local Prefix Length	The number of bits used for the prefix of the link-local IPv6 address.
Link Local Status	The status of the IPV6 link local address. The status is TENT if routing is disabled or Duplicate Address Detection (DAD) fails. The state is Active if the interface is active and DAD is successful. Otherwise, the state is Inactive.
DHCPv6 Client Mode	The administrative mode of the DHCPv6 client on the interface. An interface can acquire an IPv6 address by communicating with a network DHCPv6 server (stateful configuration). An interface can also obtain an IPv6 address through stateless address autoconfiguration or static configuration.
Stateless Address AutoConfig	The administrative mode of stateless address autoconfiguration on the interface. When enabled, the interface can configure itself by using the Neighbor Discovery Protocol.
Interface Maximum Transmit Unit	The largest IPv6 packet size the interface can transmit, in bytes. To change the MTU value, click the Edit icon to the right of the field. To reset the MTU to the default value, click the Reset icon.
Router Duplicate Address Detection Transmits	The number of duplicate address detection probes the interface transmits while doing neighbor discovery.
Router Advertisement NS Interval	The interval between router advertisements for advertised neighbor solicitations. To change the interval, click the Edit icon to the right of the field. To reset the interval to the default value, click the Reset icon.
Router Lifetime Interval	The value that is placed in the Router Lifetime field of the router advertisements sent from the interface.
Router Advertisement Reachable Time	The value that is placed in the Reachable Time field of the router advertisements The amount of time to consider a neighbor reachable after neighbor discovery confirmation.
Router Advertisement Interval	The transmission interval between router advertisements messages sent by the interface.
Router Advertisement Managed Config	The mode of the Managed Address Configuration flag in router advertisements sent from the interface. When enabled, the Managed Address Configuration flag is set in router advertisements, indicating that the interface should use stateful autoconfiguration (DHCPv6) to obtain addresses.
Router Advertisement Other Config	The mode of the Other Stateful Configuration flag in router advertisements sent from the interface. When enabled, the Other Stateful Configuration flag is set in router advertisements, indicating that the interface should use stateful autoconfiguration for information other than addresses.
Router Advertisement Suppress	The mode of router advertisement transmission suppression on an interface. When enabled, the interface does not transmit router advertisements.
IPv6 Destination Unreachable Messages	The mode for ICMPv6 Destination Unreachable messages. When enabled, the interface can send ICMPv6 Destination Unreachable messages back to the source device if a datagram is undeliverable.
IPv6 Hop Limit Unspecified	The mode that controls whether the interface transmits the hop limit value as 0 in Router Advertisements (Enabled) or transmits the global hop limit value (Disabled).

Click Refresh to update the information on the screen.

### 6.4.3 IPv6 Interface Configuration

Use this page to configure the IPv6 routing settings for each non-loopback interface.

To display the IPv6 Interface Configuration page, click Routing > IPv6 > Interface Configuration in the navigation menu.

Figure 325: IPv6 Interface Configuration

Table 309: IPv6 Interface Configuration Fields

Field	Description
Interface	The menu contains all non-loopback interfaces that can be configured for IPv6 routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page.
Operational Status	The IPv6 routing operational mode on the interface. The operational mode is enabled under the conditions in the following list; otherwise, the mode is disabled: <ul style="list-style-type: none"> <li>• The IPv6 mode is enabled on the interface.</li> <li>• The routing mode is enabled on the interface.</li> <li>• The administrative mode is enabled on the interface.</li> <li>• The link is up.</li> </ul>
Link Local Prefix	The autoconfigured link-local address which is: <ul style="list-style-type: none"> <li>• Allocated from part of the IPv6 unicast address space</li> <li>• Not visible off the local link</li> <li>• Not globally unique</li> </ul>
Link Local Prefix Length	The number of bits used for the prefix of the link-local IPv6 address.
Link Local Status	The status of the IPV6 link local address. The status is TENT if routing is disabled or Duplicate Address Detection (DAD) fails. The state is Active if the interface is active and DAD is successful. Otherwise, the state is Inactive.
Routing Mode	The administrative mode for Layer 3 routing on the interface.

Table 309: IPv6 Interface Configuration Fields (Continued)

Field	Description
IPv6 Mode	The administrative mode for IPv6 on the interface. When IPv6 mode is enabled, the interface is capable of IPv6 operation without a global address. In this case, an EUI-64 based link-local address is used.
Admin Mode	The administrative mode of the interface. If an interface is administratively disabled, it will not forward traffic.
DHCPv6 Client Mode	The administrative mode of the DHCPv6 client on the interface. An interface can acquire an IPv6 address by communicating with a network DHCPv6 server (stateful configuration). An interface can also obtain an IPv6 address through stateless address autoconfiguration or static configuration.
Stateless Address AutoConfig	When this option is selected, the interface can generate its own IPv6 address by using local interface information and prefix information advertised by routers.
Interface Maximum Transmit Unit	The largest IPv6 packet size the interface can transmit, in bytes. To change the MTU value, click the Edit icon to the right of the field. To reset the MTU to the default value, click the Reset icon.
Router Duplicate Address Detection Transmits	The number of duplicate address detection probes the interface transmits while doing neighbor discovery.
Router Advertisement NS Interval	The interval between router advertisements for advertised neighbor solicitations. To change the interval, click the Edit icon to the right of the field. To reset the interval to the default value, click the Reset icon.
Router Lifetime Interval	The value that is placed in the Router Lifetime field of the router advertisements sent from the interface.
Router Advertisement Reachable Time	The value that is placed in the Reachable Time field of the router advertisements. The amount of time to consider a neighbor reachable after neighbor discovery confirmation.
Router Advertisement Interval	The transmission interval between router advertisements messages sent by the interface.
Router Advertisement Managed Config	When this option is selected, the Managed Address Configuration flag is set in router advertisements, indicating that the interface should use stateful autoconfiguration (DHCPv6) to obtain addresses.
Router Advertisement Other Config	When this option is selected, the Other Stateful Configuration flag is set in router advertisements, indicating that the interface should use stateful autoconfiguration for information other than addresses.
Router Advertisement Suppress	When this option is selected, the interface does not transmit router advertisements.
IPv6 Destination Unreachable Messages	When this option is selected, the interface can send ICMPv6 Destination Unreachable messages back to the source device if a datagram is undeliverable.
ICMPv6 Redirects	When this option is selected, the interface is allowed to send ICMPv6 Redirect messages. An ICMPv6 Redirect message notifies a host when a better route to a particular destination is available on the network segment.
IPv6 Hop Limit Unspecified	When this option is selected, the device can send Router Advertisements on this interface with an unspecified (0) current hop limit value. This will tell the hosts on the link to ignore the hop limit from this device.

Click Refresh to update the information on the screen.

#### 6.4.4 IPv6 Loopback Configuration

Use this page to configure the IPv6 routing settings for each loopback interface. A loopback interface is a logical interface that is always up (as long as it is administratively enabled) and, because it cannot go down, allows the device to have a stable IPv6 address that other network nodes and protocols can use to reach the device. The loopback can provide the source address for sent packets. The loopback interface does not behave like a network switching port. Specifically, there are no neighbors on a loopback interface; it is a pseudo device for assigning local addresses so that the other Layer 3 hosts can communicate with the device by using the loopback IPv6 address.

To display the IPv6 Loopback Configuration page, click Routing > IPv6 > Loopback Configuration in the navigation menu.

Figure 326: IPv6 Loopback Configuration

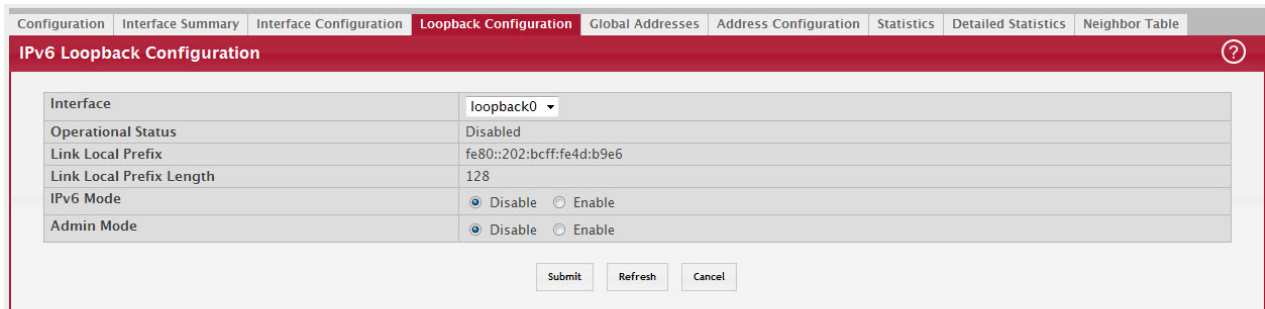


Table 310: IPv6 Loopback Configuration Fields

Field	Description
Interface	The menu contains all loopback interfaces that can be configured for IPv6 routing. To configure routing settings for an interface, select it from the menu and then configure the rest of the settings on the page. To add a new loopback interface, use the IPv6 Global Configuration page.
Operational Status	The operational status of the loopback interface. To be operational, both the IPv6 mode and administrative mode must be enabled.
Link Local Prefix	The autoconfigured link-local address which is: <ul style="list-style-type: none"> <li>Allocated from part of the IPv6 unicast address space</li> <li>Not visible off the local link</li> <li>Not globally unique</li> </ul>
Link Local Prefix Length	The number of bits used for the prefix of the link-local IPv6 address.
IPv6 Mode	The IPv6 mode on the loopback interface. When IPv6 mode is enabled, the interface is capable of IPv6 operation without a global address. In this case, an EUI-64 based link-local address is used.
Admin Mode	The administrative mode of the loopback interface.

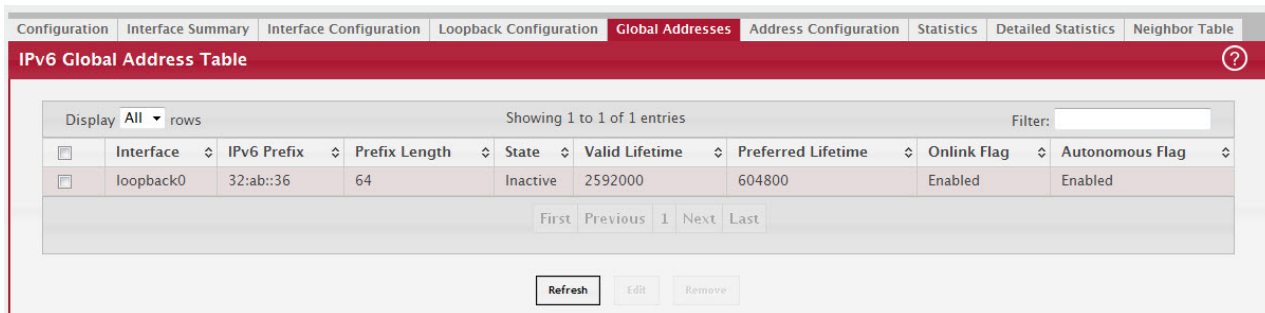
Click Refresh to update the information on the screen.

### 6.4.5 IPv6 Global Address Table

This page shows information about all global IPv6 addresses configured on all interfaces on the device. From this page, you can also remove a configured IPv6 address from an interface.

To display the IPv6 Global Address Table page, click Routing > IPv6 > Global Addresses in the navigation menu.

Figure 327: IPv6 Global Address Table



Use the buttons to perform the following tasks:

- To edit any interface, select the interface and click Edit. You are redirected to the IPv6 Global Address Configuration page for the selected interface.
- To delete the IPv6 address configuration from one or more interfaces, select each entry to remove and click Remove. You must confirm the action.

**Table 311: IPv6 Global Address Table Fields**

Field	Description
Interface	The interface associated with the rest of the data in the row.
IPv6 Prefix	The IPv6 routing prefix dynamically or manually configured on the interface. This page does not show information about link-local addresses.
Prefix Length	The number of bits used for the IPv6 prefix.
State	The link state, which is either Active or Inactive.
Valid Lifetime	The value, in seconds, to be placed in the Valid Lifetime field of the Prefix Information option in a router advertisement. The prefix is valid for on-link determination for this length of time. Hosts that generate an address from this prefix using stateless address auto-configuration can use those addresses for this length of time. An auto-configured address older than the preferred lifetime but younger than the valid lifetime are considered deprecated addresses. As defined by RFC 2462, a deprecated address is "An address assigned to an interface whose use is discouraged, but not forbidden. A deprecated address should no longer be used as a source address in new communications, but packets sent from or to deprecated addresses are delivered as expected. A deprecated address may continue to be used as a source address in communications where switching to a preferred address causes hardship to a specific upper-layer activity (e.g., an existing TCP connection)." If you specify the maximum value, an autoconfigured address may remain preferred indefinitely.
Preferred Lifetime	The value, in seconds, to be placed in the Preferred Lifetime in the Prefix Information option in a router advertisement. Addresses generated from a prefix using stateless address autoconfiguration remain preferred for this length of time. As defined by RFC 2462, a preferred address is "an address assigned to an interface whose use by upper layer protocols is unrestricted. Preferred addresses may be used as the source (or destination) address of packets sent from (or to) the interface." If you specify the maximum value, an autoconfigured address may remain preferred indefinitely.
Onlink Flag	The state of the on-link flag in the IPv6 prefix. When enabled, the prefix can be used for on-link determination by other hosts with IPv6 addresses within this prefix.
Autonomous Flag	The state of the autonomous flag in the IPv6 prefix. When enabled, the prefix can be used for autonomous address configuration by other hosts (in combination with an interface identifier on the other hosts).

Click Refresh to update the information on the screen.

### 6.4.6 IPv6 Global Address Configuration

This page shows information about all global IPv6 addresses configured on all interfaces on the device. From this page, you can also remove a configured IPv6 address from an interface.

To display the IPv6 Global Address Configuration page, click Routing > IPv6 > Address Configuration in the navigation menu.



Figure 328: IPv6 Global Address Configuration

IPv6 Global Address Configuration	
Interface	1/0/1
IPv6 Prefix	<input type="text"/> (x:x:x:x:x:x)
Prefix Length	<input type="text"/> (4 to 128)
Valid Lifetime	2592000 (0 to 4294967295)
Preferred Lifetime	604800 (0 to 4294967295)
Onlink Flag	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
Autonomous Flag	<input type="radio"/> Disable <input checked="" type="radio"/> Enable

To configure an IPv6 address on an interface that already has an IPv6 address, click Add.

Table 312: IPv6 Global Address Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
IPv6 Prefix	The IPv6 routing prefix dynamically or manually configured on the interface. This page does not show information about link-local addresses.
Prefix Length	The number of bits used for the IPv6 prefix.
Valid Lifetime	The value, in seconds, to be placed in the Valid Lifetime field of the Prefix Information option in a router advertisement. The prefix is valid for on-link determination for this length of time. Hosts that generate an address from this prefix using stateless address auto-configuration can use those addresses for this length of time. An auto-configured address older than the preferred lifetime but younger than the valid lifetime are considered deprecated addresses. As defined by RFC 2462, a deprecated address is "An address assigned to an interface whose use is discouraged, but not forbidden. A deprecated address should no longer be used as a source address in new communications, but packets sent from or to deprecated addresses are delivered as expected. A deprecated address may continue to be used as a source address in communications where switching to a preferred address causes hardship to a specific upper-layer activity (e.g., an existing TCP connection)." If you specify the maximum value, an autoconfigured address may remain preferred indefinitely.
Preferred Lifetime	The value, in seconds, to be placed in the Preferred Lifetime in the Prefix Information option in a router advertisement. Addresses generated from a prefix using stateless address autoconfiguration remain preferred for this length of time. As defined by RFC 2462, a preferred address is "an address assigned to an interface whose use by upper layer protocols is unrestricted. Preferred addresses may be used as the source (or destination) address of packets sent from (or to) the interface." If you specify the maximum value, an autoconfigured address may remain preferred indefinitely.
Onlink Flag	The state of the on-link flag in the IPv6 prefix. When enabled, the prefix can be used for on-link determination by other hosts with IPv6 addresses within this prefix.
Autonomous Flag	The state of the autonomous flag in the IPv6 prefix. When enabled, the prefix can be used for autonomous address configuration by other hosts (in combination with an interface identifier on the other hosts).

Click Refresh to update the information on the screen.



## 6.4.7 IPv6 Statistics

This page displays summary statistics about the IPv6 datagrams each interface sends and receives, successfully or unsuccessfully. It also displays summary statistics about the ICMPv6 messages each interface sends and receives. To view more information about the types of datagrams and IPv6 messages an interface has sent and received, select the interface with the information to view and click Details. You are redirected to the IPv6 Detailed Statistics page for the selected interface.

To display the IPv6 Statistics page, click Routing > IPv6 > Statistics in the navigation menu.

Figure 329: IPv6 Statistics

Interface	1/0/1
Total Datagrams Received	0
Datagrams Forwarded	0
Total ICMPv6 Messages Received	0
ICMPv6 Messages With Errors Received	0
Total ICMPv6 Messages Transmitted	0
ICMPv6 Duplicate Address Detects	0

To configure an IPv6 address on an interface that already has an IPv6 address, click Add.

Table 313: IPv6 Statistics

Field	Description
Interface	The menu contains all physical interfaces that exist on the system. Select an interface to view its IPv6 statistics.
Total Datagrams Received	The total number of input datagrams received by the interface, including those received in error.
Datagrams Forwarded	The number of output datagrams that this entity received and forwarded toward their final destinations. In entities that do not act as IPv6 routers, this counter will include only those packets which were source-routed via this entity, and the source-route processing was successful. Note that for a successfully forwarded datagram, the counter of the outgoing interface is incremented.
Total ICMPv6 Messages Received	The total number of ICMPv6 messages received by the interface, which includes all those counted by ipv6IflcmpInErrors. Note that this interface is the interface to which the ICMPv6 messages were addressed, which may not be necessarily the input interface for the messages.
ICMPv6 Messages With Errors Received	The total number of ICMPv6 messages received by the interface, which includes all those counted by ipv6IflcmpInErrors. Note that this interface is the interface to which the ICMPv6 messages were addressed, which may not be necessarily the input interface for the messages.
Total ICMPv6 Messages Transmitted	The total number of ICMPv6 messages that this interface attempted to send. Note that this counter includes all those counted by icmpOutErrors.
ICMPv6 Duplicate Address Detects	The number of duplicate IPv6 addresses detected by the interface.

Click Refresh to update the information on the screen.

## 6.4.8 IPv6 Detailed Statistics

This page displays detailed statistics about the IPv6 datagrams each interface sends and receives, successfully or unsuccessfully. It also displays detailed statistics about the ICMPv6 messages each interface sends and receives.

To display the IPv6 Detailed Statistics page, click Routing > IPv6 > Statistics in the navigation menu.

Figure 330: IPv6 Detailed Statistics

IPv6 Statistics	
Total Datagrams Received	0
Received Datagrams Locally Delivered	0
Received Datagrams Discarded Due To Header Errors	0
Received Datagrams Discarded Due To MTU	0
Received Datagrams Discarded Due To No Route	0
Received Datagrams With Unknown Protocol	0
Received Datagrams Discarded Due To Invalid Address	0
Received Datagrams Discarded Due To Truncated Data	0
Received Datagrams Discarded Other	0
Received Datagrams Reassembly Required	0
Datagrams Successfully Reassembled	0
Datagrams Failed To Reassemble	0
Datagrams Forwarded	0
Datagrams Locally Transmitted	0
Datagrams Transmit Failed	0
Datagrams Fragments Created	0
Datagrams Failed To Fragment	0
Datagrams Successfully Fragmented	0
Multicast Datagrams Received	0
Multicast Datagrams Transmitted	0

To configure an IPv6 address on an interface that already has an IPv6 address, click Add.

Table 314: IPv6 Detailed Statistics

Field	Description
Interface	The menu contains all physical interfaces that exist on the system. Select an interface to view its IPv6 statistics.
Total Datagrams Received	The total number of input datagrams received by the interface, including those received in error.
Received Datagrams Locally Delivered	The total number of datagrams successfully delivered to IPv6 user-protocols (including ICMP). This counter is incremented at the interface to which these datagrams were addressed which might not be necessarily the input interface for some of the datagrams.
Received Datagrams Discarded Due To Header Errors	The number of input datagrams discarded due to errors in their IPv6 headers, including version number mismatch, other format errors, hop count exceeded, errors discovered in processing their IPv6 options, etc.
Received Datagrams Discarded Due To MTU	The number of input datagrams that could not be forwarded because their size exceeded the link MTU of outgoing interface.
Received Datagrams Discarded Due To No Route	The number of input datagrams discarded because no route could be found to transmit them to their destination.
Received Datagrams With Unknown Protocol	The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol. This counter is incremented at the interface to which these datagrams were addressed which might not be necessarily the input interface for some of the datagrams.
Received Datagrams Discarded Due To Invalid Address	The number of input datagrams discarded because the IPv6 address in their IPv6 header's destination field was not a valid address to be received at this entity. This count includes invalid addresses, e.g., ::0, and unsupported addresses, e.g., addresses with unallocated prefixes. For entities which are not IPv6 routers and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.
Received Datagrams Discarded Due To Truncated Data	The number of input datagrams discarded because datagram frame didn't carry enough data.

Table 314: IPv6 Detailed Statistics (Continued)

Field	Description
Received Datagrams Discarded Other	The number of input IPv6 datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (e.g., for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
Received Datagrams Reassembly Required	The number of IPv6 fragments received which needed to be reassembled at this interface. Note that this counter is incremented at the interface to which these fragments were addressed which might not be necessarily the input interface for some of the fragments.
Datagrams Successfully Reassembled	The number of IPv6 datagrams successfully reassembled. Note that this counter is incremented at the interface to which these datagrams were addressed which might not be necessarily the input interface for some of the fragments.
Datagrams Failed To Reassemble	The number of failures detected by the IPv6 reassembly algorithm (for whatever reason: timed out, errors, etc.). Note that this is not necessarily a count of discarded IPv6 fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received. This counter is incremented at the interface to which these fragments were addressed which might not be necessarily the input interface for some of the fragments.
Datagrams Forwarded	The number of output datagrams that this entity received and forwarded toward their final destinations. In entities that do not act as IPv6 routers, this counter will include only those packets which were source-routed via this entity, and the source-route processing was successful. Note that for a successfully forwarded datagram, the counter of the outgoing interface is incremented.
Datagrams Locally Transmitted	The number of datagrams which this entity has successfully transmitted from this output interface.
Datagrams Transmit Failed	The number of datagrams which this entity failed to transmit successfully.
Datagrams Fragments Created	The number of IPv6 datagrams that have been successfully fragmented at this output interface.
Datagrams Failed To Fragment	The number of output datagrams that could not be fragmented at this interface.
Datagrams Successfully Fragmented	The number of output datagram fragments that have been generated as a result of fragmentation at this output interface.
Multicast Datagrams Received	The number of multicast packets received by the interface.
Multicast Datagrams Transmitted	The number of multicast packets transmitted by the interface.
Total ICMPv6 Messages Received	The total number of ICMPv6 messages received by the interface, which includes all those counted by <code>ipv6IcmpInErrors</code> . Note that this interface is the interface to which the ICMPv6 messages were addressed, which may not be necessarily the input interface for the messages.
ICMPv6 Messages With Errors Received	The number of ICMPv6 messages that the interface received but were determined to have ICMPv6-specific errors (bad ICMPv6 checksums, bad length, etc.)
ICMPv6 Destination Unreachable Messages Received	The number of ICMPv6 Destination Unreachable messages received by the interface.
ICMPv6 Messages Prohibited Administratively Received	The number of ICMPv6 destination unreachable/communication administratively prohibited messages received by the interface.
ICMPv6 Time Exceeded Messages Received	The number of ICMPv6 Time Exceeded messages received by the interface.
ICMPv6 Parameter Problem Messages Received	The number of ICMPv6 Parameter Problem messages received by the interface.

Table 314: IPv6 Detailed Statistics (Continued)

Field	Description
ICMPv6 Packet Too Big Messages Received	The number of ICMPv6 Packet Too Big messages received by the interface.
ICMPv6 Echo Request Messages Received	The number of ICMPv6 Echo (request) messages received by the interface.
ICMPv6 Echo Reply Messages Received	The number of ICMPv6 Echo Reply messages received by the interface.
ICMPv6 Router Solicit Messages Received	The number of ICMPv6 Neighbor Solicitation messages received by the interface.
ICMPv6 Router Advertisement Messages Received	The number of ICMPv6 Router Advertisement messages received by the interface.
ICMPv6 Neighbor Solicit Messages Received	The number of ICMPv6 Neighbor Solicitation messages received by the interface.
ICMPv6 Neighbor Advertisement Messages Received	The number of ICMPv6 Neighbor Advertisement messages received by the interface.
ICMPv6 Redirect Messages Received	The number of Redirect messages received.
ICMPv6 Group Membership Query Messages Received	The number of ICMPv6 Group Membership Query messages received.
ICMPv6 Group Membership Response Messages Received	The number of ICMPv6 Group Membership Response messages received.
ICMPv6 Group Membership Reduction Messages Received	The number of ICMPv6 Group Membership Reduction messages received.
Total ICMPv6 Messages Transmitted	The total number of ICMPv6 messages which this interface attempted to send. Note that this counter includes all those counted by icmpOutErrors.
ICMPv6 Messages Not Transmitted Due To Error	The number of ICMPv6 messages which this interface did not send due to problems discovered within ICMPv6 such as a lack of buffers. This value should not include errors discovered outside the ICMPv6 layer such as the inability of IPv6 to route the resultant datagram. In some implementations there may be no types of error which contribute to this counter's value.
ICMPv6 Destination Unreachable Messages Transmitted	The number of ICMPv6 Destination Unreachable Messages sent by the interface.
ICMPv6 Messages Prohibited Administratively Transmitted	The number of ICMPv6 destination unreachable/communication administratively prohibited messages sent.
ICMPv6 Time Exceeded Messages Transmitted	The number of ICMPv6 Time Exceeded messages sent by the interface.
ICMPv6 Parameter Problem Messages Transmitted	The number of ICMPv6 Parameter Problem messages sent by the interface.
ICMPv6 Packet Too Big Messages Transmitted	The number of ICMPv6 Packet Too Big messages sent by the interface.
ICMPv6 Echo Request Messages Transmitted	The number of ICMPv6 Echo (request) messages sent by the interface.

**Table 314: IPv6 Detailed Statistics (Continued)**

Field	Description
ICMPv6 Router Solicit Messages Transmitted	The number of ICMPv6 Router Solicitation messages sent by the interface.
ICMPv6 Router Advertisement Messages Transmitted	The number of ICMPv6 Router Advertisement messages sent by the interface.
ICMPv6 Neighbor Solicit Messages Transmitted	The number of ICMPv6 Neighbor Solicitation messages sent by the interface.
ICMPv6 Neighbor Advertisement Messages Transmitted	The number of ICMPv6 Neighbor Advertisement messages sent by the interface.
ICMPv6 Redirect Messages Transmitted	The number of Redirect messages sent.
ICMPv6 Group Membership Query Messages Transmitted	The number of ICMPv6 Group Membership Query messages sent.
ICMPv6 Group Membership Response Messages Transmitted	The number of ICMPv6 Group Membership Response messages sent.
ICMPv6 Group Membership Reduction Messages Transmitted	The number of ICMPv6 Group Membership Reduction messages sent.
ICMPv6 Duplicate Address Detects	The number of duplicate IPv6 addresses detected by the interface.

Click Refresh to update the information on the screen.

### 6.4.9 IPv6 Neighbor Table

This page displays the IPv6 neighbor entries in the local IPv6 neighbor cache. Neighbors are discovered by using the Neighbor Discovery Protocol via ICMPv6 messages on active IPv6 interfaces.

To display the IPv6 Neighbor Table page, click Routing > IPv6 > Neighbor Table in the navigation menu.

**Figure 331: IPv6 Neighbor Table**

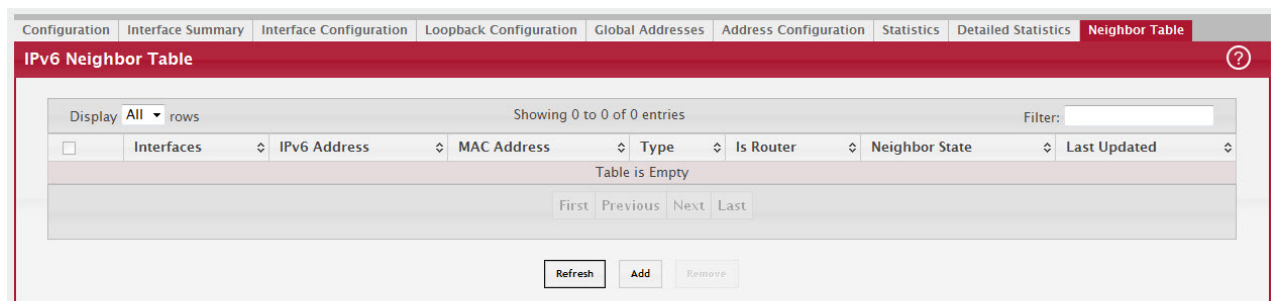


Table 315: IPv6 Neighbor Table

Field	Description
Interface	The local interface on which the neighbor was discovered.
IPv6 Address	The IPv6 prefix and prefix length of the neighbor interface.
MAC Address	The MAC address associated with the neighbor interface. If the MAC address is all zeros, the entry is a Negative NDP entry. A Negative NDP entry is added to the table when the device sends a Neighbor Solicitation Request, but it has not yet been resolved. If the request is resolved and the neighbor is reachable, its valid MAC address replaces the null address. If the request times out, the entry is removed.
Is Router	Indicates whether the neighbor is a router. If the neighbor is a router, the value is TRUE. If the neighbor is not a router, the value is FALSE.
Neighbor State	Specifies the state of the neighbor cache entry. Dynamic entries in the IPv6 neighbor discovery cache can be one of the following: <ul style="list-style-type: none"> <li>• Incmp - Address resolution is being performed on the entry. A neighbor solicitation message has been sent to the solicited-node multicast address of the target, but the corresponding neighbor advertisement message has not yet been received.</li> <li>• Reach - Positive confirmation was received within the last Reachable Time milliseconds that the forward path to the neighbor was functioning properly. While in REACH state, the device takes no special action as packets are sent.</li> <li>• Stale - More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. While in STALE state, the device takes no action until a packet is sent.</li> <li>• Delay - More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. A packet was sent within the last DELAY_FIRST_PROBE_TIME seconds. If no reachability confirmation is received within DELAY_FIRST_PROBE_TIME seconds of entering the DELAY state, send a neighbor solicitation message and change the state to PROBE.</li> <li>• Probe - A reachability confirmation is actively sought by resending neighbor solicitation messages every RetransTimer milliseconds until a reachability confirmation is received.</li> </ul>
Last Updated	The amount of time that has passed since the address was confirmed to be reachable.
Clear (Button)	Click this button to clear all entries from the table. The table is repopulated with IPv6 neighbor entries as the neighbors are discovered.

Click Refresh to update the information on the screen.

## 6.5 Configuring IPv6 Routes

The Routing > IPv6 Routes folder contains links to web pages that configure and display IP routing data.

### 6.5.1 IPv6 Route Table

This page displays the entries in the IPv6 routing table, including all dynamically learned and statically configured entries. The device uses the routing table to determine how to forward IPv6 packets. A statically-configured route does not appear in the table until it is reachable.

To display the IPv6 Global Configuration page, click Routing > IPv6 Routes > IPv6 Route Table in the navigation menu.

Figure 332: IPv6 Route Table

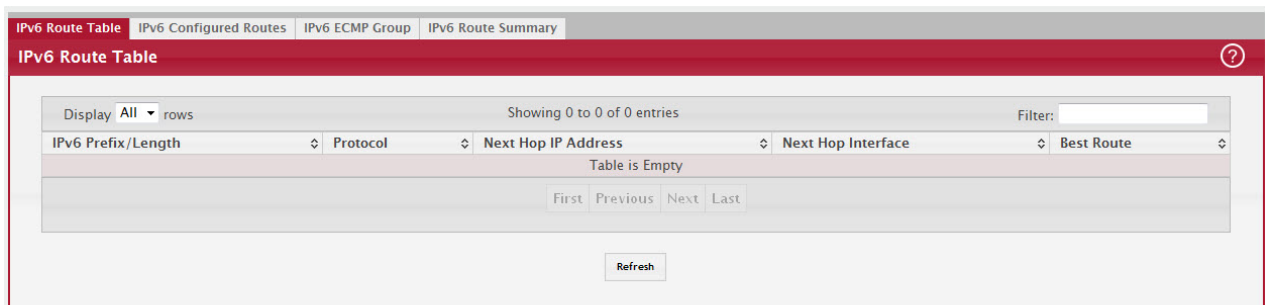


Table 316: IPv6 Route Table Fields

Field	Description
IPv6 Prefix/Length	The IPv6 address, including the prefix and prefix length, for the destination network.
Protocol	Identifies which protocol created the route. A route can be created one of the following ways: <ul style="list-style-type: none"> <li>• Dynamically learned through a supported routing protocol</li> <li>• Dynamically learned by being a directly-attached local route</li> <li>• Statically configured by an administrator</li> </ul>
Next Hop IP Address	The outgoing router IPv6 address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IPv6 address of the local interface for a directly-attached network.
Next Hop Interface	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null.
Best Route	Indicates whether the route is the preferred route to the network. If the field is blank, a better route to the same network exists in the IPv6 routing table.

Click Refresh to update the information on the screen.

### 6.5.2 IPv6 Configured Routes

Use this page to configure static IPv6 global, link local, and static reject routes in the routing table. The page shows the routes that have been manually added to the routing table. To configure a new IPv6 route, click Add.

To display the IPv6 Configured Routes page, click Routing > IPv6 Routes > IPv6 Configured Routes in the navigation menu.

Figure 333: IPv6 Configured Routes

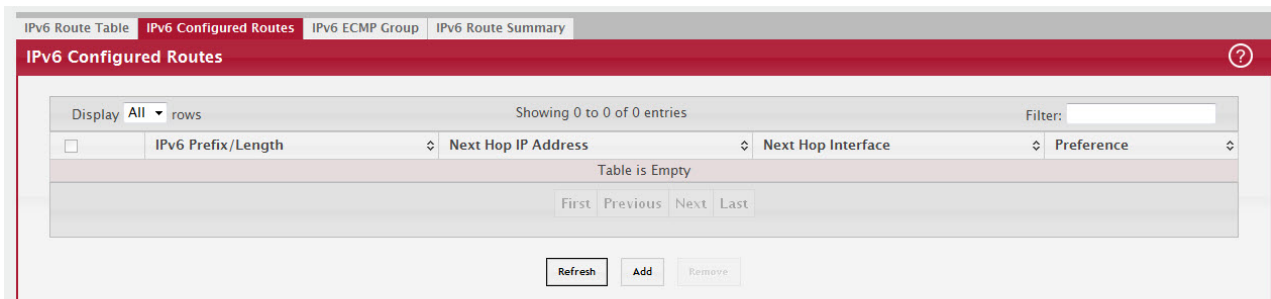


Table 317: IPv6 Configured Routes Fields

Field	Description
IPv6 Prefix/Length	The IPv6 address, including the prefix and prefix length, for the destination network.
Next Hop IP Address	The outgoing router IPv6 address to use when forwarding traffic to the next router (if any) in the path towards the destination. The next router is always one of the adjacent neighbors or the IPv6 address of the local interface for a directly-attached network.
Next Hop Interface	The outgoing interface to use when forwarding traffic to the destination. For a static reject route, the next hop is Null. The next hop is Unresolved until the device is able to reach the interface.
Preference	The preference of the route. A lower preference value indicates a more preferred route. When the routing table has more than one route to the same network, the device selects the route with the best (lowest) route preference.
After you click Add, a window opens and displays the configuration options for the new route. The following information describes the additional field in the Add Route window.	
Route Type	The type of route to configure, which is one of the following: <ul style="list-style-type: none"> <li>• Global – A route with an address that is globally routable and is recognized outside of the local network.</li> <li>• Link Local – A route with an address that is allocated from part of the IPv6 unicast address space. it is not visible off the local link and is not globally unique.</li> <li>• Static Reject – A route where packets that match the route are discarded instead of forwarded. The device might send an ICMPv6 Destination Unreachable message.</li> </ul>

If you make any changes to the page, click Submit to apply the changes to the system.

Click Refresh to update the information on the screen.

### 6.5.3 IPv6 ECMP Groups Summary

This page displays all current Equal Cost Multipath (ECMP) groups in the IPv6 routing table. An ECMP group is a set of two or more next hops used in one or more routes.

To display the IPv6 ECMP Groups Summary page, click Routing > IPv6 Routes > IPv6 ECMP Group in the navigation menu.

Figure 334: IPv6 ECMP Groups Summary

The screenshot shows a web interface for 'IPv6 ECMP Groups Summary'. At the top, there are navigation tabs: 'IPv6 Route Table', 'IPv6 Configured Routes', 'IPv6 ECMP Group' (selected), and 'IPv6 Route Summary'. Below the tabs, the page title 'IPv6 ECMP Groups Summary' is displayed. A table is shown with the following columns: 'ECMP Group', 'Number Of Next Hops', 'Route(s) Count', and 'Next Hops'. The table is currently empty, with the text 'Table is Empty' centered. Above the table, there is a 'Display' dropdown set to 'All' and 'rows', and a 'Showing 0 to 0 of 0 entries' indicator. A 'Filter:' input field is also present. Below the table, there are navigation buttons: 'First', 'Previous', 'Next', and 'Last'. At the bottom center, there is a 'Refresh' button.

Table 318: IPv6 ECMP Groups Summary Fields

Field	Description
ECMP Group	The ECMP group number associated with the rest of the data in the row. The device assigns an arbitrary ECMP group number from 1 to n to identify the ECMP group.
Number Of Next Hops	The number of next hops in the group.
Route(s) Count	The number of routes that use the set of next hops.



Table 318: IPv6 ECMP Groups Summary Fields (Continued)

Field	Description
Next Hops	The IPv6 address of each next hop in the group.
Interface	The outgoing interface of each next hop in the group.

Click Refresh to update the information on the screen.

### 6.5.4 IPv6 Route Summary

This page displays all current Equal Cost Multipath (ECMP) groups in the IPv6 routing table. An ECMP group is a set of two or more next hops used in one or more routes.

To display the IPv6 Route Summary page, click Routing > IPv6 Routes > IPv6 Route Summary in the navigation menu.

Figure 335: IPv6 Route Summary

Route Types	
Connected Routes	0
Static Routes	0
6To4 Routes	0
OSPF Routes	0
Intra Area Routes	0
Inter Area Routes	0
External Type-1 Routes	0
External Type-2 Routes	0
Total Routes	0

Route Table Counters	
Best Routes (High)	0 (0)
Alternate Routes	0
Route Adds	0
Route Deletes	0
Unresolved Route Adds	0
Invalid Route Adds	0
Failed Route Adds	0
Reserved Locals	1
Unique Next Hops (High)	0 (0)
Next Hop Groups (High)	0 (0)
ECMP Groups (High)	0 (0)
ECMP Routes	0
Truncated ECMP Routes	0
ECMP Retries	0
Number of Prefixes	

Refresh    Clear Counters

Table 319: IPv6 Route Summary Fields

Field	Description
Connected Routes	The total number of connected routes in the IPv6 routing table.
Static Routes	The total number of static routes in the IPv6 routing table.
6To4 Routes	The total number of 6to4 routes in the IPv6 routing table. A 6to4 route allows IPv6 sites to communicate with each other over an IPv4 network by treating the wide-area IPv4 network as a unicast point-to-point link layer.
OSPF Routes	The total number of routes installed by the OSPFv3 protocol.
Intra Area Routes	The total number of intra-area routes installed by the OSPFv3 protocol.
Inter Area Routes	The total number of inter-area routes installed by the OSPFv3 protocol.
External Type-1 Routes	The total number of external type-1 routes installed by the OSPFv3 protocol.

Table 319: IPv6 Route Summary Fields (Continued)

Field	Description
External Type-2 Routes	The total number of external type-2 routes installed by the OSPFv3 protocol.
Reject Routes	The total number of reject routes installed by all protocols.
Total Routes	The total number of routes in the routing table.
Best Routes (High)	The number of best routes currently in the routing table. This number counts only the best route to each destination.
Alternate Routes	The number of alternate routes currently in the routing table. An alternate route is a route that was not selected as the best route to its destination.
Route Adds	The number of routes that have been added to the routing table.
Route Deletes	The number of routes that have been deleted from the routing table.
Unresolved Route Adds	The number of route adds that failed because none of the route's next hops were on a local subnet. Note that static routes can fail to be added to the routing table at startup because the routing interfaces are not yet up. This counter gets incremented in this case. The static routes are added to the routing table when the routing interfaces come up.
Invalid Route Adds	The number of routes that failed to be added to the routing table because the route was invalid. A log message is written for each of these failures.
Failed Route Adds	The number of routes that failed to be added to the routing table because of a resource limitation in the routing table.
Reserved Locals	The number of routing table entries reserved for a local subnet on a routing interface that is down. Space for local routes is always reserved so that local routes can be installed when a routing interface bounces.
Unique Next Hops (High)	The number of distinct next hops used among all routes currently in the routing table. These include local interfaces for local routes and neighbors for indirect routes.
Next Hop Groups (High)	The current number of next hop groups in use by one or more routes. Each next hop group includes one or more next hops.
ECMP Groups (High)	The number of next hop groups with multiple next hops.
ECMP Routes	The number of routes with multiple next hops currently in the routing table.
Truncated ECMP Routes	The number of ECMP routes that are currently installed in the forwarding table with just one next hop. The forwarding table may limit the number of ECMP routes or the number of ECMP groups. When an ECMP route cannot be installed because such a limit is reached, the route is installed with a single next hop.
ECMP Retries	The number of ECMP routes that have been installed in the forwarding table after initially being installed with a single next hop.
Routes with n Next Hops	The current number of routes with each number of next hops.
Clear Counters (Button)	This button resets all IPv6 routing table event counters on this page to zero. Not that only event counters are reset; counters that report the current state of the routing table, such as the number of routes of each type, are not reset.

Click Refresh to update the information on the screen.

## 6.6 Configuring DHCPv6

The Routing > DHCPv6 folder contains links to web pages that configure and display IP routing data.

### 6.6.1 DHCPv6 Global Configuration

Use this page to configure the global Dynamic Host Configuration Protocol for IPv6 (DHCPv6) server settings on the device. The device can act as a DHCPv6 server or DHCPv6 relay agent to help assign network configuration information to IPv6 clients.

To display the DHCPv6 Global Configuration page, click Routing > DHCPv6 > Global in the navigation menu.

Figure 336: DHCPv6 Global Configuration




Table 320: DHCPv6 Global Configuration Fields

Field	Description
DHCPv6 Admin Mode	The administrative mode of the DHCPv6 server.
Conflict Logging Mode	The conflict logging mode of the bindings reported to be conflicting by the DHCPv6 Clients via the DECLINE messages
DHCPv6 Server DUID	The DHCP Unique Identifier (DUID) of the DHCPv6 server.

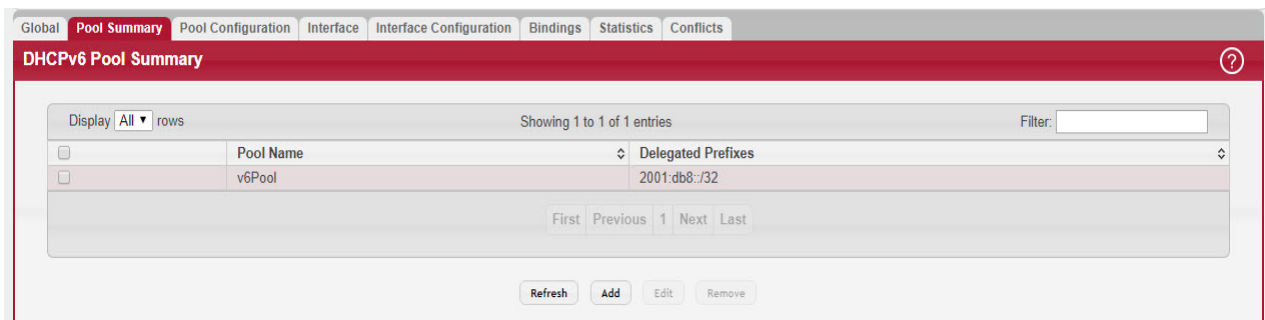
Click Refresh to update the information on the screen.

### 6.6.2 DHCPv6 Pool Summary

Use this page to view the currently configured DHCPv6 server pools and to add and remove pools. A DHCPv6 server pool is a set of network configuration information available to DHCPv6 clients that request the information.

To display the DHCPv6 Pool Summary page, click Routing > DHCPv6 > Pool Summary in the navigation menu.

Figure 337: DHCPv6 Pool Summary



Use the buttons to perform the following tasks:

- To add a pool, click Add and configure the pool information in the available fields.
- To remove a pool, select each entry to delete and click Remove. You must confirm the action before the pool is deleted.
- To change the settings for a pool, select the entry to update and click Edit. You are redirected to the DHCPv6 Pool Configuration page for the selected pool. From this page, you can configure additional bindings within the pool.

Table 321: DHCPv6 Pool Summary Fields

Field	Description
Pool Name	The name that identifies the DHCPv6 server pool.
Delegated Prefixes	The general prefix in the pool for use in allocating and assigning addresses to hosts that may be utilizing IPv6 auto-address configuration or acting as DHCPv6 clients.
After you click Add, the DHCPv6 Pool Configuration window opens. The following information describes the additional field available in the window.	
DHCPv6 Client DUID	The DHCP Unique Identifier (DUID) of the client. The DUID is a combination of the client's hardware address and client identifier.

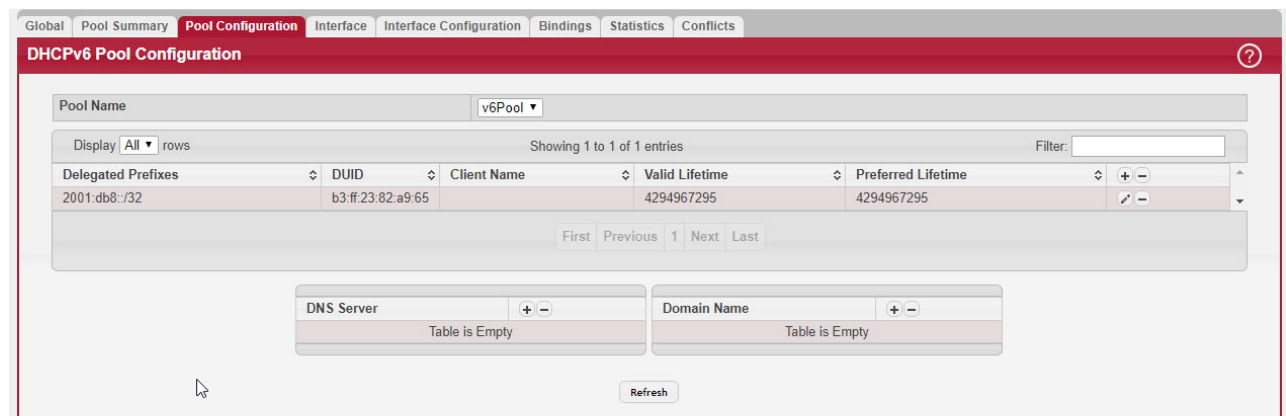
Click Refresh to update the information on the screen.

### 6.6.3 DHCPv6 Pool Configuration

Use this page to edit pool settings or to configure additional settings for existing DHCPv6 pools.

To display the DHCPv6 Pool Configuration page, click Routing > DHCPv6 > Pool Configuration in the navigation menu.

Figure 338: DHCPv6 Pool Configuration



To add, remove, or update binding entries within a pool or update other pool configuration information, you must first select the DHCPv6 pool from the Pool Name menu. After you select the pool to configure, use the icons on the page to perform the following tasks:

- To add a new binding to the selected DHCPv6 pool, click the + (plus) icon in the header row above the binding entries.
- To remove all bindings from the selected pool, click the – (minus) icon in the header row above the binding entries.
- To update the information for a binding, click the Edit icon associated with the binding.
- To remove a binding from the selected pool, click the – (minus) icon associated with the binding.
- To add DNS server or domain name information to a pool, click the + (plus) icon in the header row of the DNS Server or Domain Name field.
- To remove all configured DNS server or domain name entries from the selected pool, click the – (minus) icon in the header row of the DNS Server or Domain Name field.
- To remove a single DNS or domain name entry, click the – (minus) icon associated with the entry to remove.

Table 322: DHCPv6 Pool Configuration Fields

Field	Description
Pool Name	The menu includes all DHCPv6 server pools that have been configured on the device.
Delegated Prefixes	The IPv6 prefix and prefix length to assign the requesting client.
DUID	The DHCP Unique Identifier (DUID) of the client. The DUID is a combination of the client's hardware address and client identifier.
Client Name	The optional system name associated with the client.
Valid Lifetime	The maximum amount of time the requesting client is allowed to use the prefix.
Prefer Lifetime	The preferred amount of time the requesting client is allowed to use the prefix. The value of the Prefer Lifetime must be less than the value of the Valid Lifetime.
DNS Server	The IPv6 prefix of each DNS server each client in the pool can contact to perform address resolution.
Domain Name	The domain name configured for each client in the pool.

Click Refresh to update the information on the screen.

### 6.6.4 DHCPv6 Interface Summary

Use this page to view the per-interface settings for DHCPv6. To configure the settings, select the interface to configure and click Edit. You are redirected to the DHCPv6 Interface Configuration page for the selected interface.

To display the DHCPv6 Interface Summary page, click Routing > DHCPv6 > Interface in the navigation menu.

Figure 339: DHCPv6 Interface Summary

Interface	Interface Mode	Pool Name	Relay Interface	Destination IP Address	Remote ID
1/0/1	None	N/A	N/A	N/A	N/A
1/0/2	None	N/A	N/A	N/A	N/A
1/0/3	None	N/A	N/A	N/A	N/A
1/0/4	None	N/A	N/A	N/A	N/A
1/0/5	None	N/A	N/A	N/A	N/A
1/0/6	None	N/A	N/A	N/A	N/A
1/0/7	None	N/A	N/A	N/A	N/A
1/0/8	None	N/A	N/A	N/A	N/A
1/0/9	None	N/A	N/A	N/A	N/A
1/0/10	None	N/A	N/A	N/A	N/A

Table 323: DHCPv6 Interface Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Interface Mode	The DHCPv6 function configured on the interface, which is one of the following: <ul style="list-style-type: none"> <li>None – The interface is not configured as a DHCPv6 server or DHCPv6 relay agent.</li> <li>Server – The interface responds to requests from DHCPv6 clients.</li> <li>Relay – The interface acts as an intermediary to deliver DHCPv6 messages between clients and servers. The interface is on the same link as the client.</li> </ul>
Pool Name	(DHCPv6 server interface only) The name of the DHCPv6 pool the server uses to assign client information.

**Table 323: DHCPv6 Interface Summary Fields (Continued)**

Field	Description
Relay Interface	(DHCPv6 relay agent interface only) The interface on the device through which a DHCPv6 server is reached.
Destination IP Address	(DHCPv6 relay agent interface only) The destination IPv6 address of the DHCPv6 server to which client packets are forwarded.
Remote ID	(DHCPv6 relay agent interface only) The relay agent information option remote-ID sub-option to be added to relayed messages. This value is derived from the DHCPv6 server DUID and the relay interface number, or it can be specified as a user-defined string.

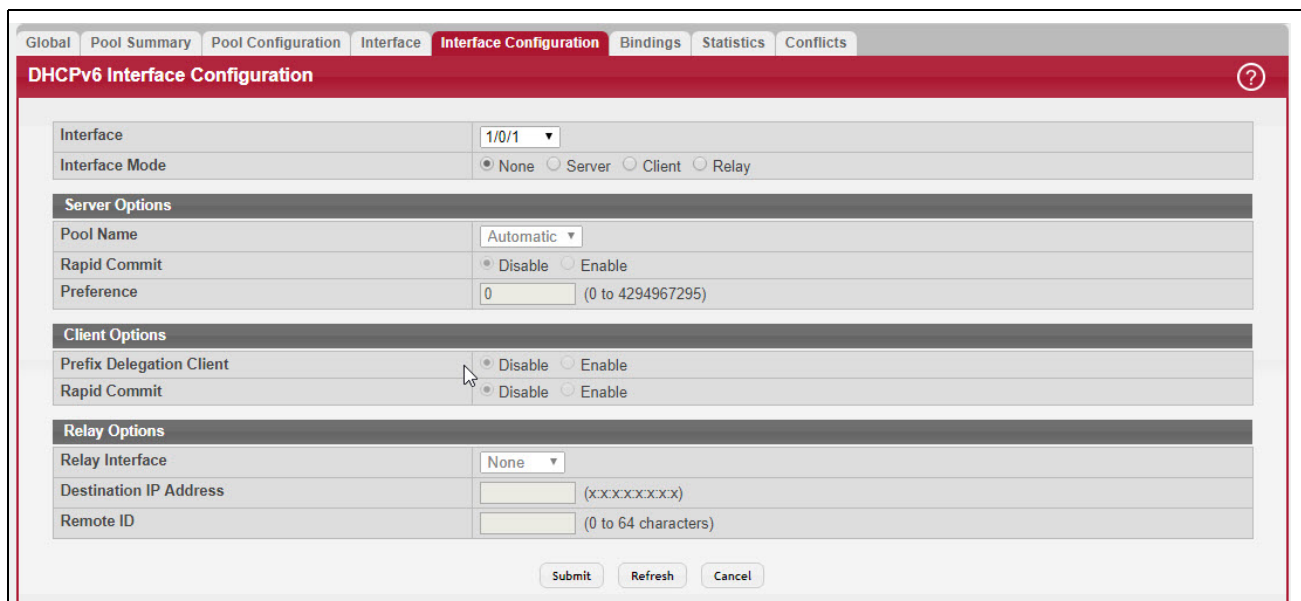
Click Refresh to update the information on the screen.

### 6.6.5 DHCPv6 Interface Configuration

Use this page to configure the per-interface settings for DHCPv6. The DHCPv6 interface modes are mutually exclusive. The fields that can be configured on this page depend on the selected mode for the interface.

To display the Interface Configuration page, click Routing > DHCPv6 > Interface Configuration in the navigation menu.

**Figure 340: DHCPv6 Interface Configuration**



**Table 324: Interface Configuration Fields**

Field	Description
Interface	Select the interface with the information to view or configure.
Interface Mode	The DHCPv6 function configured on the interface, which is one of the following: <ul style="list-style-type: none"> <li>• None – The interface is not configured as a DHCPv6 server or DHCPv6 relay agent.</li> <li>• Server – The interface responds to requests from DHCPv6 clients.</li> <li>• Client – The interface initiates requests on a link to obtain configuration parameters from one or more DHCPv6 servers.</li> <li>• Relay – The interface acts as an intermediary to deliver DHCPv6 messages between clients and servers. The interface is on the same link as the client.</li> </ul>
Server Options	(DHCPv6 server interface only) The name of the DHCPv6 pool the server uses to assign client information.
Pool Name	The name of the DHCPv6 pool the server can use to assign client information.

Table 324: Interface Configuration Fields (Continued)

Field	Description
Rapid Commit	When enabled, this option allows the DHCPv6 client to obtain configuration information by exchanging two messages with the DHCPv6 server instead of the standard four messages.
Preference	The preference value to include in DHCPv6 Advertise messages. If a DHCPv6 client receives Advertise messages from multiple DHCPv6 servers, it responds to the server with the highest preference value.
Relay Options	The information in this section can be configured only if the selected Interface Mode is Relay.
Relay Interface	The interface on the device through which a DHCPv6 server is reached.
Destination IP Address	The destination IPv6 address of the DHCPv6 server to which client packets are forwarded.
Remote ID	The relay agent information option remote-ID sub-option to be added to relayed messages. This value is derived from the DHCPv6 server DUID and the relay interface number, or it can be specified as a user-defined string.

Click Refresh to update the information on the screen.

### 6.6.6 DHCPv6 Binding Summary

Use this page to view entries in the DHCP Bindings table. After a client acquires IPv6 configuration information from the DHCPv6 server, the server adds an entry to its database. The entry is called a binding.

To display the Binding Summary page, click Routing > DHCPv6 > Bindings in the navigation menu.

Figure 341: DHCPv6 Binding Summary

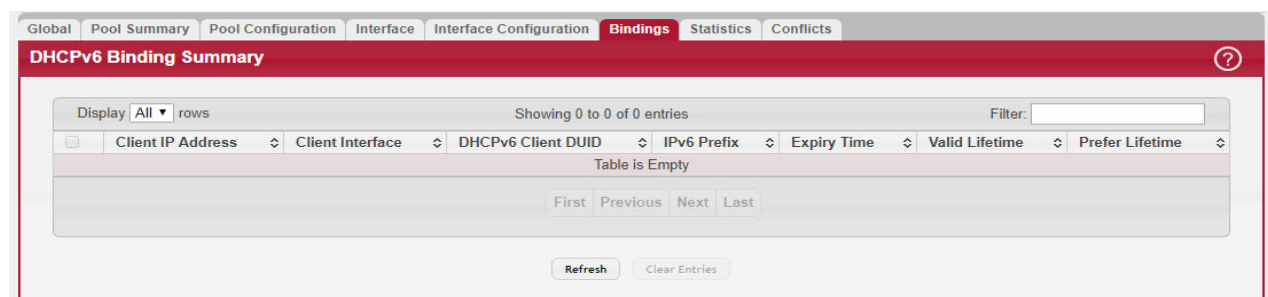


Table 325: Binding Summary Fields

Field	Description
Client IP Address	The IPv6 address associated with the client.
Client Interface	The interface number where the client binding occurred.
DHCPv6 Client DUID	The DHCP Unique Identifier (DUID) of the client. The DUID is a combination of the client's hardware address and client identifier.
IPv6 Prefix	The type of prefix associated with this binding.
Expiry Time	The number of seconds until the prefix associated with a binding expires.
Valid Lifetime	The maximum amount of time the client is allowed to use the prefix.
Prefer Lifetime	The preferred amount of time the client is allowed to use the prefix.

Click Refresh to update the information on the screen.

## 6.6.7 DHCPv6 Statistics

This page displays the DHCPv6 server statistics for the device, including information about the DHCPv6 messages sent, received, and discarded globally and on each interface. The values on this page indicate the various counts that have accumulated since they were last cleared.

To display the DHCPv6 Statistics page, click Routing > DHCPv6 > Statistics in the navigation menu.

Figure 342: DHCPv6 Statistics

Interface	Total DHCPv6 Packets Received	DHCPv6 Request Packets Received	Received DHCPv6 Packets Discarded	Total DHCPv6 Packets Transmitted	DHCPv6 Reply Packets Transmitted
All	0	0	0	0	0
1/0/1	0	0	0	0	0
1/0/2	0	0	0	0	0
1/0/3	0	0	0	0	0
1/0/4	0	0	0	0	0
1/0/5	0	0	0	0	0
1/0/6	0	0	0	0	0
1/0/7	0	0	0	0	0
1/0/8	0	0	0	0	0
1/0/9	0	0	0	0	0

Use the buttons to perform the following tasks:

- To view detailed DHCPv6 statistics for an interface, select the entry with the information to view and click Details.
- To reset the DHCPv6 counters for one or more interfaces, select each interface with the statistics to reset and click Clear.

Table 326: DHCPv6 Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. The row at the top of the table (All) contains cumulative statistics for all interfaces.
Total DHCPv6 Packets Received	The number of DHCPv6 messages received on the interface. The DHCPv6 messages sent from a DHCPv6 client to a DHCPv6 server include Solicit, Request, Confirm, Renew, Rebind, Release, Decline, and Information-Request messages. Additionally, a DHCPv6 relay agent can forward Relay-Forward messages to a DHCPv6 server.
DHCPv6 Request Packets Received	The number of DHCPv6 Request messages received on the interface. DHCPv6 Request messages are sent by a client to request IPv6 configuration information from the server.
Received DHCPv6 Packets Discarded	The number of DHCPv6 messages received on the interface that were discarded due to errors or because they were invalid.
Total DHCPv6 Packets Sent	The number of DHCPv6 messages sent by the interface. The DHCPv6 messages sent from a DHCPv6 server to a DHCPv6 client include Advertise, Reply, Reconfigure, and Relay-Reply messages.
DHCPv6 Reply Packets Transmitted	The number of DHCPv6 Reply messages sent from the interface to a DHCPv6 client in response to a Solicit, Request, Renew, Rebind, Information-Request, Confirm, Release, or Decline message.
<b>After you click Details, a window opens and shows detailed DHCPv6 statistics for the selected interface. The following information describes the additional fields that appear in the Details window.</b>	
DHCPv6 Solicit Packets Received	The number of DHCPv6 Solicit messages received on the interface. This type of message is sent by a client to locate DHCPv6 servers.



**Table 326: DHCPv6 Statistics Fields (Continued)**

Field	Description
DHCPv6 Confirm Packets Received	The number of DHCPv6 Confirm messages received on the interface. This type of message is sent by a client to all DHCPv6 servers to determine whether its configuration is valid for the connected link.
DHCPv6 Renew Packets Received	The number of DHCPv6 Renew messages received on the interface. This type of message is sent by a client to extend and update the configuration information provided by the DHCPv6 server.
DHCPv6 Rebind Packets Received	The number of DHCPv6 Rebind messages received on the interface. This type of message is sent by a client to any DHCPv6 server when it does not receive a response to a Renew message.
DHCPv6 Release Packets Received	The number of DHCPv6 Release messages received on the interface. This type of message is sent by a client to indicate that it no longer needs the assigned address.
DHCPv6 Decline Packets Received	The number of DHCPv6 Decline messages received on the interface. This type of message is sent by a client to the DHCPv6 server to indicate that an assigned address is already in use on the link.
DHCPv6 Inform Packets Received	The number of DHCPv6 Information-Request messages received on the interface. This type of message is sent by a client to request configuration information other than IP address assignment.
DHCPv6 Relay-forward Packets Received	The number of DHCPv6 Relay-Forward messages received on the interface. This type of message is sent by a relay agent to forward messages to servers.
DHCPv6 Relay-reply Packets Received	The number of DHCPv6 Relay-Reply messages received on the interface. This type of message is sent by a server to a DHCPv6 relay agent and contains the message for the relay agent to deliver to the client.
DHCPv6 Malformed Packets Received	The number of DHCPv6 messages that were received on the interface but were dropped because they were malformed.
DHCPv6 Advertisement Packets Transmitted	The number of DHCPv6 Advertisement messages sent by the interface. This type of message is sent by a server to a DHCPv6 client in response to a Solicit message and indicates that it is available for service.
DHCPv6 Reconfig Packets Transmitted	The number of DHCPv6 Reconfigure messages sent by the interface. This type of message is sent by a server to a DHCPv6 client to inform the client that the server has new or updated information. The client then typically initiates a Renew/Reply or Information-request/Reply transaction with the server to receive the updated information.
DHCPv6 Relay-forward Packets Transmitted	The number of DHCPv6 Relay-Forward messages sent by the interface. This type of message is sent by a relay agent to forward messages to servers.
DHCPv6 Relay-reply Packets Transmitted	The number of DHCPv6 Relay-Reply messages sent by the interface. This type of message is sent by a server to a DHCPv6 relay agent and contains the message for the relay agent to deliver to the client.

Click Refresh to update the information on the screen.

### 6.6.8 DHCPv6 Server Conflicts Information

This page displays information about IPv6 address conflicts detected during the DHCPv6 message exchange process between the server and client. An address conflict is created when a leased binding is declined by the DHCPv6 client.

To display the DHCPv6 Server Conflicts Information page, click Routing > DHCPv6 > Conflicts in the navigation menu.

Figure 343: DHCPv6 Server Conflicts Information

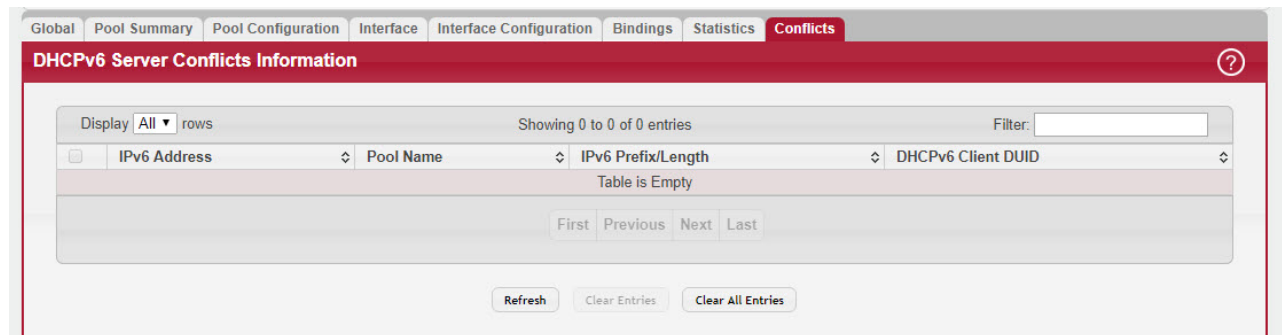


Table 327: DHCPv6 Server Conflicts Information Fields

Field	Description
IPv6 Address	The conflicting IPv6 address.
Pool Name	The name of the DHCPv6 pool the server uses to assign client information.
IPv6 Prefix/Length	The IPv6 address, including the prefix and prefix length, as a general prefix in the pool for use in allocating and assigning addresses to DHCPv6 clients.
DHCPv6 Client DUID	The DHCP Unique Identifier (DUID) of the client. The DUID is a combination of the client's hardware address and client identifier.

Use the buttons to perform the following tasks:

- To remove an entry from the table, select each entry to delete and click Clear Entries. You must confirm the action before the binding is deleted.
- To remove all entries from the table, click Clear All Entries. You must confirm the action before all bindings are deleted.
- To update the information on the screen, click Refresh.

## 6.7 Configuring Policy Based Routing

Policy based routing (PBR) enhances and modifies existing features in FASTPATH. These features are route maps and access control lists. Route maps are part of routing (see [Section 6.3: "Router"](#)) and access control lists are part of QOS (see [Section 8.1: "Configuring Access Control Lists"](#)). Because the policy-based routing feature uses services of both features mentioned previously, the FASTPATH software with a combination of the Routing and QOS packages is required to have PBR functional.

Normally, routers take forwarding decision based on routing tables to forward packets to destination addresses. Policy Based Routing is a feature that enables network administrator to define forwarding behavior based on packet contents. In brief, Policy Based Routing overrides traditional destination-based routing behavior.

The FASTPATH policy-based routing feature matches the following packet entities and overrides traditional forwarding behavior accomplished through destination-based routing:

- The size of the packet
- Protocol of the payload
- Source MAC address
- Destination MAC address
- Source IP address
- Destination IP address
- VLAN tag
- Priority

## 7/ Managing Device Security

Use the features in the Security folder on the navigation menu to set management security parameters for port, user, and server security.

### 7.1 Captive Portal

The captive portal feature allows you to prevent clients from accessing the network until user verification has been established. You can configure captive portal verification to allow access for both guest and authenticated users. Authenticated users must be validated against a database of authorized Captive Portal users before access is granted. The database can be stored locally on the device or on a RADIUS server.

#### 7.1.1 Captive Portal Global Configuration

Use this page to configure the global settings for the captive portal feature on the device.

To display the Captive Portal Global Configuration page, click Security > Captive Portal > Configuration > Global in the navigation menu.

Figure 344: Captive Portal Global Configuration

Field	Value	Range/Notes
Captive Portal	<input type="radio"/> Enable <input checked="" type="radio"/> Disable	
Operational Status	Disabled (Administratively disabled)	
IP Address	0.0.0.0	
Additional HTTP Port	0	(0 to 65535), 0 to disable
Additional HTTPS Port	0	(0 to 65535), 0 to disable
Authentication Timeout	300	(60 to 600)

Table 328: Captive Portal Global Configuration Fields

Field	Description
Captive Portal	The administrative mode of the captive portal feature.
Operational Status	The operational status of the captive portal feature, which is either Enabled or Disabled. If the captive portal is disabled, this field also displays the reason, which can be one of the following: <ul style="list-style-type: none"> <li>• None</li> <li>• Administratively disabled</li> <li>• IP address not configured</li> <li>• Routing enabled, but no routing interface</li> </ul>
IP Address	The IP address the captive portal uses.
Additional HTTP Port	The TCP port for HTTP traffic, in addition to the standard port for HTTP traffic (port 80).
Additional HTTPS Port	The TCP port for HTTP over SSL (HTTPS) traffic, in addition to the standard port for HTTPS traffic (port 443).
Authentication Timeout	The number of seconds the captive portal keeps the authentication session open with a client that is attempting to access the network through a portal. When the timeout expires, the device disconnects any active TCP or SSL connection with the client.

If you change the mode, click Submit to apply the new settings to the system.

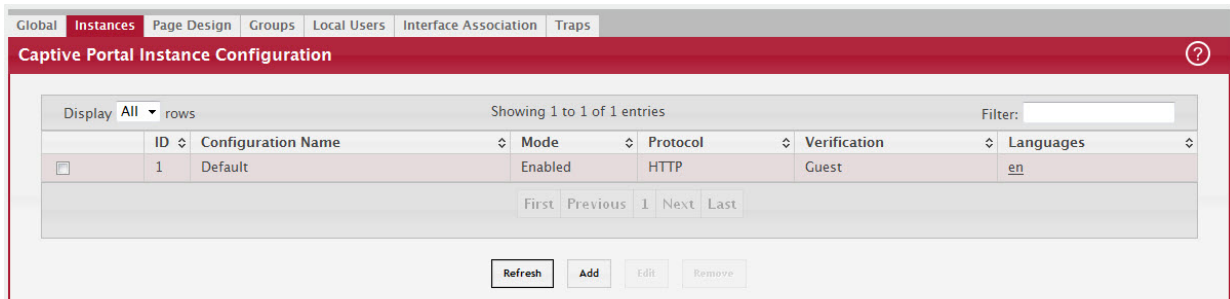
Click Refresh to update the information on the screen.

## 7.1.2 Captive Portal Instance Configuration

Use this page to configure the global settings for the captive portal feature on the device.

To display the Captive Portal Instance Configuration page, click Security > Captive Portal > Configuration > Instances in the navigation menu.

Figure 345: Captive Portal Instance Configuration



Use the buttons to perform the following tasks:

- To add a captive portal instance, click Add and configure the desired fields.
- To change the settings for an existing captive portal instance, select the instance to configure and click Edit.
- To remove one or more configured captive portal instances, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 329: Captive Portal Instance Configuration Fields

Field	Description
ID	The unique value that identifies the captive portal instance. This value is automatically assigned to the instance when it is created and cannot be changed.
Configuration	The user-configurable name that identifies the captive portal instance.
Mode	The administrative mode of the captive portal instance, which is either enabled or disabled.
Protocol	The protocol the captive portal instances uses for communication with clients, which is either HTTP or HTTPS.
Verification	The type of user verification the captive portal instance performs with clients that attempt to connect: <ul style="list-style-type: none"> <li>• Guest – The user does not need to be authenticated by a database.</li> <li>• Local – The device uses a local database to authenticated users.</li> <li>• RADIUS – The device uses a database on a remote RADIUS server to authenticate users.</li> </ul>
Languages	The IANA Language Subtag code that identifies the languages that are configured for the instance. The language code is a hyperlink to the design page for that language.
After you click Add or Edit, a window opens and allows you to configure captive portal instance settings. The following information describes the fields that are available on the Add Captive Portal Configuration and Edit Captive Portal Configuration pages.	
Configuration ID	The ID of the captive portal that is being edited. When adding a new captive portal instance, the ID is automatically assigned and appears on the main page after the instance is added.
Enable Configuration	Select this option to enable the administrative mode of the instance, or clear the option to administratively disable the captive portal instance.
Configuration Name	The user-configurable name that identifies the captive portal instance.
User Logout Mode	Select this option to allow an authenticated client to deauthenticate from the network. If this option is clear or the user does not specifically request logout, the client connection status remains authenticated until the captive portal deauthenticates the user, for example by reaching the idle timeout or session timeout values.

Table 329: Captive Portal Instance Configuration Fields (Continued)

Field	Description
Redirect Mode	Select this option to specify that the captive portal should redirect the newly authenticated client to the configured URL. If this option is clear, the user sees the locale-specific welcome page after a successful verification.
Redirect URL	The URL to which the newly authenticated client is redirected if the Redirect Mode is enabled.
Authentication Type	The settings in this section control the process used to authenticate a client that attempts to access the network through the captive portal.
Verification Mode	The type of user verification the captive portal instance performs with clients that attempt to connect: <ul style="list-style-type: none"> <li>• Guest – The user does not need to be authenticated by a database.</li> <li>• Local – The device uses a local database to authenticated users.</li> <li>• RADIUS – The device uses a database on a remote RADIUS server to authenticate users.</li> </ul>
User Group	If the Verification Mode is Local or RADIUS, assign an existing user group to the captive portal, or create a new group. All users who belong to the group are permitted to access the network through this portal. The user group list is the same for all captive portal configurations on the device. This field is unavailable if the Verification Mode is Guest.
RADIUS Auth Server	If the verification mode is RADIUS, use this field to specify the RADIUS server to use for client authentications. The device acts as the RADIUS client and performs all RADIUS transactions on behalf of the clients. Enter the RADIUS server name or IP address, or click the Select icon to the right of the field to select a RADIUS server from the list of servers configured on the device. Click the Reset icon to reset the field to the default value.
Session Parameters	The settings in this section control the communication between the authenticated client and the captive portal.
Protocol Method	The protocol the captive portal instances uses for communication with clients, which is either HTTP or HTTPS.
Session Timeout	The number of seconds to wait before terminating a session. A user is logged out once the session timeout is reached. If the value is set to 0, the timeout is not enforced.
Code	The IANA Language Subtag code for the language. All codes are listed in the IANA Language Subtag Registry. If the language is supported by the device, the code is filled in automatically when you select the language.
Language	The languages supported by the captive portal instance. Click the Select icon to display and select each language to use for the captive portal instance.

If you change the mode, click Submit to apply the new settings to the system.

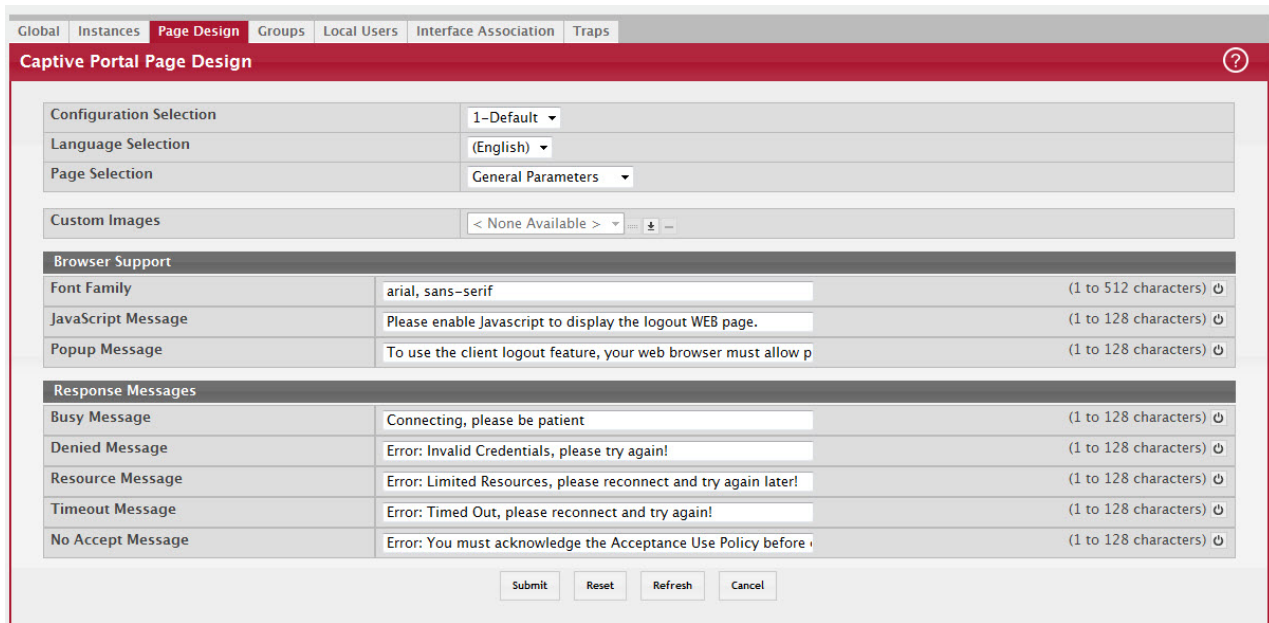
Click Refresh to update the information on the screen.

### 7.1.3 Captive Portal Page Design

Use this page to customize the appearance of the captive portal authentication, welcome, logout, and logout success pages. You can create multiple location-specific web pages for each captive portal as long as the pages for an instance all use the same verification type; either guest or authorized user web pages. This allows you to create pages in a variety of languages to accommodate a diverse group of users.

To display the Captive Portal Page Design page, click Security > Captive Portal > Configuration > Page Design in the navigation menu.

Figure 346: Captive Portal Page Design



Use the buttons to perform the following tasks:

- To add a captive portal instance, click Add and configure the desired fields.
- To change the settings for an existing captive portal instance, select the instance to configure and click Edit.
- To remove one or more configured captive portal instances, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 330: Captive Portal Page Design Fields

Field	Description
Configuration Selection	The captive portal instance to configure, identified by its ID and configuration name.
Language Selection	The menu includes each language that has been added to the captive portal instance. To configure the page design settings for a language, select it from the menu. Each language has a different set of default information.
Page Selection	The page settings to customize. The fields available on the page depend on the page selection, which is one of the following: <ul style="list-style-type: none"> <li>• General Parameters – Global settings that apply to all pages within the captive portal instance for the selected language.</li> <li>• Authentication Page – The page presented to a user who attempts to access the network through the captive portal.</li> <li>• Welcome Page – The page presented to a user after a successful authentication.</li> <li>• Logout Page – A pop-up window intended to remain open during the active session. The user clicks the Logout button on this page to terminate the session with the captive portal.</li> <li>• Logout Success Page – The page presented to a user after successfully terminating the session.</li> </ul>
<b>This section describes the fields that are available when the General Parameters option is selected from the Page Selection menu.</b>	
Custom Images	The images that are available to use for the page branding, the account image, and the backdrop for various pages. To add images, click the Download icon. Browse to the image location, select it, and click Begin Transfer. The image should be 5KB max, 200x200 pixels, GIF or JPG format. To delete an image from the list, select the file name from the menu and click the — (minus) icon.

Table 330: Captive Portal Page Design Fields (Continued)

Field	Description
Browser Support	This section includes fields that affect what the Web browser displays to captive portal clients.
Font Family	The name of the font to use for all text on the captive portal page.
JavaScript Message	The text to indicate that users must enable JavaScript to display the logout Web page. This field is applicable only when the User Logout Mode is enabled, but you can modify the text whether the feature is enabled or disabled.
Popup Message	The text to indicate that users must allow pop-up windows to display the logout Web page. This field is applicable only when the User Logout Mode is enabled, but you can modify the text whether the feature is enabled or disabled.
Response Messages	This section affects the messages the client sees when certain conditions occur.
Denied Message	The text to display when the user does not provide valid authentication information. This message displays after the user clicks the button to connect to the network.
Resource Message	The text to display when the system has rejected authentication due to system resource limitations. This message displays after the user clicks the button to connect to the network.
Timeout Message	The text to display when the system has rejected authentication because the authentication transaction took too long. This could be due to user input time, or a timeout due to the overall transaction.
Busy Message	The text to display when the user does not provide valid authentication information. This message displays after the user clicks the button to connect to the network.
No Accept Message	The text to display when the user did not accept the acceptance use policy. This message displays after the user clicks the button to connect to the network.
<b>This section describes the objects you can edit when the Authentication Page option is selected from the Page Selection menu. The screen shows an example of the authentication page presented to a captive portal user. To edit the text, color, or image associated with an object, move your mouse pointer over the object and click. A new window opens and allows you to edit the settings.</b>	
Browser Title Bar	The text that appears at the top of the web browser window or on the browser page tab for the captive portal authentication page.
Backdrop	The image file to use as the background for the captive portal authentication page. To use a non-default image, you must first upload an image to the device by using the Custom Images field, which is available when the General Parameters option is selected from the Page Selection menu.
Branding Image	The image file to use for branding, such as a company name or logo, for the captive portal authentication page. To use a non-default image, you must first upload an image to the device.
Authentication Page Title	The text that displays on the top of the authentication page.
Background Color	The background color, which outlines the acceptance use policy area and other page features.
Separator Bar Color	The color of the top and bottom bars that separate the background image from the background color.
Foreground Color	The foreground color, which is the predominant color in the area surrounding the user input field.
Account Introduction Image	The image file to use in the foreground area above the account introduction text. To use a non-default image, you must first upload an image to the device.
Account Introduction Text	The text that instructs users to authenticate, which is above the user input field.
Account Identification Label	The text that identifies the field in which the user types a name.
Account Password Label	The text to display next to the field where the user enters the password. This object is available only if the captive portal instance requires local or RADIUS verification.

Table 330: Captive Portal Page Design Fields (Continued)

Field	Description
Account Button Label	The text to display on the button the user clicks to submit the authentication information and connect to the network.
Account Instructional Text	The detailed text that instructs users to authenticate. This text appears under the button.
Acceptance Use Policy	The text to display in the Acceptance Use Policy field. The acceptance use policy instructs users about the conditions under which they are allowed to access the network.
Acknowledge AUP Text	The text that displays next to the box that the user must select to indicate that he or she accepts the terms of use.
<p>This section describes the objects you can edit when the Welcome Page option is selected from the Page Selection menu. The screen shows an example of the welcome page presented to a captive portal user after a successful authentication. To edit the text, move your mouse pointer over the appropriate object and click. A new window opens and allows you to edit the settings. Note that the branding image on the welcome page is identical to the branding image specified on the authentication page.</p>	
Welcome Title	The title text that greets the user after successfully connecting to the network.
Welcome Content Text	The optional text that displays to further identify the network to be access by the user. This message displays under the welcome title.
<p>This section describes the objects you can edit when the Logout Page option is selected from the Page Selection menu. The screen shows an example of the logout page presented to a captive portal user after a successful authentication. To edit the text, move your mouse pointer over the appropriate object and click. A new window opens and allows you to edit the settings. The logout page is intended to remain open during the active session. The user clicks the Logout button on this page to terminate the session with the captive portal.</p>	
Browser Title Bar	The text that appears at the top of the web browser window for the logout page.
Logout Title	The text that identifies the logout page.
Logout Content Text	The detailed text that confirms the user has been authenticated and instructs the user how to deauthenticate.
Logout Button Label	The text on the button the user clicks to deauthenticate.
<p>This section describes the objects you can edit when the Logout Success Page option is selected from the Page Selection menu. The screen shows an example of the logout success page presented to a captive portal user after a successful deauthentication. To edit the text, move your mouse pointer over the appropriate object and click. A new window opens and allows you to edit the settings.</p>	
Browser Title Bar	The text that appears at the top of the web browser window for the logout success page.
Backdrop	The image file to use as the background for the logout success page. To use a non-default image, you must first upload an image to the device by using the Custom Images field, which is available when the General Parameters option is selected from the Page Selection menu.
Logout Success Title	The title text that displays on the page after successfully deauthenticating from the network.
Logout Success Content Text	The optional text that provides additional information after a successful logout. This message displays under the logout success title.

If you change the mode, click Submit to apply the new settings to the system.

Click Refresh to update the information on the screen.

## 7.1.4 Captive Portal Group Configuration

Use this page to configure the Captive Portal Group settings on the device.

To display the Captive Portal Group Configuration page, click Security > Captive Portal > Configuration > Groups in the navigation menu.



Figure 347: Captive Portal Group Configuration

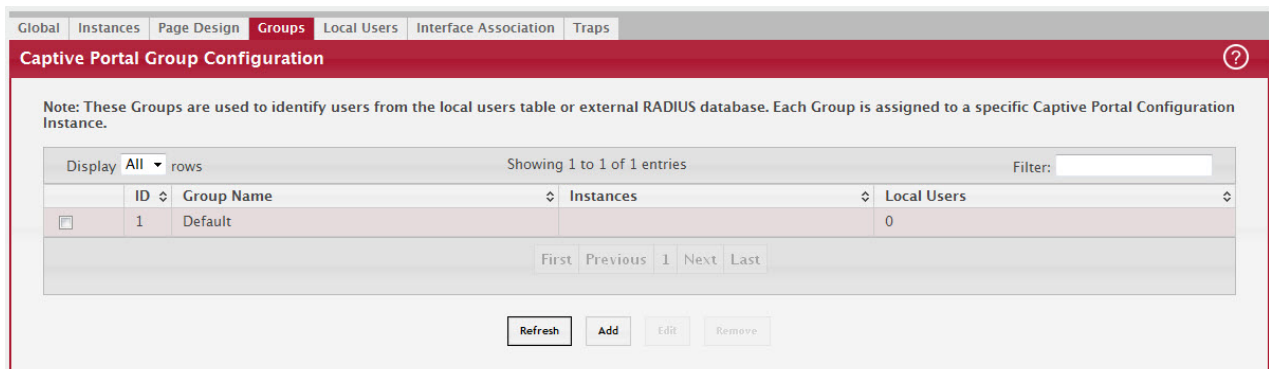


Table 331: Captive Portal Group Configuration Fields

Field	Description
ID	Shows the captive portal ID.
Group Name	Shows the captive portal group name
Instances	Shows the number of instances.
Local Users	Shows the local users.

Click Refresh to update the information on the screen.

### 7.1.5 Captive Portal Local User Configuration

Use this page to configure the Captive Portal User settings on the device.

To display the Captive Portal Local User Configuration page, click Security > Captive Portal > Configuration > Local Users in the navigation menu.

Figure 348: Captive Portal Local User Configuration

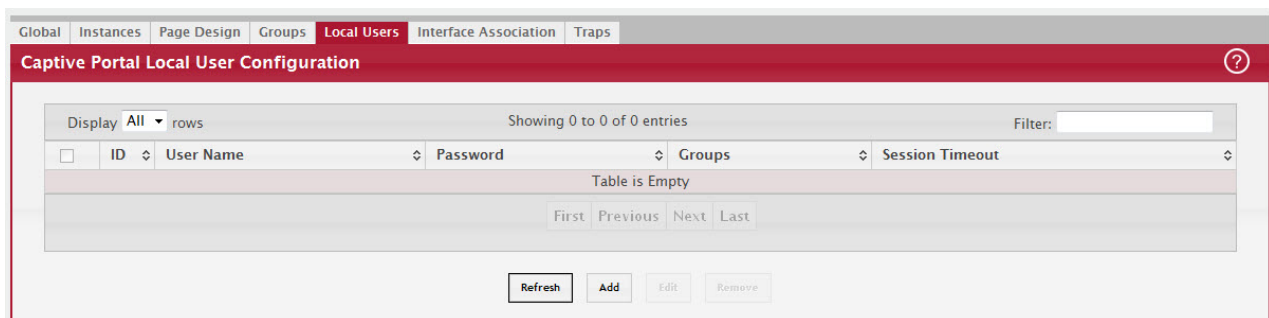


Table 332: Captive Portal Local User Configuration Fields

Field	Description
ID	Select the local user ID.
User Name	Enter the name of the user.
Password	Enter a password for the user. The password length can be from 8 to 64 characters.
Groups	Assign the user to at least one User Group. To assign a user to more than one group, press the Ctrl key and click each group. New users are assigned to the 1-Default user group by default.
Session Timeout	Shows the number of seconds a user is permitted to remain connected to the network. Once the Session Timeout value is reached, the user is logged out automatically.

Click Refresh to update the information on the screen.

## 7.1.6 Captive Portal Interface Association

Use this page to configure the Captive Portal Interface Associations on the device.

To display the Captive Portal Interface Association page, click Security > Captive Portal > Configuration > Interface Association in the navigation menu.

Figure 349: Captive Portal Interface Association

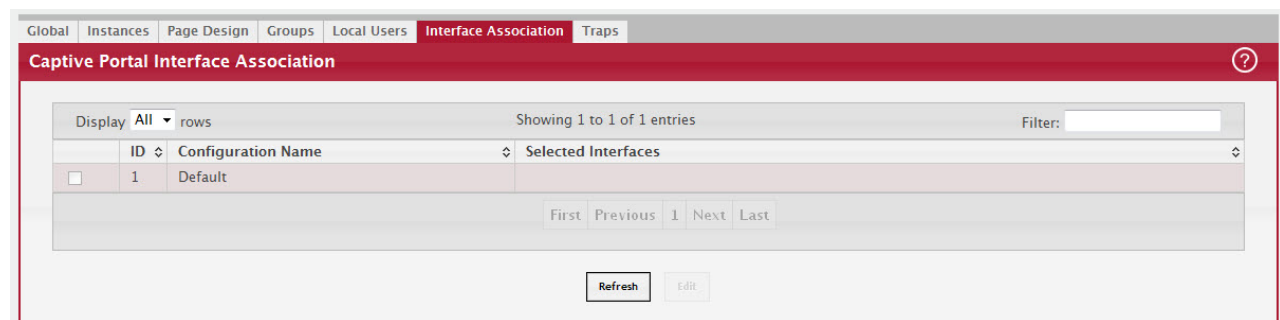


Table 333: Captive Portal Interface Association Fields

Field	Description
ID	Lists the interface association ID.
Configuration Name	Lists the captive portals configured on the switch by number and name.
Selected Interfaces	Lists the wireless interfaces that are currently associated with the selected captive portal. The interface is identified by its wireless network number and SSID

Click Refresh to update the information on the screen.

## 7.1.7 Captive Portal Traps

Use this page to configure the Captive Portal Traps on the device.

To display the Captive Portal Interface Association page, click Security > Captive Portal > Configuration > Traps in the navigation menu.

Figure 350: Captive Portal Traps

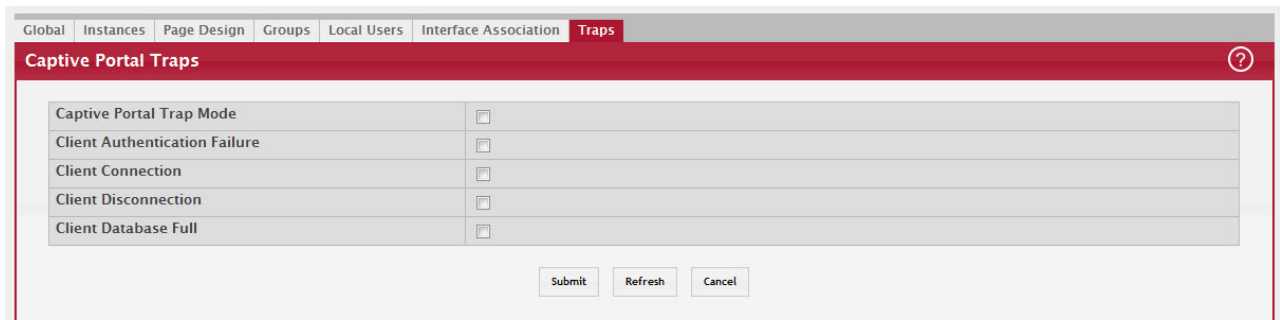


Table 334: Captive Portal Traps Fields

Field	Description
Captive Portal Trap Mode	Displays the captive portal trap mode status. To enable or disable the mode, use Captive Portal menu on the System > Trap Manager > Trap Flags page.
Client Authentication Failure	If you enable this field, the SNMP agent sends a trap when a client attempts to authenticate with a captive portal but is unsuccessful.
Client Connection	If you enable this field, the SNMP agent sends a trap when a client authenticates with and connects to a captive portal.
Client Disconnection	If you enable this field, the SNMP agent sends a trap when a client disconnects from a captive portal.
Client Database Full	If you enable this field, the SNMP agent sends a trap each time an entry cannot be added to the client database because it is full.

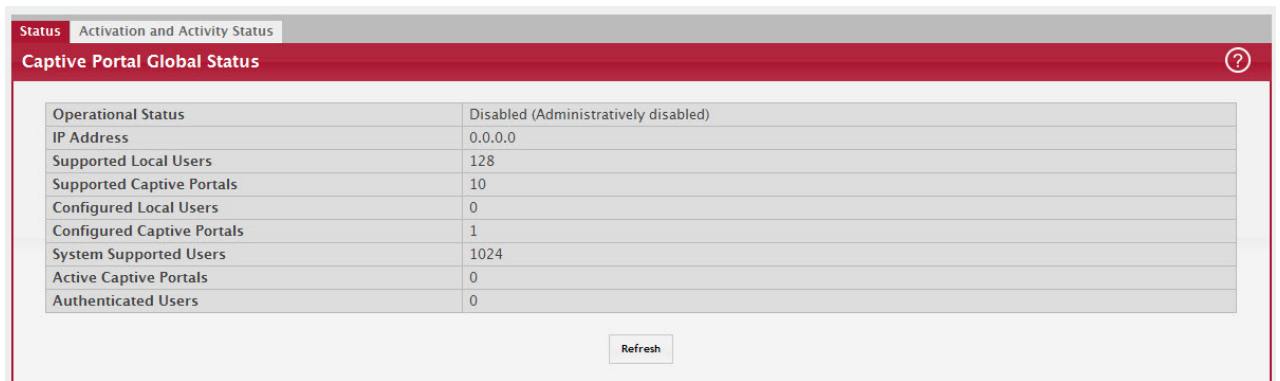
If you change the mode, click Submit to apply the new settings to the system.  
Click Refresh to update the information on the screen.

### 7.1.8 Captive Portal Global Status

This page contains a variety of information about the Captive Portal feature.

To display the Captive Portal Global Status page, click Security > Captive Portal > Global Status > Status in the navigation menu.

Figure 351: Captive Portal Global Status



**Table 335: Captive Portal Global Status Fields**

Field	Description
Operational Status	Shows whether the CP feature is enabled.
IP Address	Shows the captive portal IP address
Supported Local Users	Shows the number of authenticated users that the system can support.
Supported Captive Portals	Shows the number of captive portals configured on the switch.
Configured Local Users	Shows the number of entries that the Local User database supports.
Configured Captive Portals	Shows the number of captive portals.
System Supported Users	Shows the number of supported captive portals in the system.
Active Captive Portals	Shows the number of captive portal instances that are operationally enabled.
Authenticated Users	Shows the number of users currently authenticated to all captive portal instances on this switch.

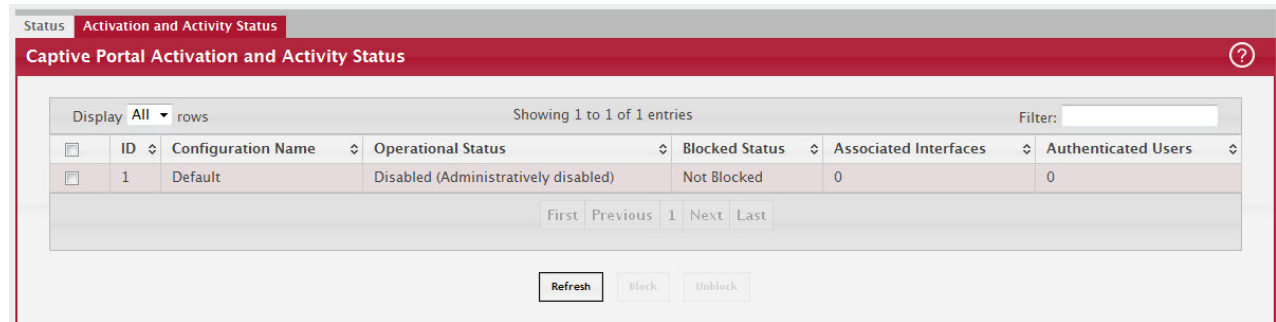
Click Refresh to update the information on the screen.

### 7.1.9 Captive Portal Activation and Activity Status

Use this page to get the information about each Captive Portal configured on the switch.

To display the Captive Portal Activation and Activity Status page, click Security > Captive Portal > Global Status > Activation and Activity Status in the navigation menu.

**Figure 352: Captive Portal Activation and Activity Status**



**Table 336: Captive Portal Activation and Activity Status Fields**

Field	Description
ID	Shows the captive portal ID.
Configuration Name	Shows the configuration name.
Operational Status	Indicates whether the captive portal is enabled or disabled.
Blocked Status	Indicates whether authentication attempts to the captive portal are currently blocked. Use the Block and Unblock buttons to control the blocked status. If the CP is blocked, users cannot gain access to the network through the CP. Use this function to temporarily protect the network during unexpected events, such as denial of service attacks. Block and Unblock are only available when the CP operational status is Enabled.

**Table 336: Captive Portal Activation and Activity Status Fields (Continued)**

Field	Description
Associated Interfaces	Shows the associated interfaces.
Authenticated Users	Shows the number of users that successfully authenticated to this captive portal and are currently using the portal.

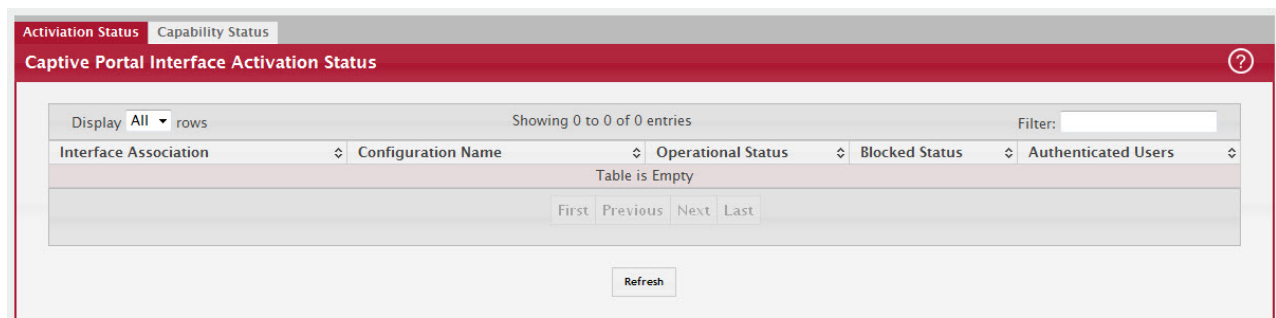
Click Refresh to update the information on the screen.

### 7.1.10 Captive Portal Interface Activation Status

Use this page to get the information for every interface assigned to a captive portal instance.

To display the Captive Portal Interface Activation Status page, click Security > Captive Portal > Global Status > Activation Status in the navigation menu.

**Figure 353: Captive Portal Interface Activation Status**



**Table 337: Captive Portal Interface Activation Status Fields**

Field	Description
Interface Association	Shows the associated interfaces.
Configuration Name	Shows the configuration name.
Operational Status	Shows whether the portal is active on the specified interface.
Blocked Status	Indicates whether the captive portal is temporarily blocked for authentications.
Authenticated Users	Displays the number of authenticated users using the captive portal instance on this interface.

Click Refresh to update the information on the screen.

### 7.1.11 Captive Portal Interface Capability Status

Use this page to get the information about interfaces that can have CPs associated with them. The page also contains status information for various capabilities. Specifically, this page indicates what services are provided through the CP to clients connected on this interface. The list of services is determined by the interface capabilities.

To display the Captive Portal Interface Capability Status page, click Security > Captive Portal > Global Status > Capability Status in the navigation menu.

Figure 354: Captive Portal Interface Capability Status

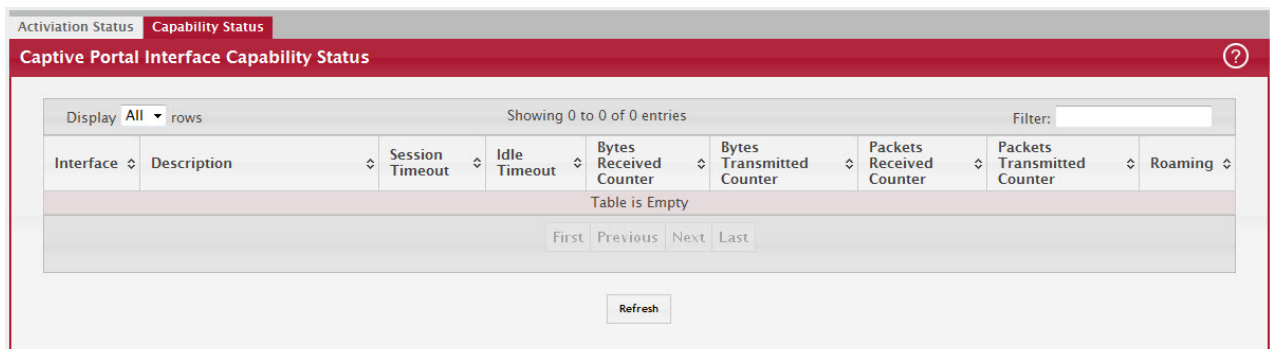


Table 338: Captive Portal Interface Capability Status Fields

Field	Description
Interface	Shows the interface.
Description	Provides a description of this interface.
Session Timeout	Shows whether the interface supports client session timeout. This attribute is supported on all interfaces.
Idle Timeout	Shows whether the interface supports a timeout when the user does not send or receive any traffic.
Bytes Received Counter	Shows whether the interface supports displaying the number of bytes received from each client.
Bytes Transmitted Counter	Shows whether the interface supports displaying the number of bytes transmitted to each client.
Packets Received Counter	Shows whether the interface supports displaying the number of packets received from each client.
Packets Transmitted Counter	Shows whether the interface supports displaying the number of packets transmitted to each client.
Roaming	Shows whether the interface supports client roaming. Only wireless interfaces support client roaming.

Click Refresh to update the information on the screen.

### 7.1.12 Captive Portal Client Summary

Use this page to view summary information about all authenticated clients that are connected through the captive portal.

To display the Captive Portal Client Summary page, click Security > Captive Portal > Global Status > Summary in the navigation menu.

Figure 355: Captive Portal Client Summary

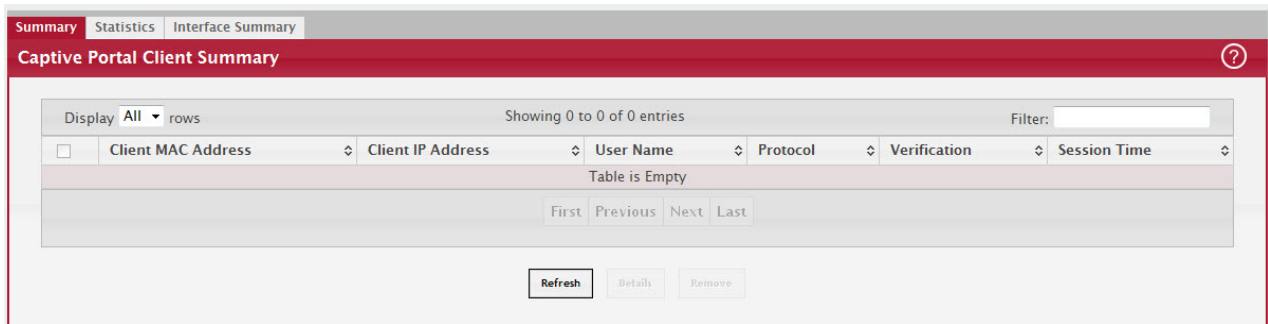


Table 339: Captive Portal Client Summary Fields

Field	Description
Client MAC Address	Identifies the MAC address of the wireless client (if applicable). If the MAC address is marked with an asterisk (*), the authenticated client is authenticated by a peer switch. In order words, the cluster controller was not the authenticator.
Client IP Address	Identifies the IP address of the wireless client (if applicable).
User Name	Displays the user name (or Guest ID) of the connected client.
Protocol	Shows the current connection protocol, which is either HTTP or HTTPS.
Verification	Shows the current account type, which is Guest, Local, or RADIUS.
Session Time	Shows the amount of time that has passed since the client was authorized.

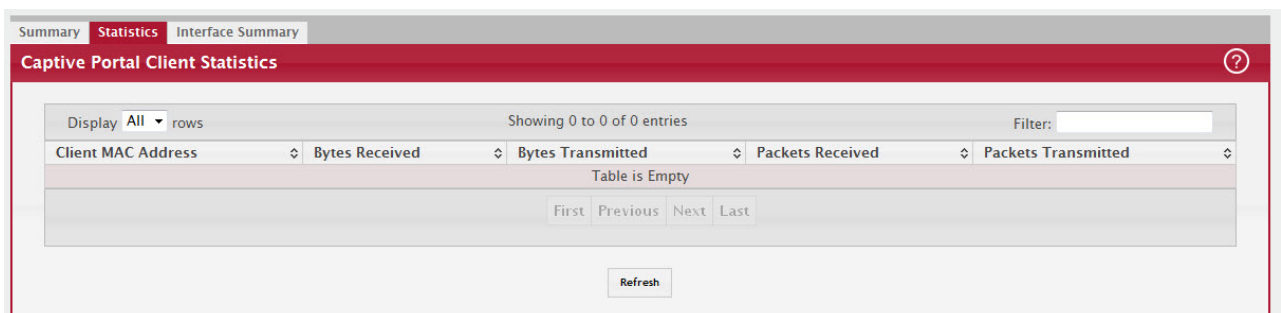
Click Refresh to update the information on the screen.

### 7.1.13 Captive Portal Client Statistics

Use this page to view information about the traffic a client has sent or received.

To display the Captive Portal Client Statistics page, click Security > Captive Portal > Global Status > Statistics in the navigation menu.

Figure 356: Captive Portal Client Statistics



**Table 340: Captive Portal Client Statistics Fields**

Field	Description
Client MAC Address	Shows the client MAC address.
Bytes Received	Total bytes the client has transmitted.
Bytes Transmitted	Total bytes the client has received.
Packets Received	Total packets the client has transmitted.
Packets Transmitted	Total packets the client has received.

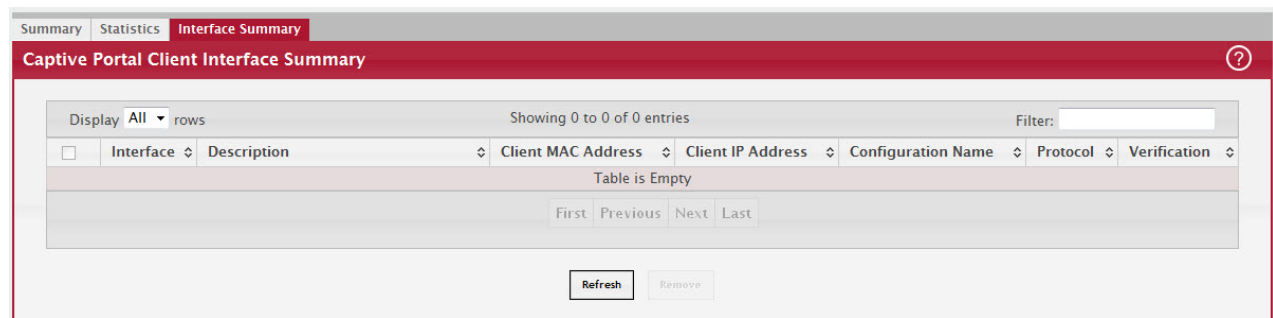
Click Refresh to update the information on the screen.

### 7.1.14 Captive Portal Client Interface Summary

Use this page to view clients that are authenticated to a specific interface.

To display the Captive Portal Client Interface Summary page, click Security > Captive Portal > Global Status > Interface Summary in the navigation menu.

**Figure 357: Captive Portal Client Interface Summary**



**Table 341: Captive Portal Client Interface Summary Fields**

Field	Description
Interface	Identifies the interface the client used to access the network.
Description	Description of the interface.
Client MAC Address	Identifies the MAC address of the wireless client. If the MAC address is marked with an asterisk (*), the authenticated client is authenticated by a peer switch. In other words, the cluster controller was not the authenticator.
Client IP Address	Identifies the IP address of the wireless client.
Configuration Name	Identifies the captive portal the client used to access the network.
Protocol	Shows the current connection protocol, which is either HTTP or HTTPS.
Verification	Shows the current account type, which is Guest, Local, or RADIUS.

Click Refresh to update the information on the screen.



## 7.2 Port Access Control

In port-based authentication mode, when 802.1x is enabled globally and on the port, successful authentication of any one supplicant attached to the port results in all users being able to use the port without restrictions. At any given time, only one supplicant is allowed to attempt authentication on a port in this mode. Ports in this mode are under bidirectional control. This is the default authentication mode.

The 802.1X network has three components:

- Authenticators: Specifies the port that is authenticated before permitting system access.
- Supplicants: Specifies host connected to the authenticated port requesting access to the system services.

Authentication Server: Specifies the external server, for example, the RADIUS server that performs the authentication on behalf of the authenticator, and indicates whether the user is authorized to access system services.

The Port Access Control folder contains links to the following pages that allow you to view and configure 802.1X features on the system.

### 7.2.1 Global Port Access Control Configuration

Use the Port Based Access Control Configuration page to enable or disable port access control on the system.

To display the Port Based Authentication page, click Security > Port Access Control > Configuration in the navigation menu.

Figure 358: Port Access Control—Port Configuration

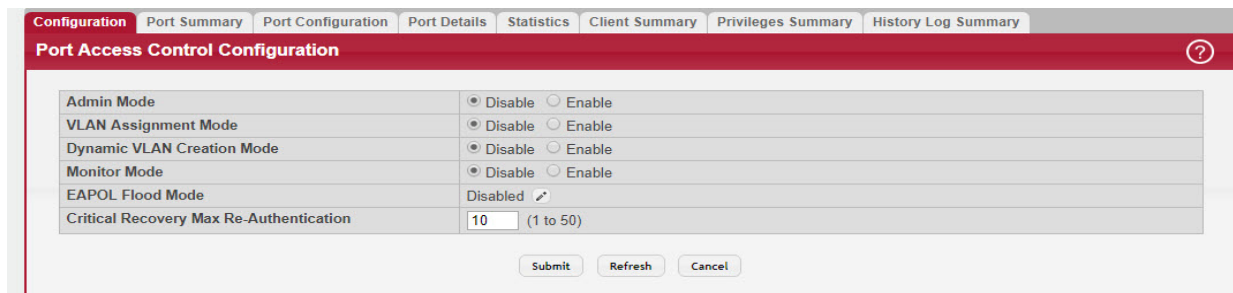


Table 342: Port Access Control—Port Configuration Fields

Field	Description
Administrative Mode	Select <code>Enable</code> or <code>Disable</code> 802.1x mode on the switch. The default is <code>Disable</code> . This feature permits port-based authentication on the switch.
VLAN Assignment Mode	If enabled, when a supplicant is authenticated by a authentication server, the port that the supplicant is connected to is placed in a particular VLAN specified by the RADIUS server. VLAN Assignment mode controls if the switch is allowed to place a port in a RADIUS-assigned VLAN. A port's VLAN assignment is determined by the first supplicant that is authenticated on the port.
Dynamic VLAN Creation Mode	Select <code>Enable</code> to allow the switch to dynamically create a RADIUS-assigned VLAN if it does not already exist in the VLAN database.
Monitor Mode	The administrative mode of the Monitor Mode feature on the device. Monitor mode is a special mode that can be enabled in conjunction with port-based access control. Monitor mode provides a way for network administrators to identify possible issues with the port-based access control configuration on the device without affecting the network access to the users of the device. It allows network access even in cases where there is a failure to authenticate, but it logs the results of the authentication process for diagnostic purposes. If the device fails to authenticate a client for any reason (for example, RADIUS access reject from the RADIUS server, RADIUS timeout, or the client itself is 802.1X unaware), the client is authenticated and is undisturbed by the failure condition(s). The reasons for failure are logged and buffered into the local logging database for tracking purposes.

Table 342: Port Access Control—Port Configuration Fields (Continued)

Field	Description
EAPOL Flood Mode	The administrative mode of the Extensible Authentication Protocol (EAP) over LAN (EAPOL) flood support on the device. EAPOL Flood Mode can be enabled when Admin Mode and Monitor Mode are disabled.
Critical Recovery Max Re-Authentication	The number of critical recovery maximum client re-authentications per second.

If you change the mode, click Submit to apply the new settings to the system.

## 7.2.2 Port Access Control Port Summary

Use this page to view summary information about the port-based authentication settings for each port.

To display the Port Access Control Port Summary page, click Security > Port Access Control > Port Summary in the navigation menu.

Figure 359: Port Access Control—Port Summary

Interface	PAE Capabilities	Control Mode	Operating Control Mode	PAE State	Backend State
1/0/1	Authenticator	Auto	N/A	Initialize	Initialize
1/0/2	Authenticator	Auto	N/A	Initialize	Initialize
1/0/3	Authenticator	Auto	N/A	Initialize	Initialize
1/0/4	Authenticator	Auto	N/A	Initialize	Initialize
1/0/5	Authenticator	Auto	N/A	Initialize	Initialize
1/0/6	Authenticator	Auto	N/A	Initialize	Initialize
1/0/7	Authenticator	Auto	N/A	Initialize	Initialize
1/0/8	Authenticator	Auto	N/A	Initialize	Initialize
1/0/9	Authenticator	Auto	N/A	Initialize	Initialize
1/0/10	Authenticator	Auto	N/A	Initialize	Initialize

Use the buttons to perform the following tasks:

- To change the port-based access control settings for a port, select the port to configure and click Edit. You are automatically redirected to the Port Access Control Port Configuration page for the selected port.
- To view additional information about the port-based access control settings for a port, select the port with the information to view and click Details. You are automatically redirected to the Port Access Control Port Details page for the selected port.

Table 343: Port Access Control—Port Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: <ul style="list-style-type: none"> <li>• Authenticator – The port enforces authentication and passes authentication information from a remote supplicant (similar to a client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access.</li> <li>• Supplicant – The port must be granted permission by the authentication server before it can access the remote authenticator port.</li> </ul>
Control Mode	The port-based access control mode configured on the port, which is one of the following: <ul style="list-style-type: none"> <li>• Auto – The port is unauthorized until a successful authentication exchange has taken place.</li> <li>• Force Unauthorized – The port ignores supplicant authentication attempts and does not provide authentication services to the client.</li> <li>• Force Authorized – The port sends and receives normal traffic without client port-based authentication.</li> <li>• MAC-Based – This mode allows multiple supplicants connected to the same port to each authenticate individually. Each host connected to the port must authenticate separately to gain access to the network. The hosts are distinguished by their MAC addresses.</li> </ul>
Operating Control Mode	The control mode under which the port is actually operating, which is one of the following: <ul style="list-style-type: none"> <li>• Auto</li> <li>• Force Unauthorized</li> <li>• Force Authorized</li> <li>• MAC-Based</li> <li>• N/A</li> </ul> <p>If the mode is N/A, port-based access control is not applicable to the port. If the port is in detached state it cannot participate in port access control. Additionally, if port-based access control is globally disabled, the status for all ports is N/A.</p>
PAE State	The current state of the authenticator PAE state machine, which is the 802.1X process that controls access to the port. The state can be one of the following: <ul style="list-style-type: none"> <li>• Initialize</li> <li>• Disconnected</li> <li>• Connecting</li> <li>• Authenticating</li> <li>• Authenticated</li> <li>• Aborting</li> <li>• Held</li> <li>• ForceAuthorized</li> <li>• ForceUnauthorized</li> </ul>
Backend State	The current state of the back-end authentication state machine, which is the 802.1X process that controls the interaction between the 802.1X client on the local system and the remote authentication server. The state can be one of the following: <ul style="list-style-type: none"> <li>• Request</li> <li>• Response</li> <li>• Success</li> <li>• Fail</li> <li>• Timeout</li> <li>• Initialize</li> <li>• Idle</li> </ul>

Table 343: Port Access Control—Port Summary Fields (Continued)

Field	Description
Initialize (Icon)	Click the Initialize icon to reset the 802.1X state machine on the associated interface to the initialization state. Traffic sent to and from the port is blocked during the authentication process. This icon can be clicked only when the port is an authenticator and the operating control mode is Auto.
Re-Authenticate (Icon)	Click the Re-Authenticate icon to force the associated interface to restart the authentication process.

If you change the mode, click Submit to apply the new settings to the system.

### 7.2.3 Port Configuration

Use the Port Access Control Port Configuration page to enable and configure port access control on one or more ports.

To access the Port Based Access Control Port Configuration page, click Security > Port Access Control > Port Configuration in the navigation menu.

Figure 360: Port Access Control Port Configuration

The screenshot displays the 'Port Access Control Port Configuration' page. At the top, there are navigation tabs: Configuration, Port Summary, Port Configuration (selected), Port Details, Statistics, Client Summary, Privileges Summary, and History Log Summary. The main content area is titled 'Port Access Control Port Configuration' and contains the following configuration fields:

- Interface:** 1/0/1
- PAE Capabilities:** Authenticator
- Authenticator Options:**
  - Control Mode:**  Force Unauthorized  Force Authorized  Auto
  - Quiet Period (Seconds):** 60 (0 to 65535)
  - Transmit Period (Seconds):** 30 (1 to 65535)
  - Guest VLAN ID:** Disabled (1 to 4093)
  - Unauthenticated VLAN ID:** Disabled (1 to 4093)
  - Supplicant Timeout (Seconds):** 30 (1 to 65535)
  - Server Timeout (Seconds):** 30 (1 to 65535)
  - Maximum Requests:** 2 (1 to 20)
  - MAB Mode:**
  - MAB Authentication Type:**  EAP-MD5  PAP  CHAP
  - Re-Authentication Period (Seconds):** Disabled (1 to 65535)
  - Maximum Users:** 48 (1 to 48)
- Supplicant Options:**
  - Control Mode:**  Force Unauthorized  Force Authorized  Auto
  - User Name:** None
  - Authentication Period (Seconds):** 30 (1 to 65535)
  - Start Period (Seconds):** 30 (1 to 65535)
  - Held Period (Seconds):** 60 (1 to 65535)
  - Maximum Start Messages:** 3 (1 to 10)

At the bottom of the form are three buttons: Submit, Refresh, and Cancel.

Use the buttons to perform the following tasks:

- To configure the port-based access control settings for one or more ports, select each port to configure and click Edit. The same settings are applied to all selected ports.
- To view additional information about the port-based access control settings for a port, select the port with the information to view and click Details.
- Click Refresh to update the information on the screen.

Table 344: Port Access Control Port Configuration Fields

Field	Description
Interface	The interface with the settings to view or configure. If you have been redirected to this page, this field is read-only and displays the interface that was selected on the Port Access Control Port Summary page.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: <ul style="list-style-type: none"> <li>• Authenticator – The port enforces authentication and passes authentication information from a remote supplicant (client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access.</li> <li>• Supplicant – The port is connected to an authenticator port and must be granted permission by the authentication server before it can send and receive traffic through the remote port. To change the PAE capabilities of a port, click the Edit icon associated with the field and select the desired setting from the menu in the Set PAE Capabilities window.</li> </ul>
Authenticator Options	The fields in this section can be changed only when the selected port is configured as an authenticator port (that is, the PAE Capabilities field is set to Authenticator).
Control Mode	The port-based access control mode on the port, which is one of the following: <ul style="list-style-type: none"> <li>• Auto – The port is unauthorized until a successful authentication exchange has taken place.</li> <li>• Force Unauthorized – The port ignores supplicant authentication attempts and does not provide authentication services to the client.</li> <li>• Force Authorized – The port sends and receives normal traffic without client port-based authentication.</li> <li>• MAC-Based – This mode allows multiple supplicants connected to the same port to each authenticate individually. Each host connected to the port must authenticate separately to gain access to the network. The hosts are distinguished by their MAC addresses.</li> </ul>
Quiet Period	The number of seconds that the port remains in the quiet state following a failed authentication exchange.
Transmit Period	The value, in seconds, of the timer used by the authenticator state machine on the port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant.
Guest VLAN ID	The value, in seconds, of the timer used for guest VLAN authentication.
Unauthenticated VLAN ID	The VLAN ID of the unauthenticated VLAN. Hosts that fail the authentication might be denied access to the network or placed on a VLAN created for unauthenticated clients. This VLAN might be configured with limited network access. To set the unauthenticated VLAN ID, click the Edit icon associated with the field and specify the ID value in the available field. To reset the unauthenticated VLAN ID to the default value, click the Reset icon associated with the field and confirm the action.
Supplicant Timeout	The amount of time that the port waits for a response before retransmitting an EAP request frame to the client.
Server Timeout	The amount of time the port waits for a response from the authentication server.
Maximum Requests	The maximum number of times that the port sends an EAP request frame (assuming that no response is received) to the client before restarting the authentication process.
MAB Mode	The MAC-based Authentication Bypass (MAB) mode on the port, which can be enabled or disabled.
MAB Authentication Type	The authentication type to be used for MAB access requests sent to the RADIUS server, which is one of the following: <ul style="list-style-type: none"> <li>• CHAP – The port uses CHAP authentication and sends a randomly generated 16-octet challenge as the CHAP-Challenge (RADIUS attribute 60) along with the CHAP-Password (RADIUS attribute 3) to the authentication server.</li> <li>• EAP-MD5 – The port uses EAP-MD5 authentication and sends the MD5 hash of the MAC address as the password in the EAP-Message (RADIUS attribute 79) to the authentication server.</li> <li>• PAP – The port uses PAP authentication and sends the MAC address of the client as the password (clear text) in the User-Password (RADIUS attribute 2) to the authentication server.</li> </ul>

Table 344: Port Access Control Port Configuration Fields (Continued)

Field	Description
Re-Authentication Period	The amount of time that clients can be connected to the port without being reauthenticated. If this field is disabled, connected clients are not forced to reauthenticate periodically. To change the value, click the Edit icon associated with the field and specify a value in the available field. To reset the reauthentication period to the default value, click the Reset icon associated with the field and confirm the action.
Maximum Users	The maximum number of clients supported on the port if the Control Mode on the port is MAC-based 802.1X authentication.
Supplicant Options	The fields in this section can be changed only when the selected port is configured as a supplicant port (that is, the PAE Capabilities field is set to Supplicant).
Control Mode	The port-based access control mode on the port, which is one of the following: <ul style="list-style-type: none"> <li>• Auto – The port is in an unauthorized state until a successful authentication exchange has taken place between the supplicant port, the authenticator port, and the authentication server.</li> <li>• Force Unauthorized – The port is placed into an unauthorized state and is automatically denied system access.</li> <li>• Force Authorized – The port is placed into an authorized state and does not require client port-based authentication to be able to send and receive traffic.</li> </ul>
User Name	The name the port uses to identify itself as a supplicant to the authenticator port. The menu includes the users that are configured for system management. When authenticating, the supplicant provides the password associated with the selected User Name.
Authentication Period	The amount of time the supplicant port waits to receive a challenge from the authentication server. If the configured Authentication Period expires, the supplicant retransmits the authentication request until it is authenticated or has sent the number of messages configured in the Maximum Start Messages field.
Start Period	The amount of time the supplicant port waits for a response from the authenticator port after sending a Start packet. If no response is received, the supplicant retransmits the Start packet.
Held Period	The amount of time the supplicant port waits before contacting the authenticator port after an active 802.1X session fails.
Maximum Start Messages	The maximum number of Start packets the supplicant port sends to the authenticator port without receiving a response before it considers the authenticator to be 802.1X-unaware.

## 7.2.4 Port Details

Use this page to view 802.1X information for a specific port.

To access the Port Access Control Port Details page, click Security > Port Access Control > Port Details in the navigation menu.

Figure 361: Port Access Control Port Details

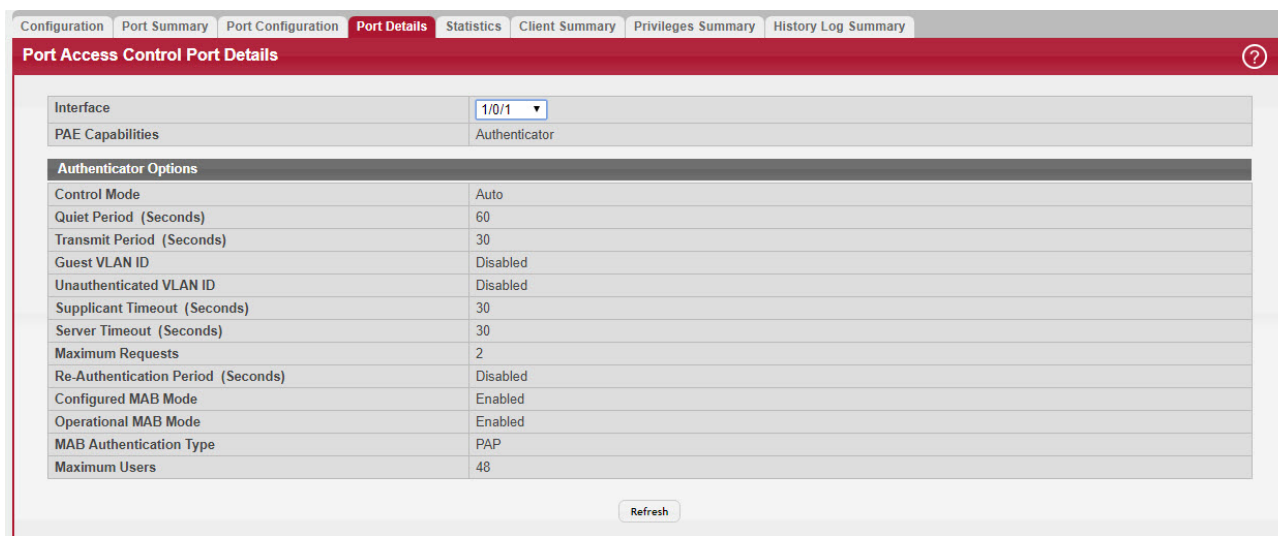


Table 345: Port Access Control Port Details Fields

Field	Description
Interface	The interface associated with the rest of the data on the page.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: <ul style="list-style-type: none"> <li>Authenticator – The port enforces authentication and passes authentication information from a remote supplicant (client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access.</li> <li>Supplicant – The port is connected to an authenticator port and must be granted permission by the authentication server before it can send and receive traffic through the remote port.</li> </ul>
Authenticator Options	The fields in this section provide information about the settings that apply to the port when it is configured as an 802.1X authenticator.
Control Mode	The port-based access control mode on the port, which is one of the following: <ul style="list-style-type: none"> <li>Auto – The port is unauthorized until a successful authentication exchange has taken place.</li> <li>Force Unauthorized – The port ignores supplicant authentication attempts and does not provide authentication services to the client.</li> <li>Force Authorized – The port sends and receives normal traffic without client port-based authentication.</li> <li>MAC-Based – This mode allows multiple supplicants connected to the same port to each authenticate individually. Each host connected to the port must authenticate separately to gain access to the network. The hosts are distinguished by their MAC addresses.</li> </ul>
Quiet Period	The number of seconds that the port remains in the quiet state following a failed authentication exchange.
Transmit Period	The value, in seconds, of the timer used by the authenticator state machine on the port to determine when to send an EAPOL EAP Request/Identity frame to the supplicant.
Guest VLAN ID	The VLAN ID for the guest VLAN. The guest VLAN allows the port to provide a distinguished service to unauthenticated users. This feature provides a mechanism to allow users access to hosts on the guest VLAN.
Unauthenticated VLAN ID	The VLAN ID of the unauthenticated VLAN. Hosts that fail the authentication might be denied access to the network or placed on a VLAN created for unauthenticated clients. This VLAN might be configured with limited network access.
Supplicant Timeout	The amount of time that the port waits for a response before retransmitting an EAP request frame to the client.

Table 345: Port Access Control Port Details Fields (Continued)

Field	Description
Server Timeout	The amount of time the port waits for a response from the authentication server.
Maximum Requests	The maximum number of times that the port sends an EAP request frame (assuming that no response is received) to the client before restarting the authentication process.
Re-Authentication Period	The amount of time that clients can be connected to the port without being reauthenticated. If this field is disabled, connected clients are not forced to reauthenticate periodically.
Configured MAB Mode	The configured MAC-based Authentication Bypass (MAB) mode on the port.
Operation MAB Mode	The operational MAB mode on the port.
MAB Authentication Type	The authentication type to be used for MAB access requests sent to the RADIUS server, which is one of the following: <ul style="list-style-type: none"> <li>• CHAP – The port uses CHAP authentication and sends a randomly generated 16-octet challenge as the CHAP-Challenge (RADIUS attribute 60) along with the CHAP-Password (RADIUS attribute 3) to the authentication server.</li> <li>• EAP-MD5 – The port uses EAP-MD5 authentication and sends the MD5 hash of the MAC address as the password in the EAP-Message (RADIUS attribute 79) to the authentication server.</li> <li>• PAP – The port uses PAP authentication and sends the MAC address of the client as the password (clear text) in the User-Password (RADIUS attribute 2) to the authentication server.</li> </ul>
Maximum Users	The maximum number of clients supported on the port if the Control Mode on the port is MAC-based 802.1X authentication.
Logical Port	The logical port number associated with the supplicant that is connected to the port.
Supplicant MAC Address	The MAC address of the supplicant that is connected to the port.
Authenticator PAE State	The current state of the authenticator PAE state machine, which is the 802.1X process that controls access to the port. The state can be one of the following: <ul style="list-style-type: none"> <li>• Initialize</li> <li>• Disconnected</li> <li>• Connecting</li> <li>• Authenticating</li> <li>• Authenticated</li> <li>• Aborting</li> <li>• Held</li> <li>• ForceAuthorized</li> <li>• ForceUnauthorized</li> </ul>
Backend Authentication State	The current state of the backend authentication state machine, which is the 802.1X process that controls the interaction between the 802.1X client on the local system and the remote authentication server. The state can be one of the following: <ul style="list-style-type: none"> <li>• Request</li> <li>• Response</li> <li>• Success</li> <li>• Fail</li> <li>• Timeout</li> <li>• Initialize</li> <li>• Idle</li> </ul>
VLAN Assigned	The ID of the VLAN the supplicant was placed in as a result of the authentication process.
VLAN Assigned Reason	The reason why the authenticator placed the supplicant in the VLAN. Possible values are: <ul style="list-style-type: none"> <li>• RADIUS</li> <li>• Default</li> <li>• Not Assigned</li> </ul>



Table 345: Port Access Control Port Details Fields (Continued)

Field	Description
Supplicant Options	The fields in this section provide information about the settings that apply to the port when it is configured as an 802.1X supplicant.
Control Mode	The port-based access control mode on the port, which is one of the following: <ul style="list-style-type: none"> <li>• Auto – The port is in an unauthorized state until a successful authentication exchange has taken place between the supplicant port, the authenticator port, and the authentication server.</li> <li>• Force Unauthorized – The port is placed into an unauthorized state and is automatically denied system access.</li> <li>• Force Authorized – The port is placed into an authorized state and does not require client port-based authentication to be able to send and receive traffic.</li> </ul>
User Name	The name the port uses to identify itself as a supplicant to the authenticator port. The menu includes the users that are configured for system management. When authenticating, the supplicant provides the password associated with the selected User Name.
Authentication Period	The amount of time the supplicant port waits to receive a challenge from the authentication server. If the configured Authentication Period expires, the supplicant retransmits the authentication request until it is authenticated or has sent the number of messages configured in the Maximum Start Messages field.
Start Period	The amount of time the supplicant port waits for a response from the authenticator port after sending a Start packet. If no response is received, the supplicant retransmits the Start packet.
Held Period	The amount of time the supplicant port waits before contacting the authenticator port after an active 802.1X session fails.
Maximum Start Messages	The maximum number of Start packets the supplicant port sends to the authenticator port without receiving a response before it considers the authenticator to be 802.1X-unaware.

#### Command Buttons

- Click Refresh to update the information on the screen.

### 7.2.5 Statistics

Use this page to view information about the Extensible Authentication Protocol over LAN (EAPOL) frames and EAP messages sent and received by the local interfaces. To view additional per-interface EAPOL and EAP message statistics, select the interface with the information to view and click Details.

To access the Port Access Control Statistics page, click Security > Port Access Control > Statistics in the navigation menu.

Figure 362: Port Access Control Statistics

Interface	PAE Capabilities	EAPOL Frames Received	EAPOL Frames Transmitted	Last EAPOL Frame Version	Last EAPOL Frame Source
1/0/1	Authenticator	0	0	0	00:00:00:00:00:00
1/0/2	Authenticator	0	0	0	00:00:00:00:00:00
1/0/3	Authenticator	0	0	0	00:00:00:00:00:00
1/0/4	Authenticator	0	0	0	00:00:00:00:00:00
1/0/5	Authenticator	0	0	0	00:00:00:00:00:00
1/0/6	Authenticator	0	0	0	00:00:00:00:00:00
1/0/7	Authenticator	0	0	0	00:00:00:00:00:00
1/0/8	Authenticator	0	0	0	00:00:00:00:00:00
1/0/9	Authenticator	0	0	0	00:00:00:00:00:00
1/0/10	Authenticator	0	0	0	00:00:00:00:00:00

Table 346: Port Access Control Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When viewing detailed information for an interface, this field identifies the interface being viewed.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: <ul style="list-style-type: none"> <li>Authenticator – The port enforces authentication and passes authentication information from a remote supplicant (similar to a client or host) to the authentication server. If the server successfully authenticates the supplicant, the port allows access.</li> <li>Supplicant – The port must be granted permission by the authentication server before it can access the remote authenticator port.</li> </ul>
EAPOL Frames Received	The total number of valid EAPOL frames received on the interface.
Last EAPOL Frame Version	The total number of EAPOL frames sent by the interface.
Last EAPOL Frame Source	The source MAC address attached to the most recently received EAPOL frame.
<p>After you click Details, a window opens and displays additional information about the EAPOL and EAP messages the interface sends and receives. The following information describes the additional fields that appear in the Details window. The fields this window displays depend on whether the interface is configured as an authenticator or supplicant, as noted in the applicable field descriptions.</p>	
EAPOL Start Frames Received	The total number of EAPOL-Start frames received on the interface. EAPOL-Start frames are sent by a supplicant to initiate the 802.1X authentication process when it connects to the interface. This field is displayed only if the interface is configured as an authenticator.
EAPOL Logoff Frames Received	The total number of EAPOL-Logoff frames received on the interface. EAPOL-Logoff frames are sent by a supplicant to indicate that it is disconnecting from the network, and the interface can return to the unauthorized state. This field is displayed only if the interface is configured as an authenticator.

Table 346: Port Access Control Statistics Fields (Continued)

Field	Description
EAP Response/ID Frames Received	The total number of EAP-Response Identity frames the interface has received. EAP-Response Identity frames are sent by a supplicant to provide user information that is used to for authentication. This field is displayed only if the interface is configured as an authenticator.
EAP Response Frames Received	The total number of EAP-Response frames the interface has received. EAP-Response frames are sent from a supplicant to an authentication server during the authentication process. This field is displayed only if the interface is configured as an authenticator.
EAP Request/ID Frames Transmitted	The total number of EAP-Request Identity frames the interface has sent. EAP-Request Identity frames are sent from an authenticator to a supplicant to request user information that is used to for authentication. This field is displayed only if the interface is configured as an authenticator.
EAPOL Start Frames Transmitted	The total number of EAPOL-Start frames the interface has sent to a remote authenticator. EAPOL-Start frames are sent by a supplicant to initiate the 802.1X authentication process when it connects to the interface. This field is displayed only if the interface is configured as a supplicant.
EAPOL Logoff Frames Transmitted	The total number of EAPOL-Logoff frames the interface has sent to a remote authenticator. EAPOL-Logoff frames are sent by a supplicant to indicate that it is disconnecting from the network, and the interface can return to the unauthorized state. This field is displayed only if the interface is configured as a supplicant.
EAP Response/ID Frames Transmitted	The total number of EAP-Response Identity frames the interface has sent. EAP-Response Identity frames are sent by a supplicant to provide user information that is used to for authentication. This field is displayed only if the interface is configured as a supplicant.
EAP Request/ID Frames Received	The total number of EAP-Request Identity frames the interface has received. EAP-Request Identity frames are sent from an authenticator to a supplicant to request user information that is used to for authentication. This field is displayed only if the interface is configured as a supplicant.
EAP Request Frames Received	The total number of EAP-Request frames the interface has received. EAP-Request frames are sent from the authentication server to the supplicant during the authentication process. This field is displayed only if the interface is configured as a supplicant.
Invalid EAPOL Frames Received	The number of unrecognized EAPOL frames received on the interface.
EAPOL Length Error Frames Received	The number of EAPOL frames with an invalid packet body length received on the interface.
Clear (Button)	Resets all statistics counters to 0 for the selected interface or interfaces.

### Command Buttons

- Click Refresh to update the information on the screen.

## 7.2.6 Client Summary

This page displays information about supplicant devices that are connected to the local authenticator ports. If there are no active 802.1X sessions, the table is empty. To view additional information about a supplicant, select the interface it is connected to and click Details.

To access the Port Access Control Client Summary page, click Security > Port Access Control > Client Summary in the navigation menu.

Figure 363: Port Access Control Client Summary

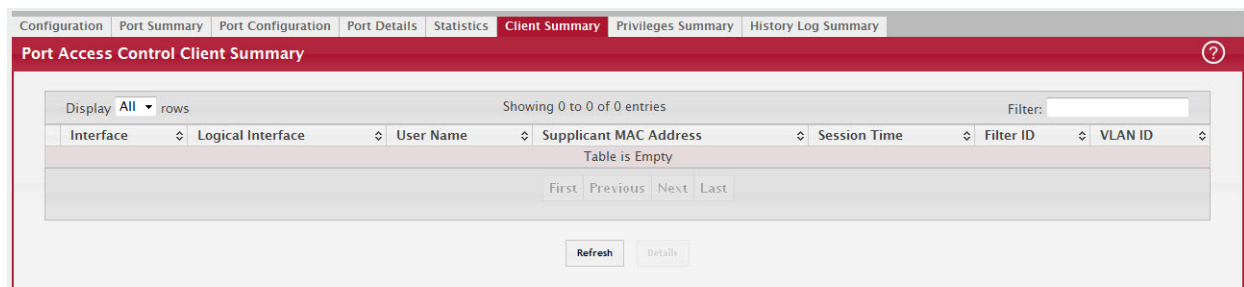


Table 347: Port Access Control Client Summary Fields

Field	Description
Interface	The local interface associated with the rest of the data in the row. When viewing detailed information for an interface, this field identifies the interface being viewed.
Logical Interface	The logical port number associated with the supplicant that is connected to the port.
User Name	The name the client uses to identify itself as a supplicant to the authentication server.
Supp MAC Address	The MAC address of the supplicant that is connected to the port.
Session Time	The amount of time that has passed since the connected supplicant was granted access to the network through the authenticator port.
Filter ID	The policy filter ID assigned by the authenticator to the supplicant device.
VLAN ID	The ID of the VLAN the supplicant was placed in as a result of the authentication process.
After you click Details, a window opens and displays additional information about the client. The following information describes the additional fields that appear in the window.	
Session Timeout	The reauthentication timeout period set by the RADIUS server to the supplicant device.
Session Termination Action	The termination action set by the RADIUS server that indicates the action that will take place once the supplicant reaches the session timeout value.

### Command Buttons

- Click Refresh to update the information on the screen.

## 7.2.7 Privileges Summary

Use this page to grant or deny port access to users configured on the system. To change the access control privileges for one or more ports, select each interface to configure and click Edit. The same settings are applied to all selected interfaces.

To access the Port Access Control Privileges Summary page, click Security > Port Access Control > Privileges Summary in the navigation menu.

Figure 364: Port Access Control Privileges Summary

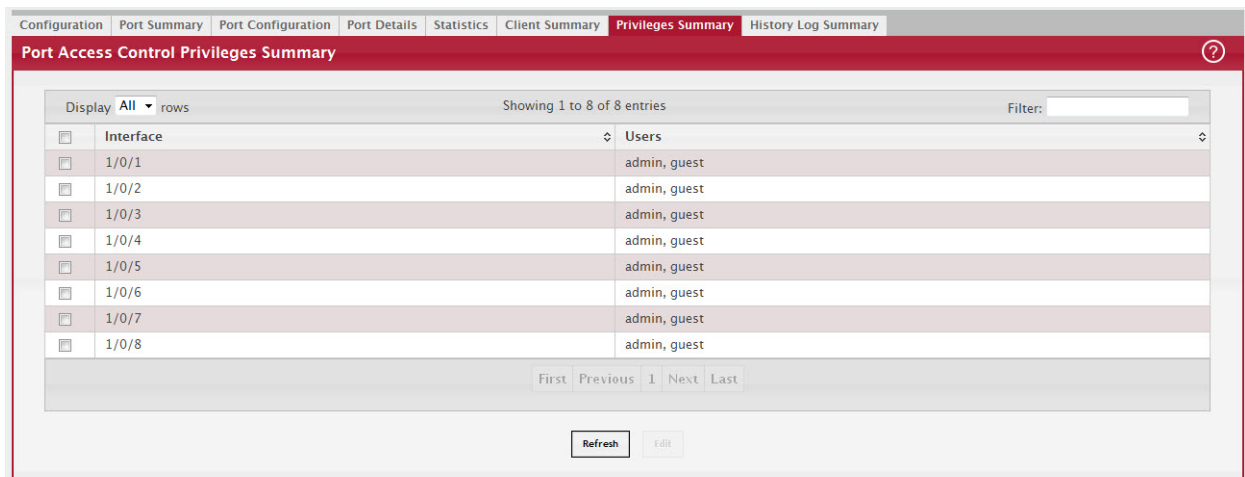


Table 348: Port Access Control Privileges Summary Fields

Field	Description
Interface	The local interface associated with the rest of the data in the row. When configuring access information for one or more interfaces, this field identifies each interface being configured.
Users	The users that are allowed access to the system through the associated port. When configuring user access for a port, the Available Users field lists the users configured on the system that are denied access to the port. The users in the Selected Users field are allowed access. To move a user from one field to the other, click the user to move (or CTL + click to select multiple users) and click the appropriate arrow.

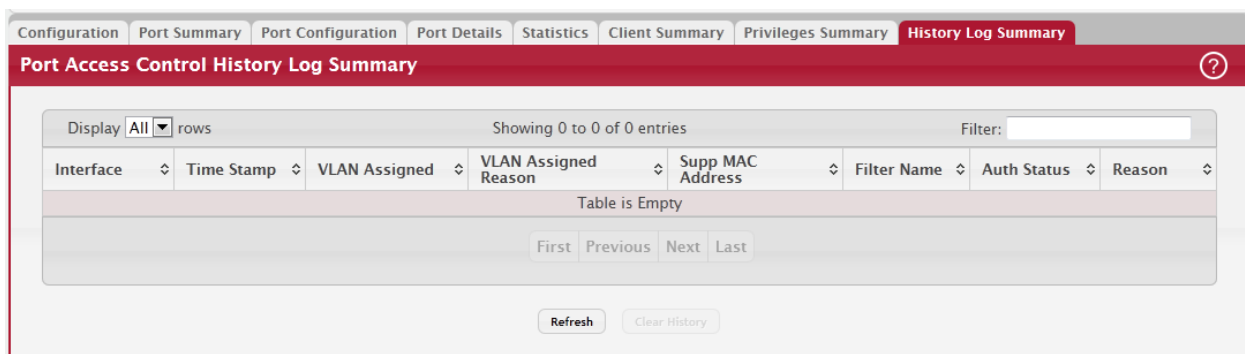
**Command Buttons**

- Click Refresh to update the information on the screen.

**7.2.8 History Log Summary**

This page displays information about the 802.1X entries that exist in the history log table for the active 802.1X sessions. To access the Port Access Control History Log Summary page, click Security > Port Access Control > History Log Summary in the navigation menu.

Figure 365: Port Access Control History Log Summary



**Table 349: Port Access Control History Log Summary Fields**

Field	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that have entries in the log history are listed.
Time Stamp	The absolute time when the authentication event took place.
VLAN Assigned	The ID of the VLAN the supplicant was placed in as a result of the authentication process.
VLAN Assigned Reason	The reason why the authenticator placed the supplicant in the VLAN. Possible values are: <ul style="list-style-type: none"> <li>• RADIUS</li> <li>• Unauth</li> <li>• Default</li> <li>• Not Assigned</li> </ul>
Supp MAC Address	The MAC address of the supplicant that is connected to the port.
Filter Name	The policy filter ID assigned by the authenticator to the supplicant device.
Auth Status	The authentication status of the client or port.
Reason	The reason for the successful or unsuccessful authentication.

Click Refresh to update the information on the screen.

## 7.3 RADIUS Settings

Remote Authorization Dial-In User Service (RADIUS) servers provide additional security for networks. The RADIUS server maintains a user database, which contains per-user authentication information. RADIUS servers provide a centralized authentication method for:

- Telnet Access
- Web Access
- Console to Switch Access
- Port Access Control (802.1X)

The RADIUS folder contains links to pages that help you view and configure system RADIUS settings.

### 7.3.1 RADIUS Configuration

Use the RADIUS Configuration page to view and configure various settings for the RADIUS servers configured on the system.

To access the RADIUS Configuration page, click Security > RADIUS > Configuration in the navigation menu.

Figure 366: RADIUS Configuration

Table 350: RADIUS Configuration Fields

Field	Description
Max Number of Retransmits	The maximum number of times the RADIUS client on the device will retransmit a request packet to a configured RADIUS server after a response is not received. If multiple RADIUS servers are configured, the max retransmit value will be exhausted on the first server before the next server is attempted. A retransmit will not occur until the configured timeout value on that server has passed without a response from the RADIUS server. Therefore, the maximum delay in receiving a response from the RADIUS server equals the sum of (retransmit * timeout) for all configured servers. If the RADIUS request was generated by a user login attempt, all user interfaces will be blocked until the RADIUS application returns a response.
Timeout Duration	The number of seconds the RADIUS client waits for a response from the RADIUS server. Consideration to maximum delay time should be given when configuring RADIUS timeout and RADIUS max retransmit values.
Accounting Mode	Specifies whether the RADIUS accounting mode on the device is enabled or disabled.

Table 350: RADIUS Configuration Fields (Continued)

Field	Description
MAB Attribute	<p>The RADIUS attribute 1 (User-Name) for sending MAC-based Authentication Bypass (MAB) requests from the client to the RADIUS server.</p> <p>The authenticator sends a request to the authentication server with the MAC address of the client (by default 'hh:hh:hh:hh:hh:hh') as the User-Name. This attribute is sent irrespective of the authentication type configured on the MAB interface.</p> <p>To configure the MAB attribute format, click the Edit icon and enter the desired settings in the available fields. To reset the MAB attribute to the default values, click the Reset icon and confirm the action</p> <p>After you click Edit, the Set MAB Attribute window appears and includes the following fields:</p> <ul style="list-style-type: none"> <li>• Group Size–The group size used by the switch to format the RADIUS attribute 1 (User-Name) of the MAB request. The size is the number of characters included in a group. <ul style="list-style-type: none"> <li>- In the following example, the group size is 1: 0:0:1:0:1:8:9:9:F:2:B:3</li> <li>- In the following example, the group size is 2: 00:10:18:99:F2:B3</li> <li>- In the following example, the group size is 4: 0010:1899:F2B3</li> <li>- In the following example, the group size is 12: 00101899F2B3</li> </ul> </li> <li>• Separator–The separator used by the switch to format the RADIUS attribute 1 (User-Name) of the MAB request. <ul style="list-style-type: none"> <li>- In the following example, the separator is - (hyphen): 00-10-18-99-F2-B3</li> <li>- In the following example, the separator is : (colon): 00:10:18:99:F2:B3</li> </ul> </li> <li>• Case–The case of any letters used by the switch to format the RADIUS attribute 1 (User-Name) of the MAB request. <ul style="list-style-type: none"> <li>- In the following example, the case is lowercase: 00:d0:18:99:f2:b3</li> <li>- In the following example, the case is uppercase: 00:D0:18:99:F2:B3</li> </ul> </li> </ul>
VSA Authentication	Specifies whether the Cisco Vendor Specific Attributes (VSA) sent by the RADIUS server are processed by the switch.
RADIUS Attributes	
NAS-IP-ADDRESS (Attribute 4)	<p>The network access server (NAS) IP address for the RADIUS server.</p> <p>To specify an address, click the Edit icon and enter the IP address of the NAS in the available field. The address should be unique to the NAS within the scope of the RADIUS server. The NAS IP address is used only in Access-Request packets. To reset the NAS IP address to the default value, click the Reset icon and confirm the action.</p>
CALLED-STATION-ID (Attribute 30)	<p>Specifies the format in which the MAC address is sent to the RADIUS server in attribute 30. To specify a format, click the Edit icon and select one of the following:</p> <ul style="list-style-type: none"> <li>• Legacy Lowercase – Format the MAC address as xx:xx:xx:xx:xx:xx</li> <li>• Legacy Uppercase – Format the MAC address as XX:XX:XX:XX:XX:XX</li> <li>• IETF Lowercase – Format the MAC address as xx-xx-xx-xx-xx-xx</li> <li>• IETF Uppercase – Format the MAC address as XX-XX-XX-XX-XX-XX</li> <li>• Unformatted Lowercase – Format the MAC address as aaaabbbbcccc</li> <li>• Unformatted Uppercase – Format the MAC address as AAAABBBBCCCC</li> </ul>
CALLING-STATION-ID (Attribute 31)	<p>Specifies the format in which the MAC address is sent to the RADIUS server in attribute 31 (Calling-Station-ID). To specify a format, click the Edit icon and select one of the following:</p> <ul style="list-style-type: none"> <li>• Legacy Lowercase – Format the MAC address as xx:xx:xx:xx:xx:xx</li> <li>• Legacy Uppercase – Format the MAC address as XX:XX:XX:XX:XX:XX</li> <li>• IETF Lowercase – Format the MAC address as xx-xx-xx-xx-xx-xx</li> <li>• IETF Uppercase – Format the MAC address as XX-XX-XX-XX-XX-XX</li> <li>• Unformatted Lowercase – Format the MAC address as aaaabbbbcccc</li> <li>• Unformatted Uppercase – Format the MAC address as AAAABBBBCCCC</li> </ul>



Table 350: RADIUS Configuration Fields (Continued)

Field	Description
NAS-IDENTIFIER (Attribute 32 MAC Format)	<p>Specifies the format in which the MAC address is sent to the RADIUS server in attribute 32 (NAS-Identifier). To specify a format, click the Edit icon and select one of the following:</p> <ul style="list-style-type: none"> <li>Legacy Lowercase – Format the MAC address as xx:xx:xx:xx:xx:xx</li> <li>Legacy Uppercase – Format the MAC address as XX:XX:XX:XX:XX:XX</li> <li>IETF Lowercase – Format the MAC address as xx-xx-xx-xx-xx-xx</li> <li>IETF Uppercase – Format the MAC address as XX-XX-XX-XX-XX-XX</li> <li>Unformatted Lowercase – Format the MAC address as aaaabbbbcccc</li> <li>Unformatted Uppercase – Format the MAC address as AAAABBBBCCCC</li> </ul>
NAS-IDENTIFIER (Attribute 32 Include in Access/Accounting Request)	<p>Determines whether the RADIUS attribute 32 (NAS-Identifier) is sent to the RADIUS server in access-request and accounting-request messages and in which format. To configure the settings, click the Edit icon and configure the following:</p> <ul style="list-style-type: none"> <li>Include in Access/Accounting Request–When selected, the attribute is sent to the RADIUS server in access-request and accounting-request messages.</li> <li>Format–Configures the format of an optional string sent in access-request and accounting-request messages in attribute 32 (NAS-Identifier). The format can be one of the following: <ul style="list-style-type: none"> <li>- %m – MAC address</li> <li>- %i – IP address</li> <li>- %h – Host name</li> <li>- %d – Domain name</li> <li>- Any String – A string including any or all of the above formatting options</li> </ul> </li> </ul> <p>If you configure the format, the string sent in attribute 32 (NAS-Identifier) includes a MAC address, an IP address, a Host name or a Domain name based on the configured format.</p>
ACCT-SESSION-ID (Attribute 44)	<p>Determines whether the RADIUS attribute 44 (ACCT-SESSION-ID) is sent to the RADIUS server in access-request and accounting-request messages. To configure the settings, click the Edit icon and select the option to indicate that the attribute should be included in the messages. Clear the option to prevent the attribute from being sent.</p>
NAS-IPV6-ADDRESS (Attribute 95)	<p>The network access server (NAS) IPv6 address for the RADIUS server. If the specific IPv6 address is configured while enabling this attribute, the RADIUS client uses that IPv6 address while sending NAS-IPV6-Address attribute in RADIUS communication. The address should be unique to the NAS within the scope of the RADIUS server.</p>

Use the buttons at the bottom of the page to perform the following actions:

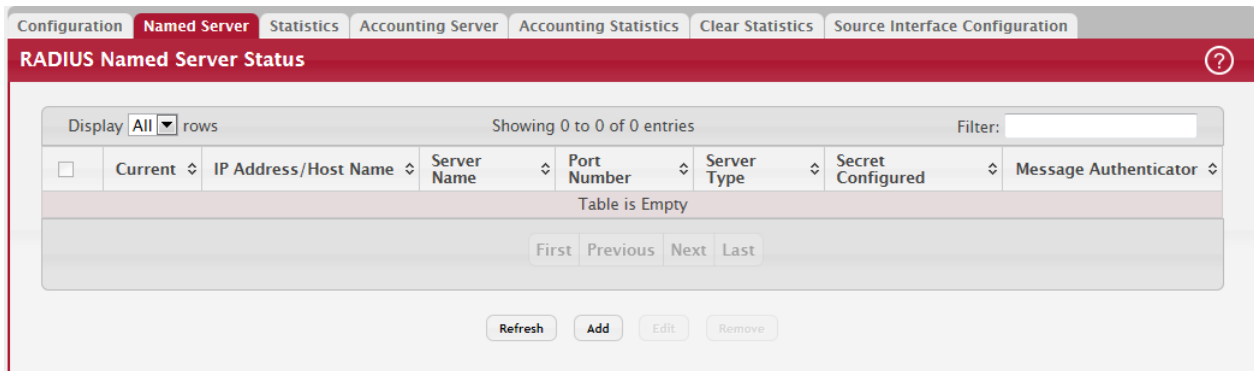
- Click Refresh to update the page with the most current information.
- If you make changes to the page, click Submit to apply the changes to the system.

### 7.3.1.1 Named Server Status

The RADIUS Named Server Status page shows summary information about the RADIUS servers configured on the system.

To access the RADIUS Named Server Status page, click Security > RADIUS > Named Server in the navigation menu.

Figure 367: Named Server Status



Use the buttons to perform the following tasks:

- To add a RADIUS authentication server to the list of servers the RADIUS client can contact, click Add.
- To change the settings for a configured RADIUS server, select the entry to modify and click Edit. You cannot change the IP address or host name for a server after it has been added.
- To remove a configured RADIUS server from the list, select the entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 351: RADIUS Server Status Fields

Field	Description
Current	An asterisk (*) in the column Indicates that the server is the current server for the authentication server group. If no asterisk is present, the server is a backup server. If more than one RADIUS server is configured with the same name, the switch selects one of the servers to be the current server from the group of servers with the same name. When the switch sends a RADIUS request to the named server, the request is directed to the server selected as the current server. Initially the primary server is selected as the current server. If the primary server fails, one of the other servers becomes the current server.
RADIUS Server Host Address	Shows the IP address of the RADIUS server.
RADIUS Server Name	Shows the RADIUS server name. Multiple RADIUS servers can have the same name. In this case, RADIUS clients can use RADIUS servers with the same name as backups for each other.
Port Number	Identifies the authentication port the server uses to verify the RADIUS server authentication. The port is a UDP port.
Server Type	Shows whether the server is a Primary or Secondary server.
Secret Configured	Indicates whether the shared secret for this server has been configured.
Message Authenticator	Shows whether the message authenticator attribute for the selected server is enabled or disabled.

Click Refresh to update the page with the most current information.

### 7.3.2 Server Statistics

Use the RADIUS Server Statistics page to view statistical information for each RADIUS server configured on the system. To access the RADIUS Server Statistics page, click Security > RADIUS > Statistics in the navigation menu.

Figure 368: RADIUS Server Statistics

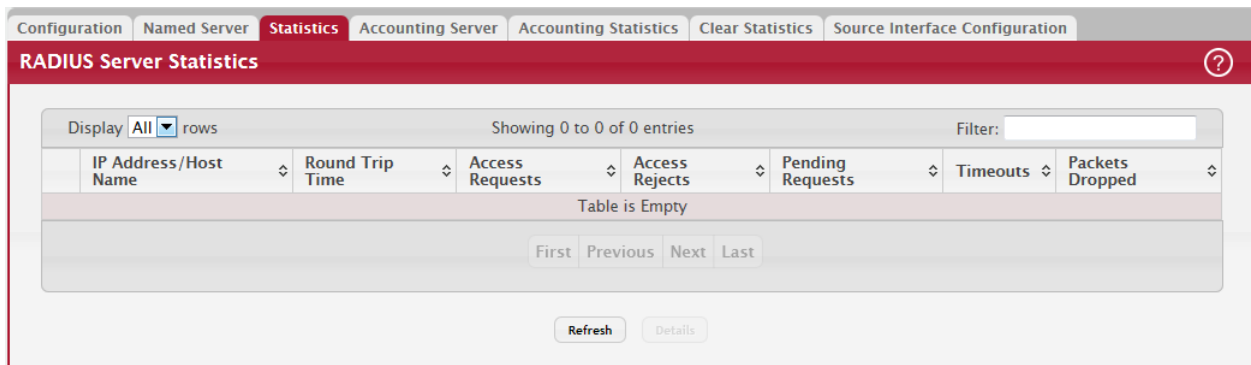


Table 352: RADIUS Server Statistics Fields

Field	Description
IP Address/Host Name	The IP address or host name of the RADIUS server associated with the rest of the data in the row. When viewing the detailed statistics for a RADIUS server, this field identifies the RADIUS server.
Round Trip Time	The time interval, in hundredths of a second, between the most recent Access-Reply/Access-Challenge and the Access-Request that matched it from the RADIUS authentication server.
Access Requests	The number of RADIUS Access-Request packets sent to the server. This number does not include retransmissions.
Access Rejects	The number of RADIUS Access-Reject packets, including both valid and invalid packets, that were received from the server.
Pending Requests	The number of RADIUS Access-Request packets destined for the server that have not yet timed out or received a response.
Timeouts	The number of times a response was not received from the server within the configured timeout value.
Packets Dropped	The number of RADIUS packets received from the server on the authentication port and dropped for some other reason.
Access Retransmissions	The number of RADIUS Access-Request packets that had to be retransmitted to the server because the initial Access-Request packet failed to be successfully delivered.
Access Accepts	The number of RADIUS Access-Accept packets, including both valid and invalid packets, that were received from the server.
Access Challenges	The number of RADIUS Access-Challenge packets, including both valid and invalid packets, that were received from the server.
Malformed Access Responses	The number of malformed RADIUS Access-Response packets received from the server. Malformed packets include packets with an invalid length. Bad authenticators, signature attributes, and unknown types are not included as malformed access responses.
Bad Authenticators	The number of RADIUS Access-Response packets containing invalid authenticators or signature attributes received from the server.
Unknown Types	The number of RADIUS packets of unknown type which were received from the server on the authentication port.

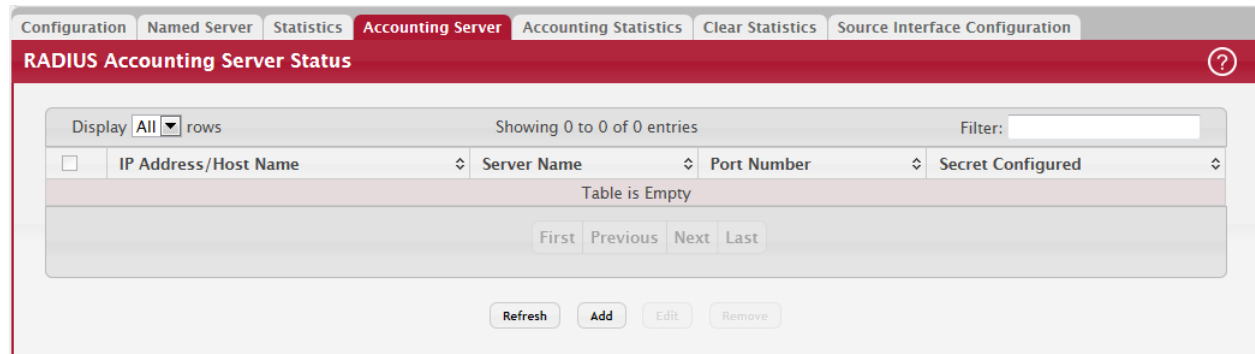
Click Refresh to update the page with the most current information.

### 7.3.3 Named Accounting Server Status

The RADIUS Named Accounting Server Status page shows summary information about the accounting servers configured on the system.

To access the RADIUS Accounting Server Status page, click Security > RADIUS > Accounting Server in the navigation menu.

Figure 369: RADIUS Accounting Server Status



Use the buttons to perform the following tasks:

- To add a RADIUS accounting server to the list of servers the RADIUS client can contact, click Add.
- To change the settings for a configured RADIUS accounting server, select the entry to modify and click Edit. You cannot change the IP address or host name for a server after it has been added.
- To remove a configured RADIUS accounting server from the list, select the entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 353: RADIUS Accounting Server Status Fields

Field	Description
IP Address/Host Name	The IP address or host name of the RADIUS accounting server. Host names must be resolvable by DNS and are composed of a series of labels separated by dots.
Server Name	The name of the RADIUS accounting server. RADIUS servers that are configured with the same name are members of the same named RADIUS server group. RADIUS accounting servers in the same group serve as backups for each other.
Port Number	The UDP port on the RADIUS accounting server to which the local RADIUS client sends request packets.
Secret Configured	Indicates whether the shared secret for this server has been configured.
Secret	The shared secret text string used for authenticating and encrypting all RADIUS communications between the RADIUS client on the device and the RADIUS accounting server. The secret specified in this field must match the shared secret configured on the RADIUS accounting server.

Click Refresh to update the page with the most current information.

### 7.3.4 Accounting Statistics

Use the RADIUS Accounting Statistics page to view statistical information for each RADIUS server configured on the system.

To access the RADIUS Accounting Statistics page, click Security > RADIUS > Accounting Statistics in the navigation menu.

Figure 370: RADIUS Accounting Statistics

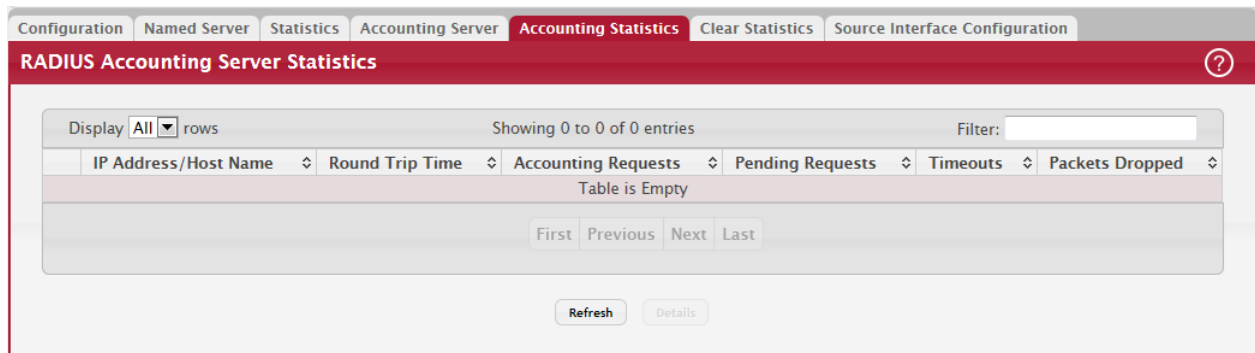



Table 354: RADIUS Accounting Statistics Fields

Field	Description
IP Address/Host Name	The IP address or host name of the RADIUS accounting server associated with the rest of the data in the row. When viewing the detailed statistics for a RADIUS accounting server, this field identifies the server.
Round Trip Time	Displays the time interval, in hundredths of a second, between the most recent Accounting-Response and the Accounting-Request that matched it from this RADIUS accounting server.
Accounting Requests	The number of RADIUS Accounting-Request packets sent to this server. This number does not include retransmissions.
Pending Requests	The number of RADIUS Accounting-Request packets destined for the server that have not yet timed out or received a response.
Timeouts	The number of times a response was not received from the server within the configured timeout value.
Packets Dropped	The number of RADIUS packets received from the server on the accounting port and dropped for some other reason.
Accounting Retransmissions	The number of RADIUS Accounting-Request packets retransmitted to the server.
Accounting Responses	The number of RADIUS packets received on the accounting port from the server.
Timeouts	The number of accounting timeouts to this server.
Malformed Access Responses	The number of malformed RADIUS Accounting-Response packets received from the server. Malformed packets include packets with an invalid length. Bad authenticators and unknown types are not included as malformed accounting responses.
Bad Authenticators	The number of RADIUS Accounting-Response packets that contained invalid authenticators received from the accounting server.
Unknown Types	The number of RADIUS packets of unknown type which were received from the server on the accounting port.

### 7.3.5 Clear Statistics

Use the RADIUS Clear Statistics page to reset all RADIUS authentication and accounting statistics to zero. To access the RADIUS Clear Statistics page, click Security > RADIUS > Clear Statistics in the navigation menu.

Figure 371: RADIUS Clear Statistics



To clear all statistics for the RADIUS authentication and accounting server, click Reset. After you confirm the action, the statistics on both the RADIUS Server Statistics and RADIUS Accounting Server Statistics pages are reset.

### 7.3.6 Source Interface Configuration

Use this page to specify the physical or logical interface to use as the RADIUS client source interface. When an IP address is configured on the source interface, this address is used for all RADIUS communications between the local RADIUS client and the remote RADIUS server. The IP address of the designated source interface is used in the IP header of RADIUS management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access the RADIUS Source Interface Configuration page, click Security > RADIUS > Source Interface Configuration in the navigation menu.

Figure 372: RADIUS Source Interface Configuration

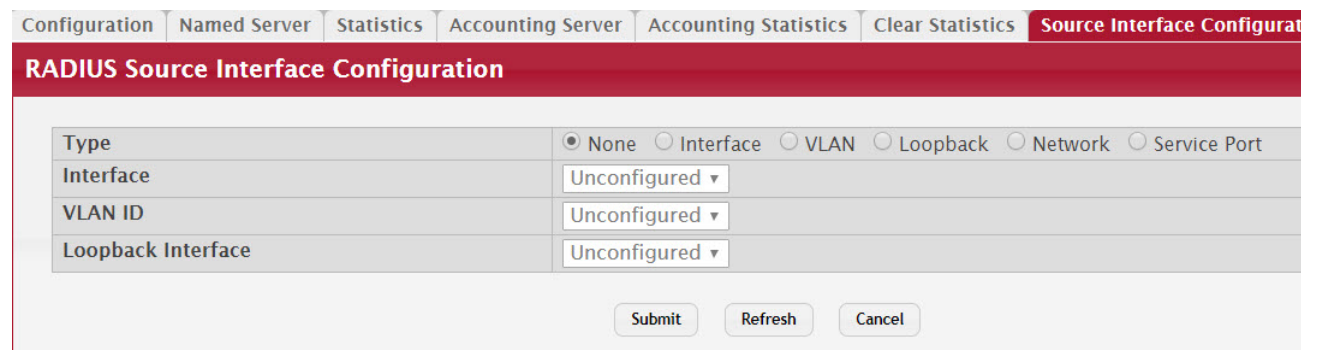


Table 355: RADIUS Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> <li>None – The primary IP address of the originating (outbound) interface is used as the source address.</li> <li>Interface – The primary IP address of a physical port is used as the source address.</li> <li>Loopback – The primary IP address of the loopback interface is used as the source address. A loopback is always reachable, as long as any routing interface is up.</li> <li>Network – The network source IP is used as the source address.</li> <li>Service Port – The management port source IP is used as the source address.</li> <li>VLAN – The primary IP address of a VLAN routing interface is used as the source address.</li> </ul>
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
Loopback Interface	When the selected Type is Loopback, select the loopback interface to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.

Click Refresh to update the page with the most current information.

## 7.4 TACACS+ Settings

To access the TACACS+ Configuration page, click Security > TACACS+ > Configuration in the navigation menu.

Figure 373: TACACS+ Configuration

Table 356: TACACS+ Configuration Fields

Field	Description
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the key configured on the TACACS+ server.
Connection Timeout	The maximum number of seconds allowed to establish a TCP connection between the device and the TACACS+ server.

Click Refresh to update the page with the most current information.

If you make any changes to the page, click Submit to apply the changes to the system.

### 7.4.1 TACACS+ Server Summary

Use this page to view and configure information about the TACACS+ Server(s).

To access the TACACS+ Server Summary page, click Security > TACACS+ > Server Summary in the navigation menu.

Figure 374: TACACS+ Server Summary

Use the buttons to perform the following tasks:

- To add a TACACS+ Server to the list of servers the TACACS+ client can contact, click Add. If maximum number of server is added, the button will be disabled
- To edit a configured TACACS+ server from the list, select the entry and click Edit.
- To remove a configured TACACS+ server from the list, select the entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 357: TACACS+ Server Summary Fields

Field	Description
Server	Specifies the TACACS+ Server IP address or Hostname.
Priority	Specifies the order in which the TACACS+ servers are used.
Port	Specifies the authentication port.
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the encryption used on the TACACS+ server.
Connection Timeout	The amount of time that passes before the connection between the device and the TACACS+ server time out.

Click Refresh to update the page with the most current information.

## 7.4.2 TACACS+ Server Configuration

Use this page to view and configure information about the TACACS+ Server(s).

To access the TACACS+ Server Configuration page, click Security > TACACS+ > Server Configuration in the navigation menu.

Figure 375: TACACS+ Server Configuration

The screenshot shows the 'TACACS+ Server Configuration' page. The form contains the following data:

Field	Value	Range/Constraint
Server	Test	
Priority	0	(0 to 65535)
Port	49	(0 to 65535)
Key String		
Connection Timeout	5	(1 to 30 secs)

Buttons: Submit, Remove, Refresh, Cancel

Table 358: TACACS+ Server Configuration Fields

Field	Description
Server	Specifies the TACACS+ Server IP address or Hostname.
Priority	Specifies the order in which the TACACS+ servers are used.
Port	Specifies the authentication port.
Key String	Specifies the authentication and encryption key for TACACS+ communications between the device and the TACACS+ server. The key must match the encryption used on the TACACS+ server.
Connection Timeout	The amount of time that passes before the connection between the device and the TACACS+ server time out.

Click Refresh to update the page with the most current information.

If you make any changes to the page, click Submit to apply the changes to the system.



### 7.4.3 TACACS+ Source Interface Configuration

Use this page to specify the physical or logical interface to use as the TACACS+ client source interface. When an IP address is configured on the source interface, this address is used for all TACACS+ communications between the local TACACS+ client and the remote TACACS+ server. The IP address of the designated source interface is used in the IP header of TACACS+ management protocol packets. This allows security devices, such as firewalls, to identify all source packets coming from a specific device.

To access the TACACS+ Source Interface Configuration page, click Security > TACACS+ > Source Interface Configuration in the navigation menu.

Figure 376: TACACS+ Source Interface Configuration

Table 359: TACACS+ Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: <ul style="list-style-type: none"> <li>None – The primary IP address of the originating (outbound) interface is used as the source address.</li> <li>Interface – The primary IP address of a physical port is used as the source address.</li> <li>VLAN – The primary IP address of a VLAN routing interface is used as the source address.</li> <li>Network – The network source IP is used as the source address.</li> <li>Service Port – The management port source IP is used as the source address.</li> </ul>
Interface	When the selected Type is Interface, select the physical port to use as the source interface.
VLAN ID	When the selected Type is VLAN, select the VLAN to use as the source interface. The menu contains only the VLAN IDs for VLAN routing interfaces.

Click Refresh to update the page with the most current information.

If you make any changes to the page, click Submit to apply the changes to the system.

## 7.5 Authentication Manager

The Authentication Manager feature allows you to configure the authentication methods used on the individual interface.

### 7.5.1 Authentication Manager Configuration

Use this page to control the administrative mode of the Authentication Manager feature, which enables configuration of the sequence and priority of the authentication methods per interface.

Authentication Manager supports the Dynamic Authorization component for Change of Authorization (CoA) requests from the DAS for the matching sessions. The following support is available:

- Change of the client VLAN
- Client re-authentication
- Change of the Filter-ID for client
- Change of the Downloadable Access Control List (DACL) for client

The following CoA requests result in termination of all sessions on the matching port:

- Bounce port
- Disable port

To access the Authentication Manager Configuration page, click Security > Authentication Manager> Configuration in the navigation menu.

Figure 377: Authentication Manager Configuration

Field	Description
Admin Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Dynamic VLAN Creation Mode	<input type="radio"/> Disable <input type="radio"/> Enable
VLAN Assignment Mode	<input type="radio"/> Disable <input type="radio"/> Enable
Authentication Monitor Mode	<input type="radio"/> Disable <input type="radio"/> Enable
Critical Recovery Max Re-Authentication	<input type="text" value="10"/> (1 to 50)
CoA Bounce Host Port	<input type="radio"/> Accept <input type="radio"/> Reject
CoA Disable Host Port	<input type="radio"/> Accept <input type="radio"/> Reject
Authenticated Clients	<input type="text" value="0"/>
Clients in Monitor Mode	<input type="text" value="0"/>

Table 360: Authentication Manager Configuration Fields

Field	Description
Admin Mode	The administrative mode of the Authentication Manager feature. When Authentication Manager is enabled globally, the authentication methods configured on an interface are executed in the order configured as the device attempts to authenticate clients on that interface.
Dynamic VLAN Creation Mode	The administrative mode of dynamic VLAN creation on the device. If RADIUS-assigned VLANs are enabled, the RADIUS server is expected to include the VLAN ID in the 802.1X tunnel attributes of its response message to the device. If dynamic VLAN creation is enabled on the device and the RADIUS-assigned VLAN does not exist, then the assigned VLAN is dynamically created. This implies that the client can connect from any port and can get assigned to the appropriate VLAN. This feature gives flexibility for clients to move around the network without much additional configuration required.
VLAN Assignment Mode	The administrative mode of RADIUS-based VLAN assignment on the device. When enabled, this feature allows a port to be placed into a particular VLAN based on the result of the authentication or type of 802.1X authentication a client uses when it accesses the device. The authentication server can provide information to the device about which VLAN to assign the client.

Table 360: Authentication Manager Configuration Fields (Continued)

Field	Description
Authentication Monitor Mode	The administrative mode of the Monitor Mode feature on the device. Monitor mode is a special mode that can be enabled in conjunction with port-based access control. Monitor mode provides a way for network administrators to identify possible issues with the port-based access control configuration on the device without affecting the network access to the users of the device. It allows network access even in cases where there is a failure to authenticate, but it logs the results of the authentication process for diagnostic purposes. If the device fails to authenticate a client for any reason (for example, RADIUS access reject from the RADIUS server, RADIUS timeout, or the client itself is 802.1X unaware), the client is authenticated and is undisturbed by the failure condition(s). The reasons for failure are logged and buffered into the local logging database for tracking purposes.
Critical Recovery Max Re-Authentication	The number of critical recovery maximum client re-authentications per second.
CoA Bounce Host Port	The administrative mode of the Change of Authorization Bounce Host Port feature on the device. When set to Reject, the device will ignore a RADIUS server <code>bounce-host-port</code> command. The <code>bounce-host-port</code> command causes a host to flap the link on an authentication port. The link flap causes DHCP renegotiation from one or more hosts connected to this port.
CoA Disable Host Port	The administrative mode of the Change of Authorization Disable Host Port feature on the device. The <code>disable-host-port</code> command puts the host port in a disabled state with the reason as CoA disabled. The disabled port can be re-enabled using one of the following methods: <ul style="list-style-type: none"> <li>If CoA Disable Host Port auto recovery is enabled on the Port Auto Recovery Configuration page, the port is re-enabled after the auto recovery timer expires. See <a href="#">Section 5.11.1: "Port Auto Recovery Configuration"</a>.</li> <li>The administrator manually re-enables the port on the Port Summary page. See <a href="#">Section 4.13.1: "Port Summary"</a>.</li> </ul>
Authenticated Clients	The total number of clients authenticated on the switch except the ones in the Monitor mode.
Clients in Monitor Mode	The number of clients authorized by the Monitor mode on the switch.

Use the buttons at the bottom of the page to perform the following actions:

- Click Submit to apply the settings to the running configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click Refresh to display the latest information from the switch.
- Click Cancel to cancel the change.

## 7.5.2 Authentication Manager Interface Configuration

Use this page to configure the Authentication Manager interface. The Open Authentication capability allows Authentication Manager to allow client traffic even before it authenticates. This is typically used to allow certain devices access to network resources prior to authenticating to obtain the IP address and download configuration or firmware upgrades. After the information is downloaded, the device will authenticate to the network.

Open Authentication is configured per interface. The Open Authentication settings are ignored for force-authorized and force-unauthorized ports. Open Authentication is supported for all switch port modes (Access, General, and Trunk). It is also supported in all Authentication Manager Host modes. The number of clients that are given open access before authentication is limited by the configured host mode on the port.

Before authentication completes for a client, it is allowed access to Open mode on the data VLAN of the port. A client authorized in Open mode is considered a data client. A client authentication will eventually trigger, based on the available and configured authentication methods on the port, and on the reception of Extensible Authentication Protocol (EAP) packets from the client. The authentication can trigger even before the client is given port access.

To access the Authentication Manager Interface Configuration page, click Security > Authentication Manager > Interface Configuration in the navigation menu.

Figure 378: Authentication Manager Interface Configuration

Table 361: Authentication Manager Interface Configuration Fields

Field	Description
Interface	The interface with the settings to view or configure.
Control Mode	The authentication control mode on the port, which is one of the following: <ul style="list-style-type: none"> <li>Force Unauthorized–The port ignores supplicant authentication attempts and does not provide authentication services to the client.</li> <li>Force Authorized–The port sends and receives normal traffic without client port-based authentication.</li> <li>Auto–The port is unauthorized until a successful authentication exchange has taken place.</li> </ul>
Host Mode	The authentication host mode on the port determines the number and type of clients that can be authenticated and authorized on the port. The port host mode can be one of the following: <ul style="list-style-type: none"> <li>Single Authentication–Only one data client can be authenticated on a port and the client is granted access to the port.</li> <li>Multiple Host–Only one data client can be authenticated on a port. However, once authentication succeeds, access is granted to all clients connected to the port.</li> <li>Multiple Domain–One data client and one voice client can be authenticated on a port and both clients are granted access to the port.</li> <li>Multiple Authentication–One voice client and multiple data clients can be authenticated on a port and these clients are granted access to the port.</li> <li>Multiple Domain/Host–One voice client and one data client can be authenticated on a port and these clients are granted access to the port. However, once a data client is authenticated, access is granted to all clients connected to the port and they are considered data clients.</li> </ul>

Table 361: Authentication Manager Interface Configuration Fields (Continued)

Field	Description
Re-Authentication	Indicates if the connected clients can re-authenticate periodically.
Re-Authentication Period (Seconds)	The amount of time that clients can be connected to the port without being re-authenticated. If Re-Authentication is disabled, connected clients are not forced to re-authenticate periodically.
Re-Authentication Timeout from Server	The amount of time, obtained from the RADIUS server, that clients can be connected to the port without being re-authenticated.
Maximum Users	The maximum number of clients supported on the port.
Guest VLAN ID	The VLAN ID for the guest VLAN. The guest VLAN allows the port to provide a distinguished service to unauthenticated users. This feature provides a mechanism to allow users access to hosts on the guest VLAN.
Authentication Retry Attempts	The maximum number of failed client authentication attempts on the port.
Unauthenticated VLAN ID	The VLAN ID of the unauthenticated VLAN. Hosts that fail the authentication might be denied access to the network or placed on a VLAN created for unauthenticated clients. This VLAN might be configured with limited network access and is used for 802.1X aware clients only.
Authentication Violation Mode	The authentication violation mode on the port. The authentication violation can occur when a device tries to connect to a port on which the maximum number of devices has exceeded. Action taken on the port when a security violation occurs can be one of the following: <ul style="list-style-type: none"> <li>• Protect</li> <li>• Restrict</li> <li>• Shutdown</li> </ul>
Authentication Server Alive Action	The action configured on the RADIUS server that is alive after all are dead. The alive-server action can be one of the following: <ul style="list-style-type: none"> <li>• Reinitialize–Dot1x triggers the re-authentication of clients authenticated on the critical VLAN.</li> <li>• None–No action is configured.</li> </ul>
Authentication Server Dead Action for Voice	The action configured to allow critical voice VLAN support on the port when all the RADIUS servers are marked dead. The dead-server action can be one of the following: <ul style="list-style-type: none"> <li>• Authorize–Allows port access on the voice VLAN when all RADIUS servers are dead.</li> <li>• None–No action is configured.</li> </ul>
Authentication Server Dead Action	The action configured on the RADIUS server that is marked dead. The dead-server action can be one of the following: <ul style="list-style-type: none"> <li>• Reinitialize–Authentication Manager triggers re-authentication of all authenticated clients on the port. Supplicants on voice VLAN, unauthenticated VLAN, and guest VLAN are not disturbed. During re-authentication if all the servers are still dead, the client is authenticated successfully and placed on the critical VLAN.</li> <li>• Authorize– Dot1x authorizes the authenticated clients to the critical VLAN. Clients on the RADIUS assigned VLAN, voice VLAN, unauthenticated VLAN, and guest VLAN are not disturbed. Clients authorized on the port PVID are re-authorized on the critical VLAN.</li> <li>• None–No action is configured.</li> </ul>
Critical VLAN ID	The VLAN ID of the critical VLAN. Critical VLAN allows supplicants to authenticate on the VLAN when all RADIUS servers are dead.
MAB Mode	The MAC-based Authentication Bypass (MAB) mode on the port, which can be enabled or disabled.
Operational MAB Mode	The operational MAB mode on the port.

Table 361: Authentication Manager Interface Configuration Fields (Continued)

Field	Description
MAB Authentication Type	The authentication type to be used for MAB access requests sent to the RADIUS server, which is one of the following: <ul style="list-style-type: none"> <li>• CHAP—The port uses CHAP authentication and sends a randomly generated 16-octet challenge as the CHAP-Challenge (RADIUS attribute 60) along with the CHAP-Password (RADIUS attribute 3) to the authentication server.</li> <li>• EAP-MD5—The port uses EAP-MD5 authentication and sends the MD5 hash of the MAC address as the password in the EAP-Message (RADIUS attribute 79) to the authentication server.</li> <li>• PAP—The port uses PAP authentication and sends the MAC address of the client as the password (clear text) in the User-Password (RADIUS attribute 2) to the authentication server.</li> </ul>
Open Authentication	Enable or disable open authentication on the specified interface. Open authentication permits client traffic on the data VLAN prior to port authentication. This is typically used to allow access to network resources, such as DHCP, configuration download, or firmware upgrade. After the information is downloaded, the device will proceed with normal authentication. Open authentication settings are ignored for force-authorized and force-unauthorized ports.
Allow Protocols When Unauthorized	Allows the specified protocol on the port when this field is enabled and the port is unauthorized. If enabled, DHCP packets entering the port are sent to the CPU to be processed to determine if the packet is allowed to ingress from that port. If the DHCP packet is not allowed, it is ignored. If allowed, the packet is forwarded based on a matching entry in the forwarding database or flooded to the VLAN. If DHCP snooping is enabled on this port, processing defers to DHCP snooping rules.

Use the buttons at the bottom of the page to perform the following actions:

- Click Submit to apply the settings to the running configuration and cause the change to take effect. These changes will not be retained across a power cycle unless a Save configuration is performed.
- Click Refresh to display the latest information from the switch.
- Click Cancel to cancel the change.

### 7.5.3 Authentication Tiering

Use this page to configure the sequence and priority of the authentication methods for the interfaces on the device. When Authentication Manager is enabled globally, the authentication methods configured on an interface are executed in the order configured as the device attempts to authenticate clients on that interface. The default method order is Dot1x, MAC Authentication Bypass (MAB), and Captive Portal.

To access the Authentication Tiering page, click Security > Authentication Manager > Authentication Tiering in the navigation menu.

Figure 379: Authentication Tiering

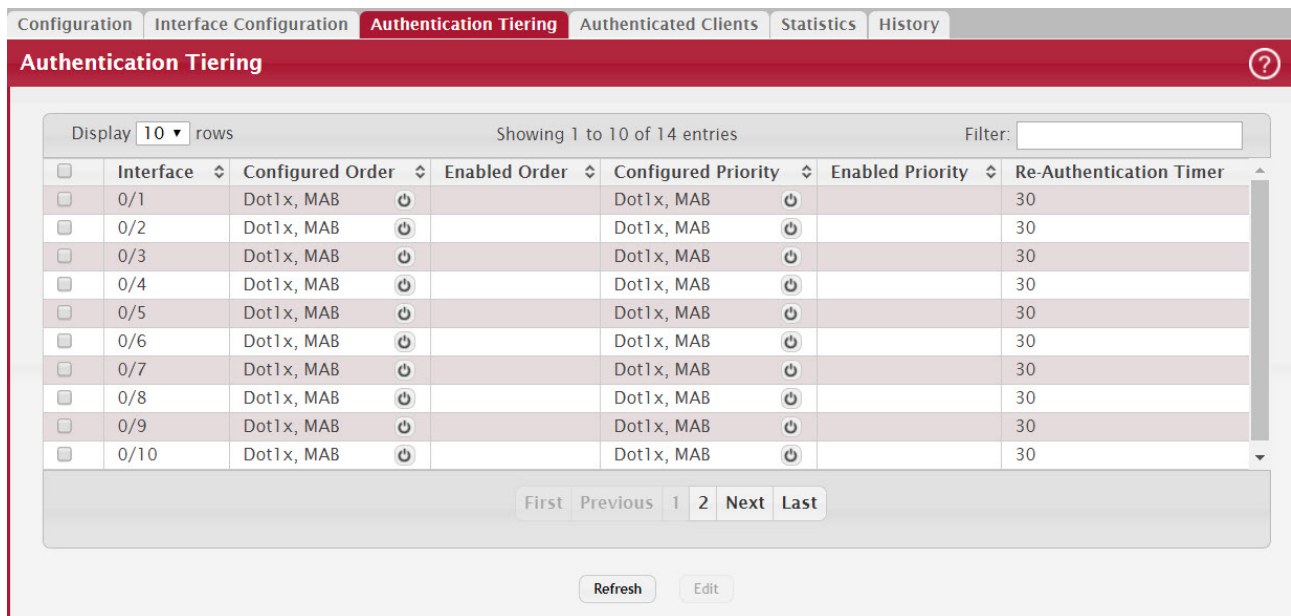


Table 362: Authentication Tiering Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When editing information for one or more interfaces, this field identifies the interfaces that are being configured.
Configured Order	The order in which the authentication methods are used to authenticate a client connected to an interface, which can be one or more of the following: <ul style="list-style-type: none"> <li>• Dot1x – The port-based authentication method.</li> <li>• MAB – MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide.</li> <li>• Captive Portal – The authentication method that prevents clients from accessing the network until user verification has been established.</li> </ul> Captive portal must always be the last method in the list.
Enabled Order	The methods from the list of authentication methods configured on an interface which are administratively enabled in the device.
Configured Priority	The priority of the authentication methods. The default priority of a method is equivalent to its position in the order of the authentication list configured per interface. If the priority of the methods is changed, all clients authenticated using a lower priority method are forced to re-authenticate.
Enabled Priority	The methods from the list of authentication method priorities configured on an interface which are administratively enabled in the device.
Re-Authentication Timer	Interval, in seconds, after which an attempt is made to authenticate an unauthorized port.

Use the buttons at the bottom of the page to perform the following actions:

- Click Refresh to display the latest information from the switch.
- Click Edit to configure the settings for one or more interfaces, select each entry to modify. The settings are applied to all selected interfaces.



### 7.5.4 Authenticated Clients

Use this page to view information about the clients connected on the interfaces. If there are no clients connected, the table is empty.

To access the Authentication Clients page, click Security > Authentication Manager> Authenticated Clients in the navigation menu.

Figure 380: Authenticated Clients

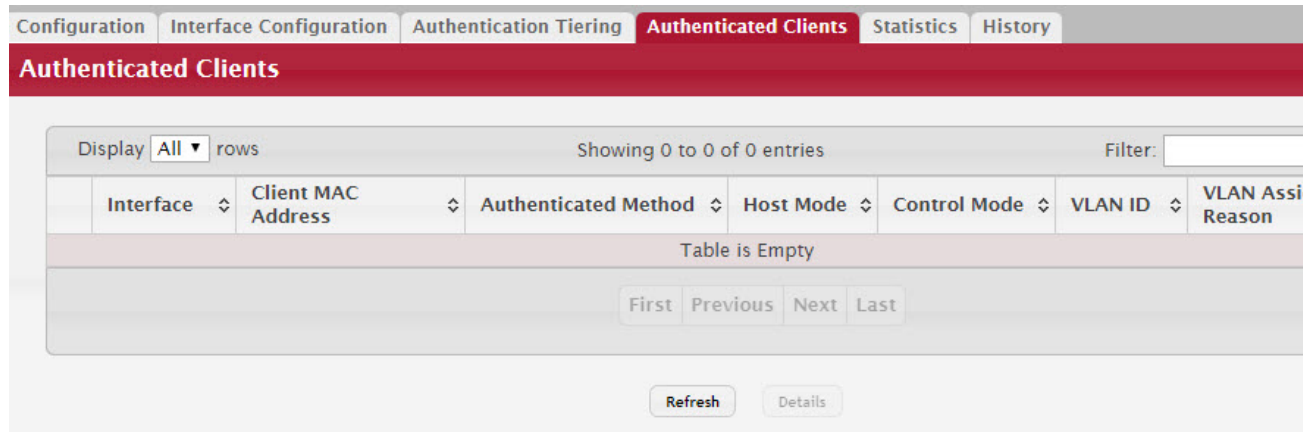


Table 363: Authenticated Clients Fields

Field	Description
Interface	The local interface associated with the rest of the data in the row.
Client MAC Address	The MAC address of the client that is connected to the port.
Authenticated Method	The authentication method used to authenticate a client connected to an interface, which can be one of the following: <ul style="list-style-type: none"> <li>• Dot1x – The port-based authentication method.</li> <li>• MAB – MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide.</li> <li>• None – The authentication method is undefined.</li> <li>• Captive Portal – The authentication method that prevents clients from accessing the network until user verification has been established.</li> </ul>
Host Mode	The authentication host mode on the port determines the number and type of clients that can be authenticated and authorized on the port. The port's host mode can be one of the following: <ul style="list-style-type: none"> <li>• Single Authentication – Only one data client or one voice client can be authenticated on a port, and the client is granted access to the port.</li> <li>• Multiple Host – Only one data client can be authenticated on a port. However, when authentication succeeds, access is granted to all clients connected to the port.</li> <li>• Multiple Domain – One data client and one voice client can be authenticated on a port, and both clients are granted access to the port.</li> <li>• Multiple Authentication – One voice client and multiple data clients can be authenticated on a port, and these clients are granted access to the port.</li> <li>• Multiple Domain/Host – One voice client and one data client can be authenticated on a port, and these clients are granted access to the port. However, when a data client is authenticated, access is granted to all clients connected to the port and they are considered data clients.</li> </ul>



**Table 363: Authenticated Clients Fields (Continued)**

Field	Description
Control Mode	The authentication control mode on the port, which is one of the following: <ul style="list-style-type: none"> <li>Force Unauthorized – The port ignores client authentication attempts and does not provide authentication services to the client.</li> <li>Force Authorized – The port sends and receives normal traffic without client port-based authentication.</li> <li>Auto – The port is unauthorized until a successful authentication exchange has taken place.</li> </ul>
VLAN ID	The ID of the VLAN in which the client was placed as a result of the authentication process.
VLAN Assigned Reason	The reason that the VLAN identified in the VLAN ID field has been assigned to the port, which can be one of the following: <ul style="list-style-type: none"> <li>Default Assigned VLAN – The client is authenticated on the port in the default VLAN, and the authentication server is not RADIUS.</li> <li>RADIUS Assigned VLAN – RADIUS is used for authenticating the client.</li> <li>Unauthenticated VLAN – The client is authenticated on the unauthenticated VLAN.</li> <li>Guest VLAN – The client is authenticated on the guest VLAN.</li> <li>Voice VLAN – The client is authenticated on the voice VLAN.</li> <li>Monitor Mode VLAN – The client is authenticated by the monitor mode.</li> <li>Critical VLAN – The client is authenticated on the critical VLAN.</li> <li>Not Assigned – No VLAN is assigned to the port.</li> </ul>

After you click Details, a window opens and displays additional information about the client. The following information describes the additional fields that appear in the window.

**Table 364: Authenticated Client Details**

Fields	Description
User Name	The name the client uses to identify itself as a supplicant to the authentication server.
Type	The type of the VLAN the client was placed in as a result of the authentication process, which can be either Data or Voice VLAN.
Session Time	The amount of time that has passed since the connected supplicant was granted access to the network through the authenticator port.
Session Timeout	The reauthentication timeout period set by the RADIUS server to the supplicant device.
Session Termination Action	The termination action set by the RADIUS server that indicates the action that will take place once the supplicant reaches the session timeout value.
Filter ID	The policy filter ID assigned by the authenticator to the supplicant device. This is a configured DiffServ policy name on the switch.
Accounting Session ID	The Accounting Session Id associated with the client session.
Downloadable Access Control List	Identifies the Downloadable Access Control List (DACL) returned by the RADIUS server when the client was authenticated.

Use the buttons at the bottom of the page to perform the following actions:

- Click Refresh to display the latest information from the switch.

## 7.5.5 Authenticated Statistics

Use this page to view information about the Authentication Manager client authentication attempts and failures per interface.

To access the Authentication Statistics page, click Security > Authentication Manager > Statistics in the navigation menu.

Figure 381: Authentication Statistics

The screenshot displays the 'Authentication Statistics' page. At the top, there are tabs for Configuration, Interface Configuration, Authentication Tiering, Authenticated Clients, **Statistics**, and History. Below the tabs, the page title 'Authentication Statistics' is shown with a help icon. The main content area features a table with the following data:

Interface	Dot1x Attempts	Dot1x Failures	MAB Attempts	MAB Failures
0/1	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
0/9	0	0	0	0
0/10	0	0	0	0

Below the table, there are navigation buttons: First, Previous, 1, 2, Next, Last. At the bottom, there are 'Refresh' and 'Clear' buttons. The interface also shows 'Display 10 rows' and 'Showing 1 to 10 of 14 entries'.

Table 365: Authentication Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Dot1x Attempts	The number of attempts made to authenticate a client using the Dot1x authentication method.
Dot1x Failures	The number of attempts that failed when Dot1x method is used for client authentication.
MAB Attempts	The number of attempts made to authenticate a client using the MAC Authentication Bypass (MAB) authentication method.
MAB Failures	The number of attempts that failed when MAB method is used for client authentication.
Captive Portal Attempts	The number of attempts made to authenticate a client using the Captive Portal authentication method.
Captive Portal Failures	The number of attempts that failed when the Captive Portal method is used for client authentication.
ACS ACL Name	The downloadable ACL returned by the RADIUS server when the client was authenticated.
Downloadable Access Control List	Identifies the Dynamic Access Control List returned by the RADIUS server when the client was authenticated.
Redirect ACL	The static ACL sent in the RADIUS attribute redirect-acl. It is used to redirect matching packets to the CPU for further action.
Redirect URL	The URL sent in the RADIUS attribute redirect-url. It is used to redirect matching packets to the redirect URL by using HTTP 302 response code.

Use the buttons at the bottom of the page to perform the following actions:

- Click Refresh to display the latest information from the switch.
- Click Clear to reset all statistics counters to 0 for the selected interfaces.

## 7.5.6 Authenticated History

Use this page to view the Authentication Manager history log per interface.

To access the Authentication History page, click Security > Authentication Manager> History in the navigation menu.

Figure 382: Authentication History

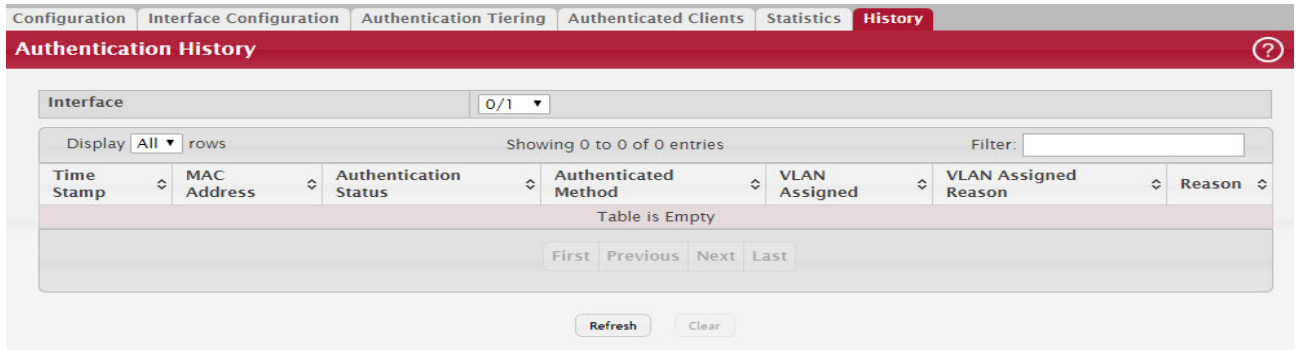


Table 366: Authentication History Fields

Field	Description
Interface	The menu contains all interfaces in the device. To view the history log on a specific interface, select the interface from the menu.
Time Stamp	The absolute time when the authentication event took place.
MAC Address	The MAC address of the client that is connected to the port.
Authentication Status	The client authentication status, which can be one of the following: <ul style="list-style-type: none"> <li>• Authorized – Indicates client is authorized on the port.</li> <li>• Unauthorized – Indicates client is not authorized on the port.</li> </ul>
Authenticated Method	The authentication method used to authenticate a client connected to an interface, which can be one of the following: <ul style="list-style-type: none"> <li>• Dot1x – The port-based authentication method.</li> <li>• MAB – MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide.</li> <li>• Captive Portal – The authentication method that prevents clients from accessing the network until user verification has been established.</li> <li>• None - The authentication method is undefined.</li> </ul>
VLAN Assigned	The ID of the VLAN in which the client is placed as a result of the authentication process.

**Table 366: Authentication History Fields (Continued)**

Field	Description
VLAN Assigned Reason	<p>The reason why the authenticator placed in the client in the VLAN, which can be one of the following:</p> <ul style="list-style-type: none"> <li>• Default Assigned VLAN – The client is authenticated on the port in the default VLAN, and the authentication server is not RADIUS.</li> <li>• RADIUS Assigned VLAN – RADIUS is used for authenticating the client.</li> <li>• Unauthenticated VLAN – The client is authenticated on the unauthenticated VLAN.</li> <li>• Guest VLAN – The client is authenticated on the guest VLAN.</li> <li>• Voice VLAN – The client is authenticated on the voice VLAN.</li> <li>• Monitor Mode VLAN – The client is authenticated by the monitor mode.</li> <li>• Critical VLAN – The client is authenticated on the critical VLAN.</li> <li>• Not Assigned – No VLAN is assigned to the port.</li> </ul>
Reason	The reason for the successful or unsuccessful authentication.

Use the buttons at the bottom of the page to perform the following actions:

- Click Refresh to display the latest information from the switch.
- Click Clear to clear the Authentication Manager history log on the selected interface.

## 7.6 Media Access Control Security

The Media Access Control Security (MACsec) feature provides secure communications between stations that are attached to the same LAN. It uses symmetric key cryptography so that communication cannot be monitored or altered on the wire. Traffic traversing the link is MACsec-secured through the use of data integrity checks and encryption.

MACsec is standardized in IEEE 802.1AE, which specifies the forwarding plane for MACsec. The key management and its integration with 802.1x is standardized in IEEE 802.1X-2010.

The data integrity checks verify the integrity of the data. MACsec appends an 8-byte header called Sectag and a 16-byte tail called ICV to all secured Ethernet frames. These are checked by the receiving interface to ensure that the data was not compromised while traversing the link.

MACsec can also be used to encrypt all traffic on the Ethernet link. The encryption used by MACsec ensures that the data in the Ethernet frame cannot be viewed by anyone monitoring traffic on the link. MACsec encryption is optional and user-configurable—you can enable MACsec to ensure the data integrity checks are performed while still sending unencrypted data in the clear over the MACsec-secured link, if desired.

FASTPATH MACsec supports confidentiality through encryption, integrity, and replay protection.

For information about the CLI commands you use to configure MACsec features, refer to the FASTPATH CLI Command Reference in [Related Documents](#).

### 7.6.1 MACsec Key Agreement Policy Summary

Use the MKA Policy Summary page to view and configure the MACsec Key Agreement (MKA) policies. For mutual authentication between peers, the peer switch ports must have MKA policies and matching Connectivity Association Keys (CAKs) configured and applied. Once peers are mutually authenticated, the MKA exchange commences. From this page, you can also configure the key server priority, confidentiality-offset, and MACsec cipher suite for a MACsec policy.

To access the MKA Policy Summary page, click Security > MAC Security > MKA Policy in the navigation menu. MKA Policy Summary

Figure 383: MKA Policy Summary

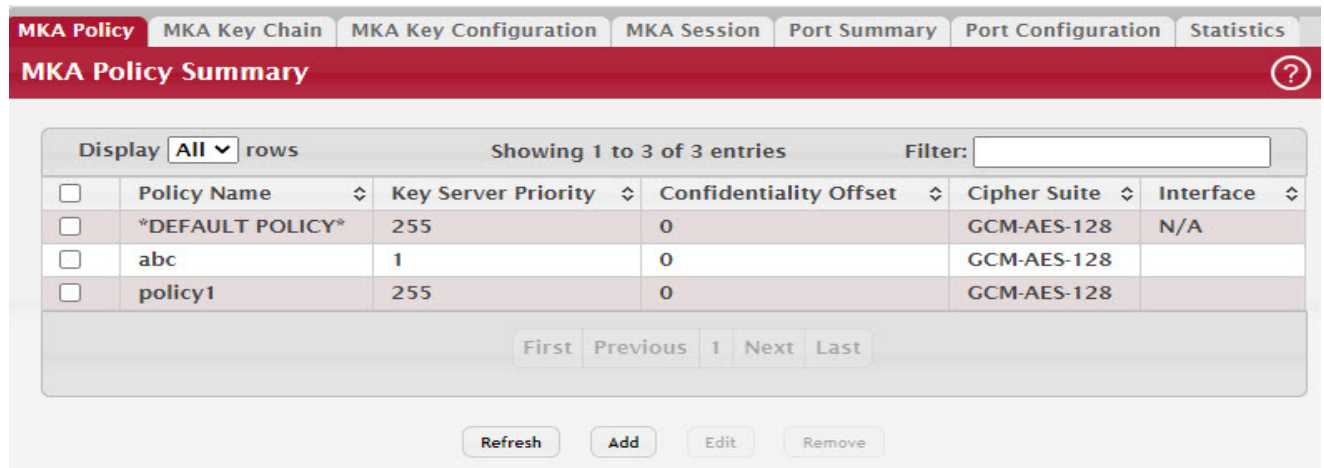


Table 367: MKA Policy Summary Fields

Field	Description
Policy Name	The MKA policy name.
Key Server Priority	The key server priority in the MKA policy. The key server is elected by the MKA protocol. The MKA peer with a higher priority is designated as the key server. If peer priorities match, the peer with the lower MAC address is elected as the key server. The key server has the responsibility to distribute Secure Association Keys (SAKs) to its peer.
Confidentiality Offset	Confidentiality offset can be 0, 30 or 50.
Cipher Suite	The cipher suite used for MACsec encryption, which can be one of the following: <ul style="list-style-type: none"> <li>GCM-AES-128</li> <li>GCM-AES-256</li> <li>GCM-AES-XPN-128</li> <li>GCM-AES-XPN-256</li> </ul>

To add a new policy, click the Add button. The Add MKA Policy window opens and allows you to create policies. Specify a policy name and other desired policy settings in the available fields.

Figure 384: Add MKA Policy

Policy Name	<input type="text"/>	(1 to 17 alphanumeric characters)
Key Server Priority	<input type="text" value="255"/>	(0 to 255)
Confidentiality Offset	<input type="text" value="None"/>	
Cipher Suite	<input type="text" value="GCM-AES-128"/>	

To edit a configured policy, select the entry to modify and click the Edit button. The Edit MKA Policy window opens and allows you to edit the selected policy.

Figure 385: Edit MKA Policy

Policy Name	<input type="text" value="p1"/>	
Key Server Priority	<input type="text" value="254"/>	(0 to 255)
Confidentiality Offset	<input type="text" value="50"/>	
Cipher Suite	<input type="text" value="GCM-AES-XPN-256"/>	

To remove one or more configured policies, select each entry in the Summary page to delete and click Remove. You must confirm the action before the entry is deleted.

Figure 386: Remove MKA Policy

Are you sure you want to remove each selected entry?

To proceed, click OK.  
To return to the web page, click Cancel.

**NOTICE**

Applied changes update the device running configuration. If you want the device to retain these changes across a reset, you must first save the configuration.

## 7.6.2 MKA Key Chain Summary

Use the MKA Key Chain Summary page to view and configure key chains. Key chains group together a set of keys configured statically on peer switches. Each key is associated with a lifetime. The CAKs in the key chain should be configured preferably with overlapping lifetimes to ensure a subsequent CAK is ready for use before the active one expires.

To access the MKA Key Chain Summary page, click Security > MAC Security > MKA Key Chain in the navigation menu.

Figure 387: MKA Key Chain Summary

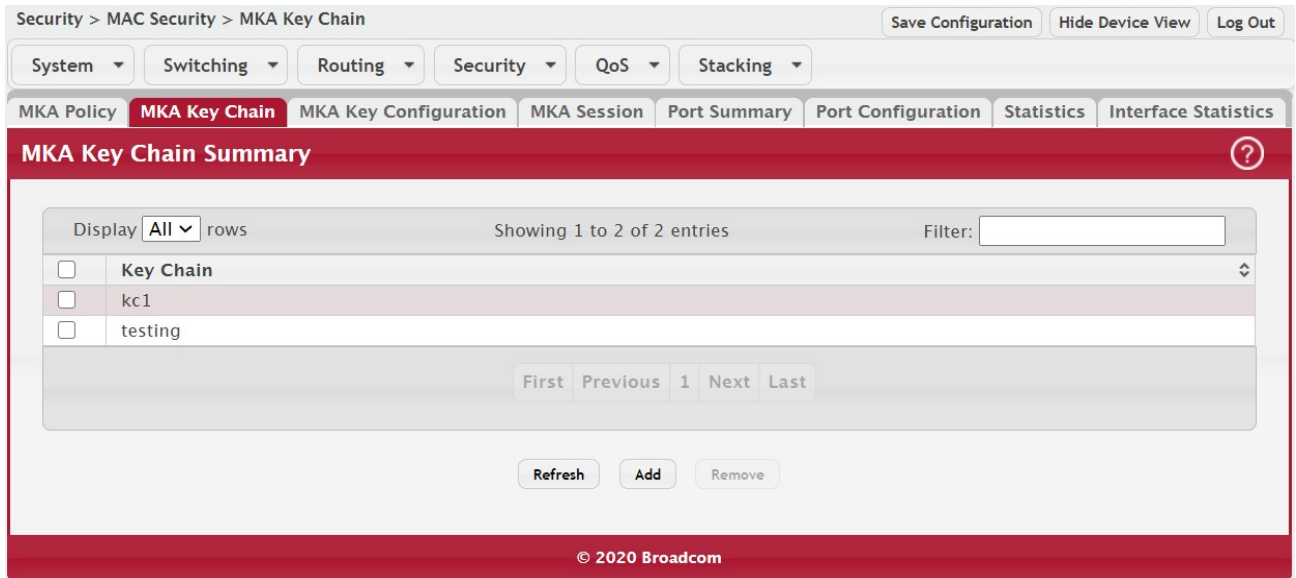
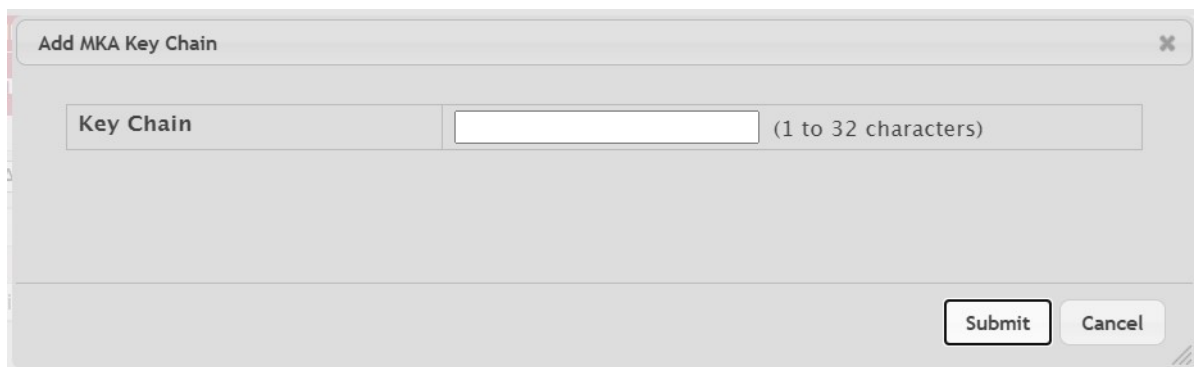


Table 368: MKA Key Chain Summary Fields

Field	Description
Key Chain	The unique key chain name.

To add a key chain, click the Add button. The Add MKA Key Chain window opens and allows you to create key chains. Specify a key chain name in the available field.

Figure 388: Add MKA Key Chain



To delete one or more configured key chains, select each entry to delete and click the Remove button. You must confirm the action before the entries are deleted.

**NOTICE**

Applied changes update the device running configuration. If you want the device to retain these changes across a reset, you must first save the configuration.

### 7.6.3 MKA Key Configuration

Use the MKA Key Configuration page to view and configure the Connectivity Association Keys (CAKs) for the selected key chain. Each CAK is identified by a unique Connectivity Association Key Name (CKN). From this page, you can also configure a key string, the supported cryptographic algorithm, and the lifetime for a CAK. The CAKs in the key chain should be configured preferably with overlapping lifetimes to ensure a subsequent CAK is ready for use before the active one expires, to avoid traffic loss. Once the peer switch ports have MKA policies and matching CAKs configured and applied, it signifies mutual authentication.

To access the MKA Key Configuration page, click Security > MAC Security > MKA Key Configuration in the navigation menu.

Figure 389: MKA Key Configuration

The screenshot shows the MKA Key Configuration page. At the top, there are tabs for 'MKA Policy', 'MKA Key Chain', 'MKA Key Configuration' (selected), 'MKA Session', 'Port Summary', 'Port Configuration', and 'Statistics'. Below the tabs, the page title 'MKA Key Configuration' is displayed. A dropdown menu for 'Key Chain' is set to 'keychain1'. Below this, there is a table with the following columns: 'Key', 'Key String', 'Cryptographic Algorithm', and 'Lifetime'. The table shows one entry with the following values: Key: 1000, Key String: 135bd758b0ee5c11c55ff6ab19fdb199, Cryptographic Algorithm: GCM-AES-128, and Lifetime: Start Time: 00:00:00 01 Jan 1970, End Time: Infinite. Below the table, there are navigation buttons: 'First', 'Previous', '1', 'Next', 'Last'. At the bottom, there are three buttons: 'Refresh', 'Add', and 'Remove'.

Table 369: MKA Key Configuration Fields

Field	Description
Key	A unique hexadecimal character string that identifies a key in the key chain. The key can be up to 32 characters in length when 128-bit encryption is specified, and up to 64 characters in length when 256-bit encryption is specified. However, the key identifier must contain an even number of characters.
Key String	A unique hexadecimal character string associated with the key in the key chain. The key string can be up to 32 characters in length when 128-bit encryption is specified, and up to 64 characters in length when 256-bit encryption is specified.
Cryptographic Algorithm	The type of encryption to be used with the key, which can be one of the following: <ul style="list-style-type: none"> <li>GCM-AES-128 – 128-bit encryption</li> <li>GCM-AES-256 – 256-bit encryption</li> </ul>
Lifetime	Indicates the key lifetime period. Once the CAK lifetime is over, the connectivity association between the peers is torn down by the key server.

To add a new key for a key chain, click the Add button. The Add MKA Key window opens. Specify a key name and other



desired key settings in the available fields.

Figure 390: Add MKA Key

Key Chain	kc1
Key	<input type="text"/> (1 to 32 hex)
Key String	<input type="text"/> (0 to 32 hex)
Cryptographic Algorithm	<input checked="" type="radio"/> GCM-AES-128 <input type="radio"/> GCM-AES-256
Start Time	<input type="text"/> (00:00:00 to 23:59:59)
Start Date	<input type="text"/>
Duration	<input checked="" type="checkbox"/> Infinite <input type="text"/> (0 to 864000)
End Time	<input type="text"/> (00:00:00 to 23:59:59)
End Date	<input type="text"/>

Table 370: Add MKA Key Fields

Field	Description
Key Chain	The unique key chain name.
Key	A unique hexadecimal character string that identifies a key in the key chain. The key can be up to 32 characters in length when 128-bit encryption is specified, and up to 64 characters in length when 256-bit encryption is specified. However, the key identifier must contain an even number of characters.
Key String	A unique hexadecimal character string associated with the key in the key chain. The key string can be up to 32 characters in length when 128-bit encryption is specified, and up to 64 characters in length when 256-bit encryption is specified.
Cryptographic Algorithm	The type of encryption to be used with the key, which can be one of the following: <ul style="list-style-type: none"> <li>GCM-AES-128 – 128-bit encryption</li> <li>GCM-AES-256 – 256-bit encryption</li> </ul>
Start Time	The start time for the key in hh:mm:ss format.
Start Date	The start date for the key.
Duration	The time in seconds from the key start time until the key expires. If the duration is configured, the key is valid from the configured start time for the duration specified. The default duration is infinite.
End Time	The end time for the key in hh:mm:ss format.
End Date	The end date for the key.

To remove one or more configured keys from a key chain, select each entry to delete and click the Remove button. You must confirm the action before the entries are deleted.

**NOTICE**

Applied changes update the device running configuration. If you want the device to retain these changes across a reset, you must first save the configuration.

### 7.6.4 MKA Session Summary

Use the MKA Session Summary page to view the active MKA sessions and session details. To access the MKA Session Summary page, click Security > MAC Security > MKA Session in the navigation menu.

Figure 391: MKA Session Summary

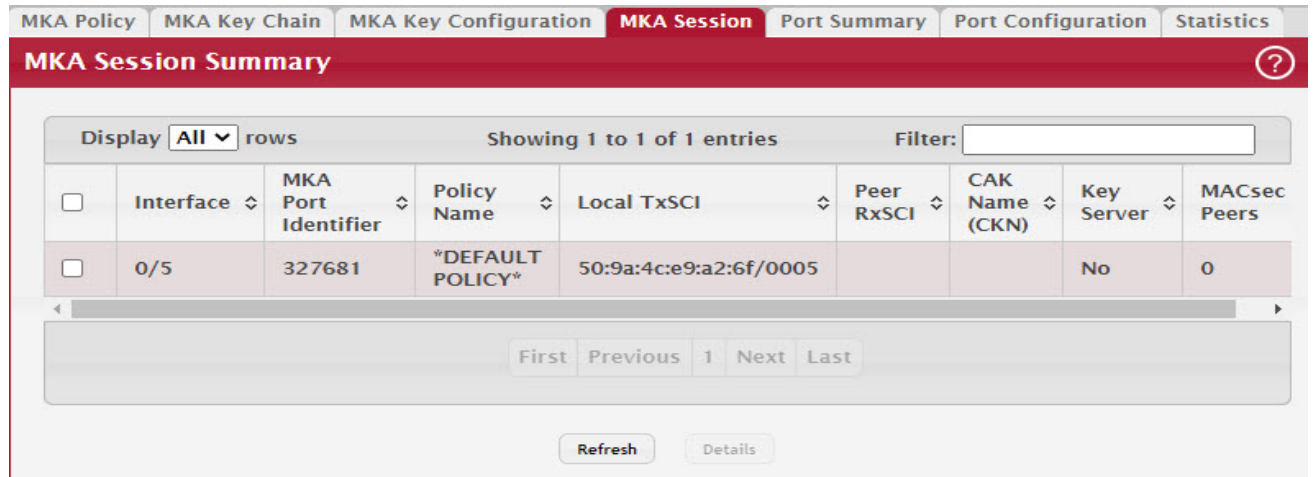


Table 371: MKA Session Summary Fields

Field	Description
Interface	The physical interface on which the MKA session is active.
MKA Port Identifier	The logical port identifier.
Policy Name	The name of the MKA policy applied on the interface.
Local TxSCI	The MAC address of the physical interface concatenated with the 16-bit port ID.
Peer RxSCI	The MAC address of the peer interface concatenated with the peer 16-bit port ID.
CAK Name (CKN)	The connectivity association key name.
Key Server	Indicates the key server status. If the MKA session is the key server, the status is Yes, otherwise the status is No.
MACsec Peers	The number of live peers.

To view session details, select an interface and click the Details button. The MKA Session Details window opens with the details of the selected session.

Figure 392: MKA Session Details

MKA Session Details	
Interface	0/5
Status	UnSecured
Local TxSCI	50:9a:4c:e9:a2:6f/0005
MKA Port Identifier	327681
CAK Name (CKN)	
Key Server	No
Interface MAC Address	50:9a:4c:e9:a2:6f
Member Identifier (MI)	8e018d12fb6c760b329e5394
Message Number (MN)	
Replay Protection	No
Replay Window Size	0
Algorithm Agility	0080c201
MACsec Desired	Yes
MKA Policy	*DEFAULT POLICY*
Cipher Suite	GCM-AES-128
MACsec Capability	3 (Integrity, Confidentiality Offset of 0,30,50)
Key Server Priority	255
Confidentiality Offset	0
Latest SAK Status	No Tx & No Rx
Latest SAK AN	0
Latest SAK KI (KN)	
# of MACsec Capable Live Peers	0
# of MACsec Capable Potential Peers	0

Close

Table 372: MKA Session Details Fields

Field	Description
Interface	The physical interface on which the MKA session is active.
Status	The secured status of the MKA session.
Local TxSCI	The MAC address of the physical interface concatenated with the 16-bit port ID.
MKA Port Identifier	The logical port identifier.
CAK Name (CKN)	The connectivity association key name.
Key Server	Indicates the key server status. If the MKA session is the key server, the status is Yes, otherwise the status is No.

**Table 372: MKA Session Details Fields (Continued)**

Field	Description
Interface MAC Address	The MAC address of the local interface.
Member Identifier (MI)	The local MKA participant identifier.
Message Number (MN)	The local MKA participant message number.
Replay Protection	Indicates the replay protection configuration status.
Replay Window Size	The configured replay protection window size.
Algorithm Agility	The algorithm agility parameter for the CA.
MACsec Desired	The MACsec desired parameter for the CA.
MKA Policy	The name of the MKA policy applied on the interface.
Cipher Suite	The operational cipher suite for the CA.
MACsec Capability	The MACsec capability for the CA.
Key Server Priority	The local MKA instance key server priority.
Confidentiality Offset	The configured confidentiality offset.
Latest SAK Status	Indicates whether the current SAK is used for Tx and Rx.
Latest SAK AN	The current SAK association number.
Latest SAK KI (KN)	The current SAK key identifier.
# of MACsec Capable Live Peers	The number of live MKA peers.
# of MACsec Capable Potential Peers	The number of potential MKA peers.

**Figure 393: MKA Session Details – Live Peer List, Potential Peer List**

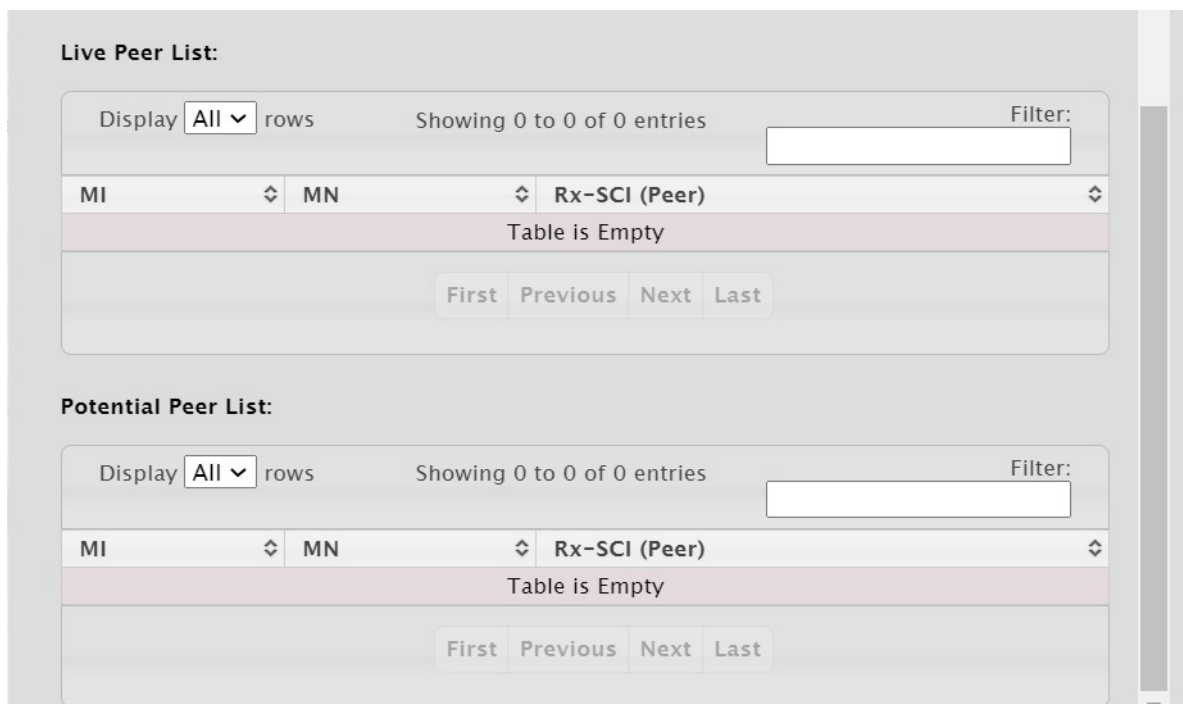


Table 373: Live Peer List and Potential Peer List Fields

Field	Description
MI	MKA participant identifier.
MN	MKA participant message number.
Rx-SCI (Peer)	MAC address of the peer interface concatenated with the peer 16-bit port ID.

### 7.6.5 MACsec Port Summary

Use the MACsec Port Summary page to view port-specific data. To access the MACsec Port Summary page, click Security > MAC Security > Port Summary in the navigation menu.

Figure 394: MACsec Port Summary

Table 374: MACsec Port Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Transmit SCI	The transmit secure channel identifier.
Transmit SCI Transmitting	Indicates the transmitting status of the transmit secure channel.
Transmit SA Next PN	The next packet number of the transmit security association.
Receive SCI	The receive secure channel identifier.
Receive SC Receiving	Indicates the receiving status of the receive secure channel.
Receive SA Next PN	The next packet number of the receive security association.
Receive SA AN	The association number of the receive security association.

### 7.6.6 MACsec Port Configuration

Use the MACsec Port Configuration page to view and configure interfaces for MACsec. MACsec is supported only on physical interfaces. Once the peer switch ports have MKA policies and matching CAKs configured and applied, it signifies mutual authentication.

To access the MACsec Port Configuration page, click Security > MAC Security > Port Configuration in the navigation menu.

Figure 395: MACsec Port Configuration

Interface	MACsec Mode	MKA Policy	MACsec PSK Keychain	Replay Protection	Replay Protection Window	Cipher(s) Supported
0/1	Disabled	*DEFAULT POLICY*		Disabled	0	GCM-AES-128
0/2	Disabled	*DEFAULT POLICY*		Disabled	0	GCM-AES-128
0/3	Disabled	*DEFAULT POLICY*		Disabled	0	GCM-AES-128
0/4	Disabled	*DEFAULT POLICY*		Disabled	0	GCM-AES-128
0/5	Switch to Switch	*DEFAULT POLICY*	keychain1	Disabled	0	GCM-AES-128

Table 375: MACsec Port Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When editing information for one or more interfaces, this field identifies the interfaces that are being configured.
MACsec Mode	The MACsec mode on the interface, which can be one of the following: <ul style="list-style-type: none"> <li>Disable: MACsec is disabled on the interface.</li> <li>Switch to Switch: MACsec operates in switch to switch mode on the interface using PSK CAKs configured statically on the peer switches.</li> <li>Host to Switch: MACsec operates in host to switch mode on the interface via a secure link between a MACsec capable host and the switch.</li> </ul>
MKA Policy	The MKA policy applied on the interface for peer authentication.
MACsec PSK Keychain	The key chain applied on the interface for peer authentication. Key chain can be applied only on interfaces in the switch to switch mode.
Replay Protection	Indicates the replay protection status on the interface as enabled or disabled.
Replay Protection Window	The number of out of sequence frames that are accepted on the interface. Window size "0" indicates all frames that arrive out of sequence are dropped.
Cipher(s) Supported	The supported cipher suites on the interface, which can be one of the following: <ul style="list-style-type: none"> <li>GCM-AES-128</li> <li>GCM-AES-256</li> <li>GCM-AES-XPB-128</li> <li>GCM-AES-XPB-256</li> </ul>

**NOTICE**

Applied changes update the device running configuration. If you want the device to retain these changes across a reset, you must first save the configuration.

To configure a port, select a port and click the Edit button. The Edit MACsec Port Configuration window opens. Specify a key name and other desired key settings in the available fields.

Figure 396: Edit MACsec Port Configuration

Interface	0/1
MACsec Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Switch to Switch <input type="radio"/> Host to Switch
MKA Policy	*DEFAULT POLICY* ▼
MACsec PSK Keychain	▼
Replay Protection	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Replay Protection Window	0 (0 to 4294967295)

Submit Cancel

Table 376: Edit MACsec Port Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When editing information for one or more interfaces, this field identifies the interfaces that are being configured.
MACsec Mode	The MACsec mode on the interface, which can be one of the following: <ul style="list-style-type: none"> <li>• Disable: MACsec is disabled on the interface.</li> <li>• Switch to Switch: MACsec operates in switch to switch mode on the interface using PSK CAKs configured statically on the peer switches.</li> <li>• Host to Switch: MACsec operates in host to switch mode on the interface via a secure link between a MACsec capable host and the switch.</li> </ul>
MKA Policy	The MKA policy applied on the interface for peer authentication.
MACsec PSK Keychain	The key chain applied on the interface for peer authentication. Key chain can be applied only on interfaces in the switch to switch mode.
Replay Protection	Disable or enable replay protection.
Replay Protection Window	The number of out of sequence frames that are accepted on the interface. Window size "0" indicates all frames that arrive out of sequence are dropped.

### 7.6.7 MKA Global Statistics

Use the MKA Global Statistics page to view the MKA protocol global statistical information. To access the MKA Global Statistics page, click Security > MAC Security > Statistics in the navigation menu.

Figure 397: MKA Global Statistics

MKA Global Statistics	
<b>MKA Session Totals</b>	
Secured	0
Deleted (Secured)	0
<b>SA Statistics</b>	
SAKs Generated	0
SAKs Rekeyed	0
SAKs Received	0
SAKs Responses Received	0
<b>MKPDU Statistics</b>	
MKPDUs Validated & Rx	0
MKPDUs Transmitted	13461
Distributed SAK	0

Figure 398: MKA Global Statistics Part 2

<b>SAK Failures</b>	
SAK Generation	0
SAK Encryption	0
SAK Decryption	0
<b>CA Failures</b>	
ICK Derivation	0
KEK Derivation	0
Invalid Peer MACsec Capability	0
<b>MACsec Failures</b>	
Rx SC Creation	0
Tx SC Creation	0
Rx SA Installation	0
Tx SA Installation	0
<b>MKPDU Failures</b>	
MKPDU Tx	0
MKPDU Rx Validation	0
MKPDU Rx Bad Peer MN	0



Table 377: MKA Global Statistics Fields

Field	Description
<b>MKA Session Tools</b>	
Secured	The number of MKA sessions secured.
Deleted (Secured)	The number of MKA sessions deleted.
<b>SA Statistics - The Secure Association statistical information.</b>	
SAKs Generated	The number of Secure Association Keys (SAKs) generated.
SAKs Rekeyed	The number of SAKs refreshed or re-keyed.
SAKs Received	The number of SAKs received.
SAKs Responses Received	The number of SAKs received from the key server.
<b>MKPDU Statistics</b>	
MKPDUs Validated and Rx	The number of valid MKPDUs received.
MKPDUs Transmitted	The number of MKPDUs transmitted.
Distributed SAK	The number of SAKs distributed.
<b>SAK Failures</b>	
SAK Generation	The number of SAK generation failures.
SAK Encryption	The number of SAK encryption failures.
SAK Decryption	The number of SAK decryption failures.
<b>CA Failures - The Connectivity Association (CA) statistical information.</b>	
ICK Derivation	The number of Integrity Check Value Key (ICK) derivation failures.
KEK Derivation	The number of Key Encryption Key (KEK) derivation failures.
Invalid Peer MACsec Capability	The number of invalid peer MACsec capability.
<b>MACsec Failures</b>	
Rx SC Creation	The number of Rx Secure Channel (SC) creation failures.
Tx SC Creation	The number of Tx SC creation failures.
Rx SA Installation	The number of Rx SC installation failures.
Tx SA Installation	The number of Tx SC installation failures.
<b>MKPDU Failures</b>	
MKPDU Tx	The number of MKPDU Tx failures.
MKPDU Rx Validation	The number of MKPDU Rx validation failures.
MKPDU Rx Bad Peer MN	The number of MKPDU Rx bad peer message number.

Click the Clear Counters button to clear all of the statistics displayed on this page by resetting them to their default values.

## 7.6.8 MKA Interface Statistics

Use the MKA Interface Statistics page to view the MKA statistical information for an interface. To access the MKA Interface Statistics page, click Security > MAC Security > Interface Statistics in the navigation menu.

Figure 399: MKA Interface Statistics

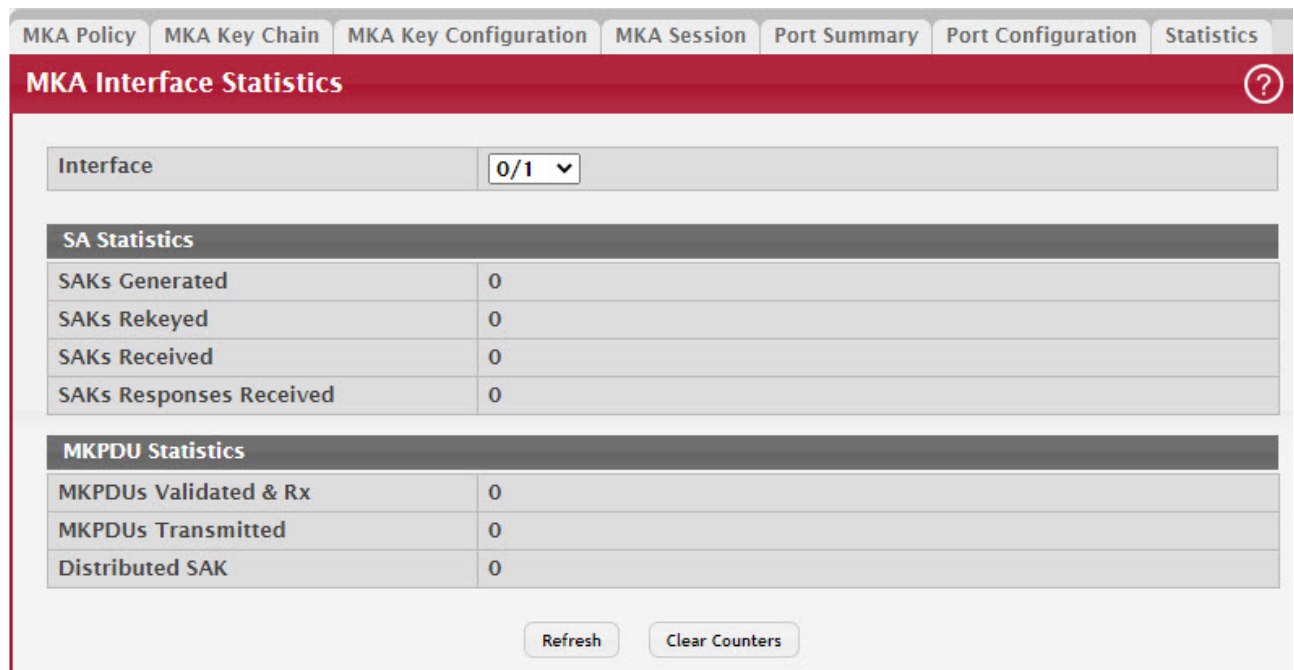


Table 378: MKA Interface Statistics Fields

Field	Description
Interface	Select the interface to view the MKA statistics.
<b>SA Statistics: The Secure Association statistical information.</b>	
SAKs Generated	The number of Secure Association Keys (SAKs) generated.
SAKs Rekeyed	The number of SAKs refreshed or rekeyed.
SAKs Received	The number of SAKs received.
SAKs Responses Received	The number of SAKs received from the key server.
<b>MKPDU Statistics</b>	
MKPDUs Validated & Rx	The number of valid MKPDUs received.
MKPDUs Transmitted	The number of MKPDUs transmitted.
Distributed SAK	The number of SAKs distributed.

To clear the interface statistics and reset them to their default values, select the interface from the Interface list, and click the Clear Counters button.

## 8/ Configuring Quality of Service

This section gives an overview of Quality of Service (QoS) and explains the QoS features available from the Quality of Service navigation menu.

In a typical switch, each physical port consists of one or more queues for transmitting packets on the attached network. Multiple queues per port are often provided to give preference to certain packets over others based on user-defined criteria. When a packet is queued for transmission in a port, the rate at which it is serviced depends on how the queue is configured and possibly the amount of traffic present in the other queues of the port. If a delay is necessary, packets get held in the queue until the scheduler authorizes the queue for transmission. As queues become full, packets have no place to be held for transmission and get dropped by the switch.

QoS is a means of providing consistent, predictable data delivery by distinguishing between packets that have strict timing requirements from those that are more tolerant of delay. Packets with strict timing requirements are given special treatment in a QoS capable network. With this in mind, all elements of the network must be QoS-capable. The presence of at least one node which is not QoS-capable creates a deficiency in the network path and the performance of the entire packet flow is compromised.

---

### NOTICE

Some of the features described in this section may not be supported in FASTPATH software releases for particular hardware platforms.

---

## 8.1 Configuring Access Control Lists

Access Control Lists (ACLs) ensure that only authorized users have access to specific resources while blocking off any unwarranted attempts to reach network resources. ACLs provide traffic flow control, restrict contents of routing updates, decide which types of traffic are forwarded or blocked, and provide security for the network. FASTPATH software supports IPv4, IPv6, and MAC ACLs. The total number of MAC and IP ACLs supported by FASTPATH software is platform specific.

You first create an IPv4-based, IPv6-based, or MAC-based rule and assign a unique ACL ID. Then, you define the rules, which can identify protocols, source and destination IP and MAC addresses, and other packet-matching criteria. Finally, you use the ID number to assign the ACL to a port or to a VLAN interface.

### 8.1.1 IP Access Control Lists

IP access control lists (ACL) allow network managers to define classification actions and rules for specific ports. ACLs are composed of access control entries (ACE), or rules, that consist of the filters that determine traffic classifications. The total number of rules that can be defined for each ACL is platform-specific. These rules are matched sequentially against a packet. When a packet meets the match criteria of a rule, the specified rule action (Permit/Deny) is taken, including dropping the packet or disabling the port, and the additional rules are not checked for a match. For example, a network administrator defines an ACL rule that says port number 20 can receive TCP packets. However, if a UDP packet is received the packet is dropped.

The IP Access Control List folder contains links to web pages that allow you to configure and view IP ACLs.

To configure an IP ACL:

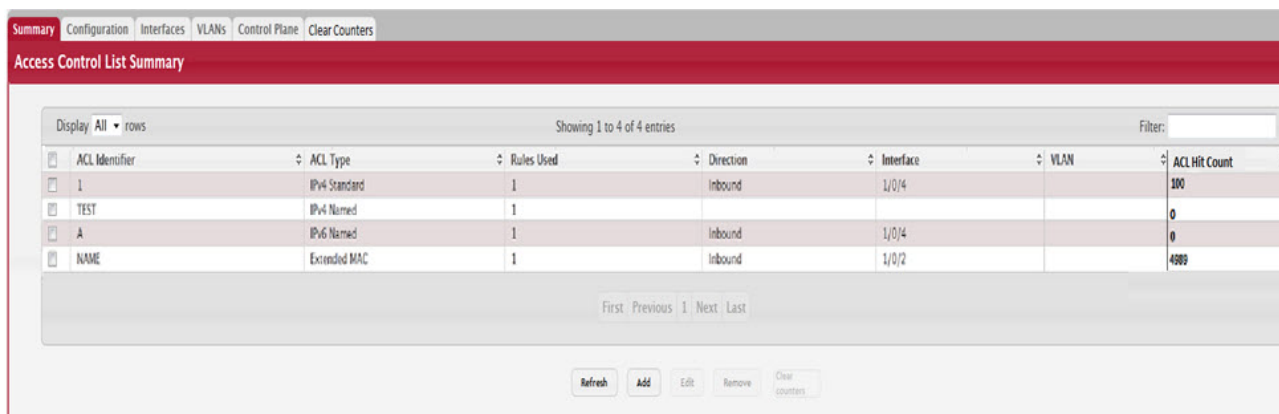
1. Use the [Section 8.1.1.1: "IP ACL Configuration"](#) page to define the IP ACL type and assign an ID to it.
2. Use the [Section 8.1.1.3: "Access Control List Interface Summary"](#) page to create rules for the ACL.
3. Use the [Section 8.1.1.2: "Access Control List Configuration"](#) page to view the configuration.

#### 8.1.1.1 IP ACL Configuration

Use the IP ACL Configuration page to add or remove IP-based ACLs and to enable or disable the ACL counters. On this menu the interfaces to which an IP ACL applies must be specified, as well as whether it applies to inbound or outbound traffic. Rules for the IP ACL are specified/created using the [Section 8.1.1.3: "Access Control List Interface Summary"](#) page.

To display the Access List Summary page, click QoS > Access Control Lists > Summary in the navigation menu.

Figure 400: Access Control List Summary



Use the buttons at the bottom of the page to perform the following tasks:

- To add an ACL, click Add and configure the ACL type and ID.
- To remove one or more configured ACLs, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.
- To configure rules for an ACL, select the ACL to configure and click Edit. You are redirected to the Access Control List Configuration page for the selected ACL.

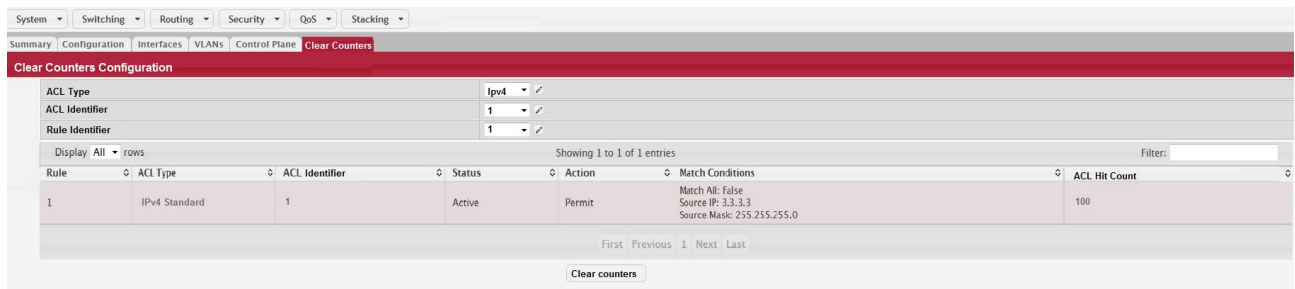
Table 379: Access List Summary Fields

Field	Description
ACL Counters	The administrative status of the ACL counters. This field controls the status of the counters for all ACL types.
ACL Identifier	The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4 and MAC ACLs use alphanumeric characters. Dynamic ACLs are identified with an additional #d appended to the ACL Identifier. The Summary page displays the 255-character length ACL names for named IPv4, IPv6, and MAC ACLs.
ACL Type	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows: <ul style="list-style-type: none"> <li>• IPv4 Standard – Match criteria is based on the source address of IPv4 packets.</li> <li>• IPv4 Extended – Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets.</li> <li>• IPv4 Named – Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number.</li> <li>• IPv6 Named – Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets.</li> <li>• Extended MAC – Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.</li> </ul>
Rules Used	The number of rules currently configured for the ACL.
Direction	Indicates whether the packet is checked against the rules in an ACL when it is received on an interface (Inbound) or after it has been received, routed, and is ready to exit an interface (Outbound).
Interface	The interface(s) to which the ACL has been applied.
VLAN	Each VLAN to which the ACL has been applied.
ACL Hit Count	Displays the ACL Rule hit count value. Click the Clear Counters button to clear the hit count value.

### 8.1.1.1.1 Clear Counters Configuration

Click the Clear Counters button to clear the hit count value.

Figure 401: Clear Counters Configuration

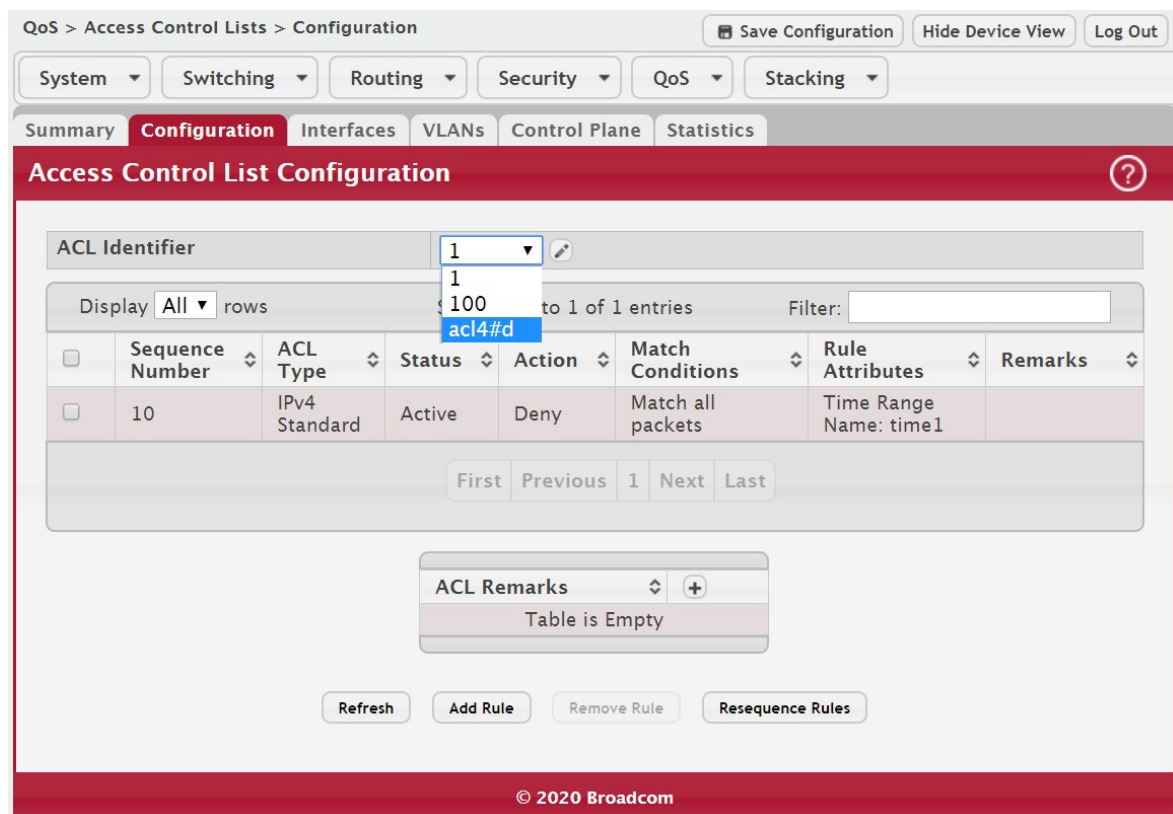


### 8.1.1.2 Access Control List Configuration

Use this page to configure rules for the existing Access Control Lists (ACLs) on the system and to view summary information about the rules that have been added to an ACL. Each ACL rule is configured to match one or more aspects of traffic on the network. When a packet matches the conditions in a rule, it is handled according to the configured action (permit or deny) and attributes. Each ACL can have multiple rules, but the final rule for every ACL is an implicit deny all rule. For each rule, a packet must match all the specified criteria in order for the specified rule action (Permit/Deny) to take place.

To display the Access Control List Configuration page, click QoS > Access Control Lists > Configuration in the navigation menu.

Figure 402: Access Control List Configuration



Use the buttons to perform the following tasks:

- To add an Access List Rule entry, select the ID of the ACL that will include the rule from the ACL Identifier menu. Then, click Add Rule and configure the rule criteria and attributes. New rules cannot be created if the maximum number of rules has been reached.

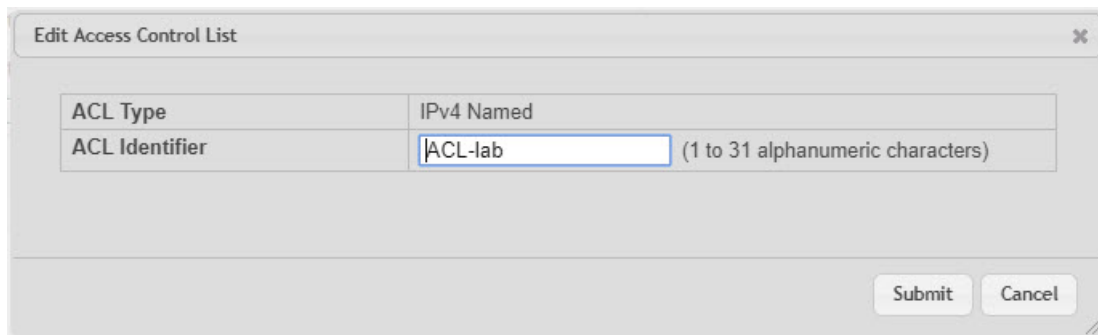
- To remove the most recently configured rule for an ACL, select the ID of the appropriate ACL from the ACL Identifier menu and click Remove Last Rule. You must confirm the action before the entry is deleted.
- To resequence rules for an ACL, select the ID of the appropriate ACL from the ACL Identifier menu and click Resequence Rules.

**Table 380: ACL Identifier Field**

Field	Description
ACL Identifier	The menu contains the ID for each ACL that exists on the system. Dynamic ACLs have an additional tag #d appended to the ACL Identifier in the ACL Identifier list. Before you add or remove a rule, you must select the ID of the ACL from the menu.

For ACLs with alphanumeric names, click the Edit icon to change the ACL ID. The Edit Access Control List dialog box is displayed. The ID of a Named IPv4 ACL must begin with a letter, and not a number. The ACL identifier for IPv4 standard and IPv4 extended ACLs cannot be changed.

**Figure 403: Edit Access Control List**



The Edit Access Control List page prevents the administrator from renaming a named IPv4, IPv6, or MAC ACL beginning with the case-insensitive names reserved for Dynamic ACLs, for example, **IP-DACL-IN-** and **IPV6-DACL-IN-**.

The user can perform the resequence action on the Edit Access Control List page.

**Figure 404: Edit Access Control List - Dynamic ACLs**

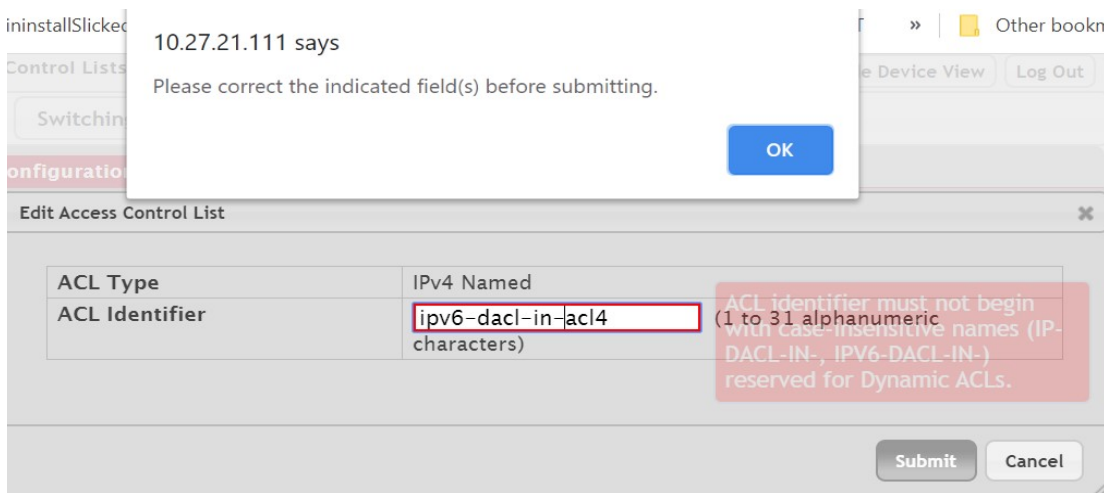


Table 381: Edit Access Control List Configuration Fields

Field	Description
Sequence Number	The number that indicates the position of a rule within the ACL. If the sequence number is not specified during rule creation, the rule is automatically assigned a sequence number after it is successfully added to the ACL. The rules are displayed based on their position within the ACL, but can also be renumbered. Packets are checked against the rule criteria in order, from the lowest-numbered rule to the highest. When the packet matches the criteria in a rule, it is handled according to the rule action and attributes. If no rule matches a packet, the packet is discarded based on the implicit deny all rule, which is the final rule in every ACL.
ACL Type	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows: <ul style="list-style-type: none"> <li>• IPv4 Standard – Match criteria is based on the source address of IPv4 packets.</li> <li>• IPv4 Extended – Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets.</li> <li>• IPv4 Named – Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number.</li> <li>• IPv6 Named – Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets.</li> <li>• Extended MAC – Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.</li> </ul>
Status	Indicates whether the ACL is active. If the ACL is a time-based ACL that includes a time range, the ACL is active only during the periods specified within the time range. If an ACL does not include a time range, the status is always active.
Action	The action to take when a packet or frame matches the criteria in the rule: <ul style="list-style-type: none"> <li>• Permit – The packet or frame is forwarded.</li> <li>• Deny – The packet or frame is dropped.</li> </ul> <p><b>Note:</b> When configuring ACL rules in the Add Access Control List Rule window, the selected action determines which fields can be configured. Not all fields are available for both Permit and Deny actions.</p>
Match Conditions	The criteria used to determine whether a packet or frame matches the ACL rule.
Rule Attributes	Each action (beyond the basic Permit and Deny actions) to perform on the traffic that matches the rule.
Remarks	One or more remarks configured for the selected ACL and associated with the rule during rule creation. To delete a remark associated with the rule, click the – (minus) button preceding remark. You must confirm the action before the rule associated remark is removed.

Use the buttons available in the ACL Remarks table to perform the following tasks:

- To add a remark, click the + (plus) button and enter the remark to add.
- To delete a remark from the list, click the – (minus) button associated with the entry to remove. You must confirm the action before the entry is removed.

Table 382: ACL Remarks Fields

Field	Description
ACL Remarks	Lists the configured remarks for the selected ACL. All remarks present in this table are applied to the next rule created with the Add Rule button.

After you click Add Rule, the Add Access Control List Rule window opens and allows you to add a rule to the ACL that was selected from the ACL Identifier field. The fields available in the window depend on the ACL Type. The following informa-

tion describes the fields in this window. The Match Criteria tables that apply to IPv4 ACLs, IPv6 ACLs, and MAC ACLs are described separately.

Figure 405: Add IPv4 ACL Rule

Add IPv4 ACL Rule
✕

Sequence Number	<input type="text"/> (1 to 2147483647)
Action	<input type="radio"/> Permit <input checked="" type="radio"/> Deny
<b>Match Criteria</b>	
Every	<input type="checkbox"/>
Protocol	<input type="text"/> (0 to 255, or keyword) ?
Fragments	<input type="checkbox"/>
Source IP Address / Wildcard Mask	<input type="text"/> / <input type="text"/> (x.x.x.x)
Source L4 Port	<input checked="" type="radio"/> Equal <input type="radio"/> Not Equal <input type="radio"/> Less Than <input type="radio"/> Greater Than <input type="radio"/> Range <input type="text"/> - <input type="text"/> (0 to 65535, or keyword) ?
Destination IP Address / Wildcard Mask	<input type="text"/> / <input type="text"/> (x.x.x.x)
Destination L4 Port	<input checked="" type="radio"/> Equal <input type="radio"/> Not Equal <input type="radio"/> Less Than <input type="radio"/> Greater Than <input type="radio"/> Range <input type="text"/> - <input type="text"/> (0 to 65535, or keyword) ?
TTL Field Value	<input type="text"/> (0 to 255)
IGMP Type	<input type="text"/> (0 to 255)
ICMP Type	<input type="text"/> (0 to 255)
ICMP Code	<input type="text"/> (0 to 255)
ICMP Message	<input type="text"/>
TCP Flags	<div style="border: 1px solid #ccc; padding: 2px;">           +FIN            -FIN            +SYN            -SYN            +RST            -RST            +PSH            -PSH         </div>
Service Type	<input type="checkbox"/>
IP DSCP	<input type="checkbox"/> <input type="text"/> (0 to 63, or keyword) ?
IP Precedence	<input type="checkbox"/> <input type="text"/> (0 to 7)
IP TOS Bits / Wildcard Mask	<input type="checkbox"/> <input type="text"/> / <input type="text"/> (0 to FF hex)
<b>Rule Attributes</b>	
Assign Queue	<input type="text"/> (0 to 7)
Interface	<input type="text"/> <input type="radio"/> Redirect <input type="radio"/> Mirror
Log	<input type="checkbox"/>
Redirect External Agent	<input type="text"/> (1 to 100)
Time Range Name	<input type="text"/> (1 to 31 alphanumeric characters)
Committed Rate / Burst Size	<input type="text"/> (1 to 4294967295) <input type="text"/> (1 to 128)



Table 383: Add IPv4 ACL Rule Fields

Field	Description
Sequence Number	The user can apply a sequence number for the rule. The sequence number is removed with the removing of the rule. The sequence number can be changed as a result of the resequence action using the Edit Access Control List page.
Action	Select the option to permit or deny associating the sequence number to the ACL rule.
Match Criteria (IPv4 ACLs)	The fields in this section specify the criteria to use to determine whether an IP packet matches the rule. The fields described below apply to IPv4 Standard, IPv4 Extended, and IPv4 Named ACLs unless otherwise noted.
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria, so if Every is selected, no other match criteria can be configured. To configure specific match criteria, this option must be clear.
Protocol	(IPv4 Extended and IPv4 Named ACLs) The IANA-assigned protocol number to match within the IP packet. You can also specify one of the following keywords: EIGRP, GRE, ICMP, IGMP, IP, IPINIP, OSPF, PIM, TCP, or UDP.
Fragments	(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on fragmented IP packets.
Source IP Address / Wildcard Mask	The source port IP address in the packet and source IP wildcard mask (in the second field) to compare to the IP address in a packet header. Wild card masks determines which bits in the IP address are used and which bits are ignored. A wild card mask of 255.255.255.255 indicates that no bit is important. A wildcard of 0.0.0.0 indicates that all of the bits are important. Wildcard masking for ACLs operates differently from a subnet mask. A wildcard mask is in essence the inverse of a subnet mask. With a subnet mask, the mask has ones (1's) in the bit positions that are used for the network address, and has zeros (0's) for the bit positions that are not used. In contrast, a wildcard mask has (0's) in a bit position that must be checked. A '1' in a bit position of the ACL mask indicates the corresponding bit can be ignored. This field is required when you configure a source IP address.
Source L4 Port	(IPv4 Extended and IPv4 Named ACLs) The TCP/UDP source port to match in the packet header. The Source L4 Port and Destination L4 port are configurable only if protocol is either TCP or UDP. Equal to, Not Equal to, Greater than, and Less than options are available. For TCP protocol: BGP, Domain, Echo, FTP, FTP-Data, HTTP, SMTP, Telnet, WWW, POP2, or POP3 For UDP protocol: Domain, Echo, NTP, RIP, SNMP, TFTP, Time, or WHO
Destination IP Address / Wildcard Mask	The destination port IP address in the packet and destination IP wildcard mask (in the second field) to compare to the IP address in a packet header. Wild card masks determines which bits in the IP address are used and which bits are ignored. A wild card mask of 255.255.255.255 indicates that no bit is important. A wildcard of 0.0.0.0 indicates that all of the bits are important. Wildcard masking for ACLs operates differently from a subnet mask. A wildcard mask is in essence the inverse of a subnet mask. With a subnet mask, the mask has ones (1's) in the bit positions that are used for the network address, and has zeros (0's) for the bit positions that are not used. In contrast, a wildcard mask has (0's) in a bit position that must be checked. A 1 in a bit position of the ACL mask indicates the corresponding bit can be ignored. This field is required when you configure a destination IP address.
Destination L4 Port	(IPv4 Extended and IPv4 Named ACLs) The TCP/UDP destination port to match in the packet header. The Source L4 Port and Destination L4 port are configurable only if protocol is either TCP or UDP. Equal to, Not Equal to, Greater than, and Less than options are available. For TCP protocol: BGP, Domain, Echo, FTP, FTP-Data, HTTP, SMTP, Telnet, WWW, POP2, or POP3 For UDP protocol: Domain, Echo, NTP, RIP, SNMP, TFTP, Time, or WHO
TTL Field Value	(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the specified Time-to-Live (TTL) field value.
IGMP Type	(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the specified IGMP message type. This option is available only if the protocol is IGMP.
ICMP Type	(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the specified ICMP message type. This option is available only if the protocol is ICMP.

Table 383: Add IPv4 ACL Rule Fields (Continued)

Field	Description
ICMP Code	(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the specified ICMP message code. This option is available only if the protocol is ICMP.
ICMP Message	(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the ICMP message type and code. Specify one of the following supported ICMP messages: Echo, Echo-Reply, Host-Redirect, Mobile-Redirect, Net-Redirect, Net-Unreachable, Redirect, Packet-Too-Big, Port-Unreachable, Source-Quench, Router-Solicitation, Router-Advertisement, Time-Exceeded, TTL-Exceeded, and Unreachable. This option is available only if the protocol is ICMP.
TCP Flags	(IPv4 Extended and IPv4 Named ACLs) IP ACL rule to match on the TCP flags. When a + flag is specified, a match occurs if the flag is set in the TCP header. When a - flag is specified, a match occurs if the flag is not set in the TCP header. When Established is specified, a match occurs if either RST or ACK bits are set in the TCP header. This option is available only if the protocol is TCP.
Service Type	(IPv4 Extended and IPv4 Named ACLs) The service type to match in the IP header. The options in this menu are alternative ways of specifying a match condition for the same Service Type field in the IP header, but each service type uses a different user notation. After you select the service type, specify the value for the service type in the appropriate field. Only the field associated with the selected service type can be configured. The service types are as follows: <ul style="list-style-type: none"> <li>IP DSCP – Matches the packet IP DiffServ Code Point (DSCP) value to the rule. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header.</li> <li>IP Precedence – Matches the IP Precedence value to the rule. The IP Precedence field in a packet is defined as the high-order three bits of the Service Type octet in the IP header.</li> <li>IP TOS Bits – Matches on the Type of Service (TOS) bits in the IP header. The IP TOS field in a packet is defined as all eight bits of the Service Type octet in the IP header. For example, to check for an IP TOS value having bits 7 and 5 set and bit 1 clear, where bit 7 is most significant, use a TOS Bits value of 0xA0 and a TOS Mask of 0xFF. <ul style="list-style-type: none"> <li>TOS Bits – Requires the bits in a packet's TOS field to match the two-digit hexadecimal number entered in this field.</li> <li>TOS Mask – The bit positions that are used for comparison against the IP TOS field in a packet.</li> </ul> </li> </ul>
Time Range Name	The name of the time range that will impose a time limitation on the ACL rule. If a time range with the specified name does not exist, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied immediately. If a time range with specified name exists, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied when the time-range with specified name becomes active. The ACL rule is removed when the time-range with specified name becomes inactive.
Committed Rate / Burst Size	The allowed transmission rate for packets on the interface (Committed Rate), and the number of bytes allowed in a temporary traffic burst (Burst Rate).
Match Criteria (IPv6 ACLs)	The fields in this section specify the criteria to use to determine whether an IP packet matches the rule. The fields described below apply to IPv6 ACLs.
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria, so if Every is selected, no other match criteria can be configured. To configure specific match criteria, this option must be clear.
Protocol	The IANA-assigned protocol number to match within the IP packet. You can also specify one of the following keywords: ICMP, IGMP, TCP, UDP, ICMPv6, or IP.
Fragments	IPv6 ACL rule to match on fragmented IP packets.
Source Prefix/Prefix Length	The IPv6 prefix combined with IPv6 prefix length of the network or host from which the packet is being sent.
Source L4 Port	The TCP/UDP source port to match in the packet header. Select one of the following options: Equal, Not Equal, Less Than, Greater Than, or Range and specify the port number or keyword. TCP port keywords include BGP, Domain, Echo, FTP, FTP Data, HTTP, SMTP, Telnet, WWW, POP2, and POP3. UDP port keywords include Domain, Echo, NTP, RIP, SNMP, TFTP, TIME, and WHO.

Table 383: Add IPv4 ACL Rule Fields (Continued)

Field	Description
Destination Prefix/ Prefix Length	The IPv6 prefix combined with the IPv6 prefix length to be compared to a packet's destination IPv6 address as a match criteria for the IPv6 ACL rule. To indicate a destination host, specify an IPv6 prefix length of 128.
Destination L4 Port	The TCP/UDP destination port to match in the packet header. Select one of the following options: Equal, Not Equal, Less Than, Greater Than, or Range and specify the port number or keyword. TCP port keywords include BGP, Domain, Echo, FTP, FTP Data, HTTP, SMTP, Telnet, WWW, POP2, and POP3. UDP port keywords include Domain, Echo, NTP, RIP, SNMP, TFTP, TIME, and WHO.
ICMP Type	IPv6 ACL rule to match on the specified ICMP message type. This option is available only if the protocol is ICMPv6.
ICMP Code	IPv6 ACL rule to match on the specified ICMP message code. This option is available only if the protocol is ICMPv6.
ICMP Message	IPv6 ACL rule to match on the ICMP message type and code. Specify one of the following supported ICMPv6 messages: Destination-Unreachable, Echo-Request, Echo-Reply, Header, Hop-Limit, MLD-Query, MLD-Reduction, MLD-Report, ND-NA, ND-NS, Next-Header, No-Admin, No-Route, Packet-Too-Big, Port-Unreachable, Router-Solicitation, Router-Advertisement, Router-Renumbering, Time-Exceeded, and Unreachable. This option is available only if the protocol is ICMPv6.
TCP Flags	IPv6 ACL rule to match on the TCP flags. When a + flag is specified, a match occurs if the flag is set in the TCP header. When a - flag is specified, a match occurs if the flag is not set in the TCP header. When Established is specified, a match occurs if either RST or ACK bits are set in the TCP header. This option is available only if the protocol is TCP.
Flow Label	A 20-bit number that is unique to an IPv6 packet, used by end stations to signify quality-of-service handling in routers.
IP DSCP	The IP DSCP value in the IPv6 packet to match to the rule. The DSCP value is defined as the high-order six bits of the Service Type octet in the IPv6 header.
Routing	IPv6 ACL rule to match on routed packets.
Match Criteria (MAC ACLs)	The fields in this section specify the criteria to use to determine whether an Ethernet frame matches the rule. The fields described below apply to MAC ACLs.
Every	When this option is selected, all packets will match the rule and will be either permitted or denied. This option is exclusive to all other match criteria, so if Every is selected, no other match criteria can be configured. To configure specific match criteria, this option must be clear.
CoS	The 802.1p user priority value to match within the Ethernet frame.
Secondary CoS	The secondary 802.1p user priority value to match within the Ethernet frame.
Ethertype	The EtherType value to match in an Ethernet frame. Specify the number associated with the EtherType or specify one of the following keywords: AppleTalk, ARP, IBM SNA, IPv4, IPv6, IPX, MPLS, Unicast, NETBIOS, NOVELL, PPPoE, or RARP.
Source MAC Address / Mask	The MAC address to match to an Ethernet frame's source port MAC address. If desired, enter the MAC Mask associated with the source MAC to match. The MAC address mask specifies which bits in the source MAC to compare against an Ethernet frame. Use F's and zeros in the MAC mask, which is in a wildcard format. An F means that the bit is not checked, and a zero in a bit position means that the data must equal the value given for that bit. For example, if the MAC address is aa_bb_cc_dd_ee_ff, and the mask is 00_00_ff_ff_ff_ff, all MAC addresses with aa_bb_xx_xx_xx_xx result in a match (where x is any hexadecimal number).
Destination MAC Address / Mask	The MAC address to match to an Ethernet frame's destination port MAC address. If desired, enter the MAC Mask associated with the destination MAC to match. The MAC address mask specifies which bits in the destination MAC to compare against an Ethernet frame. Use F's and zeros in the MAC mask, which is in a wildcard format. An F means that the bit is not checked, and a zero in a bit position means that the data must equal the value given for that bit. For example, if the MAC address is aa_bb_cc_dd_ee_ff, and the mask is 00_00_ff_ff_ff_ff, all MAC addresses with aa_bb_xx_xx_xx_xx result in a match (where x is any hexadecimal number).

**Table 383: Add IPv4 ACL Rule Fields (Continued)**

Field	Description
VLAN	The VLAN ID to match within the Ethernet frame.
Secondary VLAN	The secondary VLAN ID to match within the Ethernet frame.
Rule Attributes	The fields in this section provide information about the actions to take on a frame or packet that matches the rule criteria. The attributes specify actions other than the basic Permit or Deny actions.
Assign Queue	The number that identifies the hardware egress queue that will handle all packets matching this rule.
Interface	The interface to use for the action: <ul style="list-style-type: none"> <li>Redirect – Allows traffic that matches a rule to be redirected to the selected interface instead of being processed on the original port. The redirect function and mirror function are mutually exclusive.</li> <li>Mirror – Provides the ability to mirror traffic that matches a rule to the selected interface. Mirroring is similar to the redirect function, except that in flow-based mirroring a copy of the permitted traffic is delivered to the mirror interface while the packet itself is forwarded normally through the device.</li> </ul>
Log	When this option is selected, logging is enabled for this ACL rule (subject to resource availability in the device). If the Access List Trap Flag is also enabled, this will cause periodic traps to be generated indicating the number of times this rule went into effect during the current report interval. A fixed 5 minute report interval is used for the entire system. A trap is not issued if the ACL rule hit count is zero for the current interval.
Redirect External Agent	The number that identifies the external agent that will receive all packets matching this rule.
Time Range Name	The name of the time range that will impose a time limitation on the ACL rule. If a time range with the specified name does not exist, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied immediately. If a time range with specified name exists, and the ACL containing this ACL rule is associated with an interface, the ACL rule is applied when the time-range with specified name becomes active. The ACL rule is removed when the time-range with specified name becomes inactive.
Committed Rate / Burst Size	The allowed transmission rate for frames on the interface (Committed Rate), and the number of bytes allowed in a temporary traffic burst (Burst Rate).

After you click the Resequence Rules button, the Resequence ACL Rules window opens and allows you to resequence rules of the ACL selected from the ACL Identifier field. The following information describes the fields in this window.

**Table 384: Resequence ACL Rules**

Field	Description
Sequence Start	The starting sequence number for resequencing the existing rules.
Sequence Step	The increment of sequence numbers for resequencing the existing rules.

Click Refresh to update the information on the screen.

After you click the + (plus) button next to ACL Remarks, the Add ACL Remark window opens and allows you to add a remark.

Figure 406: Add ACL Remark

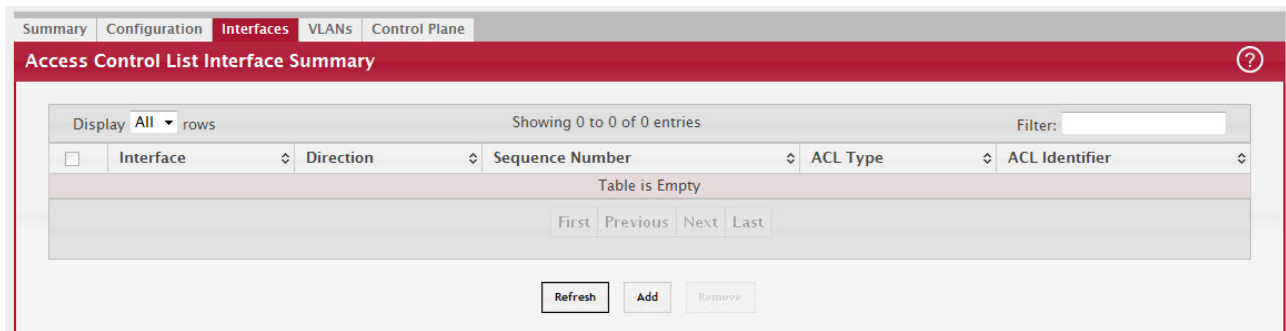


### 8.1.1.3 Access Control List Interface Summary

Use this page to associate one or more ACLs with one or more interfaces on the device. When an ACL is associated with an interface, traffic on the port is checked against the rules defined within the ACL until a match is found. If the traffic does not match any rules within an ACL, it is dropped because of the implicit deny all rule at the end of each ACL.

To display the Access Control List Interface Summary page, click QoS > Access Control Lists > Interfaces in the navigation menu.

Figure 407: Access Control List Interface Summary



Use the buttons to perform the following tasks:

- To apply an ACL to an interface, click Add and configure the settings in the available fields.
- To remove the association between an interface and an ACL, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 385: Access Control List Interface Summary Fields

Field	Description
Interface	The interface that has an associated ACL.
Direction	Indicates whether the packet is checked against the rules in an ACL when it is received on an interface (Inbound) or after it has been received, routed, and is ready to exit an interface (Outbound).
Sequence Number	The order the ACL is applied to traffic on the interface relative to other ACLs associated with the interface in the same direction. When multiple ACLs are applied to the same interface in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.
ACL Type	The type of ACL, which is either IPv4, IPv6, or MAC.
ACL Identifier	The name or number that identifies the ACL. When applying an ACL to an interface, the ACL Identifier menu includes only the ACLs within the selected ACL Type.

### 8.1.1.4 Access Control List VLAN Summary

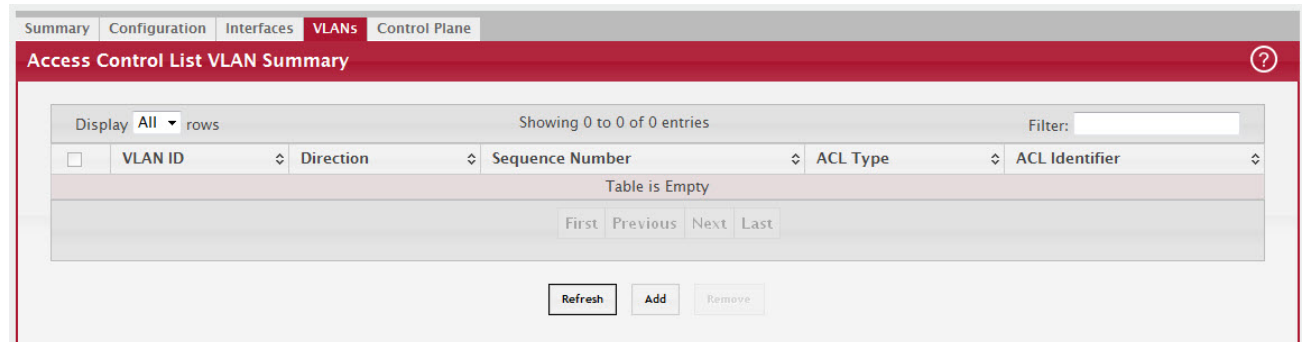
Use this page to associate one or more ACLs with one or more VLANs on the device.

You can also associate an ACL with a VLAN routing interface.

#### NOTICE

To display the Access Control List VLAN Summary page, click QoS > Access Control Lists > VLANs in the navigation menu.

Figure 408: Access Control List VLAN Summary



Use the buttons to perform the following tasks:

- To associate an ACL with a VLAN, click Add and configure the settings in the available fields.
- To remove the association between a VLAN and an ACL, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 386: Access Control List VLAN Summary Fields

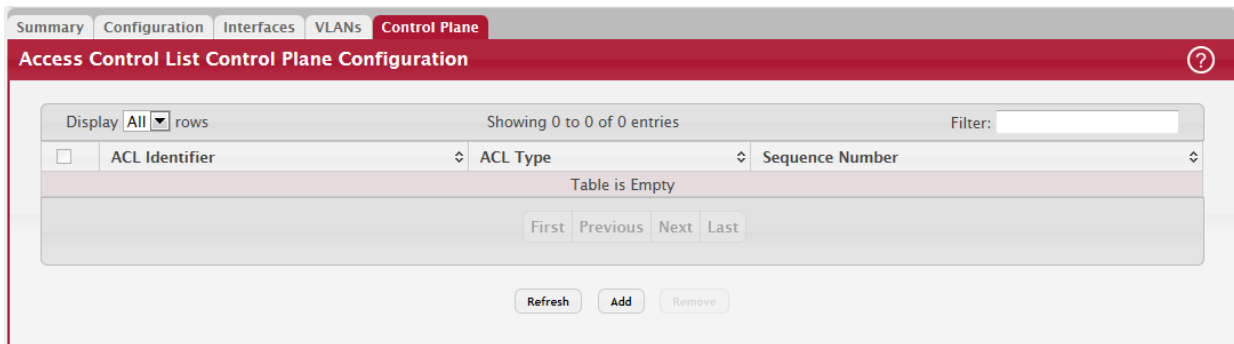
Field	Description
VLAN ID	The ID of the VLAN associated with the rest of the data in the row. When associating a VLAN with an ACL, use this field to select the desired VLAN.
Direction	Indicates whether the packet is checked against the rules in an ACL when it is received on a VLAN (Inbound) or after it has been received, routed, and is ready to exit a VLAN (Outbound).
Sequence Number	The order the ACL is applied to traffic on the VLAN relative to other ACLs associated with the VLAN in the same direction. When multiple ACLs are applied to the same VLAN in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.
ACL Type	The type of ACL, which is either IPv4, IPv6, or MAC.
ACL Identifier	The name or number that identifies the ACL. The permitted identifier depends on the ACL type. Standard and Extended IPv4 ACLs use numbers within a set range, and Named IPv4, IPv6, and MAC ACLs use alphanumeric characters.

### 8.1.1.5 Access Control List Control Plane Configuration

Use this page to define controlled management access to the device. Control plane ACLs allow you to determine which addresses or protocols are allowed to access the management interface on the device. The control plane ACLs are applied to management access through the in-band (production network) ports only. Inbound traffic on the CPU port is checked against the rules defined within the ACL until a match is found. If the traffic does not match any rules within an ACL, it is dropped because of the implicit deny all rule at the end of each ACL.

To display the Access Control List Control Plane Configuration Page, click QoS > Access Control Lists > Control Plane in the navigation menu.

Figure 409: Access Control List Control Plane Configuration



Use the buttons to perform the following tasks:

- To apply an ACL to the CPU interface, click Add and configure the settings in the available fields.
- To remove the association between the CPU interface and an ACL, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 387: Access Control List Control Plane Configuration Fields

Field	Description
ACL Identifier	The name or number that identifies the ACL.
ACL Type	The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic. The ACL types are as follows: <ul style="list-style-type: none"> <li>• IPv4 Standard – Match criteria is based on the source address of IPv4 packets.</li> <li>• IPv4 Extended – Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of IPv4 packets.</li> <li>• IPv4 Named – Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number.</li> <li>• IPv6 Named – Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within IPv6 packets.</li> <li>• Extended MAC – Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within Ethernet frames.</li> </ul>
Sequence Number	The order the ACL is applied to traffic on the interface relative to other ACLs associated with the interface in the same direction. When multiple ACLs are applied to the same interface in the same direction, the ACL with the lowest sequence number is applied first, and the other ACLs are applied in ascending numerical order.

## 8.1.2 IPv6 ACL Rules

The maximum number of IPv6 rules depends on the following factors (also refer to the FASTPATH Scaling Parameters and Values for the maximum number of rules per device type):

- If both SRC IPv6 and DST IPv6 are part of the ACL rule, then the maximum number of rules is one quarter the possible number for that device type.
- If DSCP is part of the rule along with any other qualifier, then the maximum number of rules possible are one quarter the possible number for that device type.
- In all other cases, the maximum number of rules are equal to half the maximum possible for that device type or 1021, whichever is smaller.



### 8.1.2.1 Scenarios

In the following scenarios, the BCM56334 device is used (1789 rules maximum).


- Scenario #1: If the rules have both SRC IPv6 and DST IPv6, then maximum rules possible are  $1789/4 = 447$ .
- Scenario #2: If the rules have DSCP along with any other qualifier, then the maximum number of rules possible are  $1789/4 = 447$ .
- Scenario #3: In all the other cases, 894 rules can be accommodated.

### 8.1.2.2 Access Control List Statistics

Use this page to display the statistical information about the packets forwarded or discarded by the port that matches the configured rules within an Access Control List (ACL). Each ACL rule is configured to match one or more aspects of traffic on the network. When a packet matches the conditions in a rule, the counter associated with the rule gets incremented, until it reaches the rollover value of the counter. ACL counters do not interact with DiffServ policies or Policy-based Routing counters.

To display the Access Control List Statistics page, click QoS > Access Control Lists > Statistics in the navigation menu.

**Figure 410: Access Control List Statistics**



Use the buttons to perform the following tasks:

- To clear the hit count for one or more configured rules within an ACL, select the rule entry and click Clear Rule Counter. You must confirm the action before the hit count is cleared for the selected rule(s).
- To clear the hit count for an ACL, select the ACL ID from the ACL Identifier menu and click Clear ACL Counters. You must confirm the action before the hit count is cleared for the selected ACL.
- To clear the hit count for an ACL type, select the type from the ACL Type menu and select All from the ACL Identifier menu and then click Clear ACL Counters. You must confirm the action before the hit count is cleared for the selected ACL type.

**Table 388: Access Control List Statistics Fields**

Field	Description
ACL Type	<p>The type of ACL. The ACL type determines the criteria that can be used to match packets. The type also determines which attributes can be applied to matching traffic. IPv4 ACLs classify Layer 3 and Layer 4 IPv4 traffic, IPv6 ACLs classify Layer 3 and Layer 4 IPv6 traffic, and MAC ACLs classify Layer 2 traffic.</p> <p>The ACL types are as follows:</p> <ul style="list-style-type: none"> <li>• IPv4 Standard – Match criteria is based on the source address of the IPv4 packets.</li> <li>• IPv4 Extended – Match criteria can be based on the source and destination addresses, source and destination Layer 4 ports, and protocol type of the IPv4 packets.</li> <li>• IPv4 Named – Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number.</li> <li>• IPv6 Named – Match criteria can be based on information including the source and destination IPv6 addresses, source and destination Layer 4 ports, and protocol type within the IPv6 packets.</li> <li>• Extended MAC – Match criteria can be based on the source and destination MAC addresses, 802.1p user priority, VLAN ID, and EtherType value within the Ethernet frames.</li> </ul>
ACL Identifier	A list of ACL IDs that exist on the system for a given ACL type. To view the rule(s) within an ACL, you must select the ID of the ACL from the list. The ACL rules are not displayed when option All is selected. Option All lets you clear the hit count for an ACL type.
Sequence Number	The number that indicates the position of a rule within the ACL.



Table 388: Access Control List Statistics Fields (Continued)

Field	Description
Action	The action to take when a packet or frame matches the criteria in the rule: <ul style="list-style-type: none"> <li>Permit – The packet or frame is forwarded.</li> <li>Deny – The packet or frame is dropped.</li> </ul>
Match Conditions	The criteria used to determine whether a packet or frame matches the ACL rule.
Rule Attributes	Each action — beyond the basic Permit and Deny actions — to perform on the traffic that matches the rule.
Hit Count	Indicates the number of packets that match the configured rule in an ACL. If a rule is configured without rate limit, then the hit count is the number of matched packets forwarded or discarded by the port. If a rule is configured with rate limit, then if the sent traffic rate exceeds the configured rate, the hit count displays the matched packet count equal to the sent rate, despite packets getting dropped beyond the configured limit. If the sent traffic rate is less than the configured rate, the hit count displays only the matched packet count.

## 8.2 Configuring Auto VoIP

Voice over Internet Protocol (VoIP) allows you to make telephone calls using a computer network over a data network like the Internet. With the increased prominence of delay-sensitive applications (voice, video, and other multimedia applications) deployed in networks today, proper QoS configuration will ensure high-quality application performance. The Auto VoIP feature is intended to provide an easy classification mechanism for voice packets so that they can be prioritized above data packets to provide better QoS.

The Auto-VoIP feature explicitly matches VoIP streams in Ethernet switches and provides them with a better class of service than ordinary traffic. If you enable the Auto-VoIP feature on an interface, the interface scans incoming traffic for the following call-control protocols:

- Session Initiation Protocol (SIP)
- H.323
- Skinny Client Control Protocol (SCCP)

When a call-control protocol is detected the switch assigns the traffic in that session to the highest CoS queue, which is generally used for time-sensitive traffic.

### 8.2.1 Auto VoIP Global Configuration

Use this page to configure the VLAN ID for the Auto VoIP VLAN or to reset the current Auto VoIP VLAN ID to the default value. Voice over Internet Protocol (VoIP) enables telephone calls over a data network. Because voice traffic is typically more time-sensitive than data traffic, the Auto VoIP feature helps provide a classification mechanism for voice packets so that they can be prioritized above data packets to provide better Quality of Service (QoS). With the Auto VoIP feature, voice prioritization is provided based on call-control protocols (SIP, SCCP, H.323) and/or OUI bits. When the device identifies voice traffic, it is placed in the VLAN specified on this page. The Auto VoIP feature does not rely on LLDP-MED support in connected devices.

To display the Auto VoIP Global Configuration page, click Quality of Service > Auto VoIP > Global in the navigation menu.

Figure 411: Auto VoIP Global Configuration

Table 389: Auto VoIP Global Configuration Fields

Field	Description
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic.
Reset (Button)	Click this button to reset the voice VLAN to the default value.

## 8.2.2 OUI Table Summary

Use this page to add and remove Organizationally Unique Identifiers (OUIs) from the OUI database the device maintains. Device hardware manufacturers can include an OUI in a network adapter to help identify the device. The OUI is a unique 24-bit number assigned by the IEEE registration authority. Several default OUIs have been preconfigured in the OUI database on the device.

To display the Auto VoIP OUI Table page, click Quality of Service > Auto VoIP > OUI Table in the navigation menu.

Figure 412: OUI Table Summary

Telephone OUI	Status	Description
00:01:E3	Default	SIEMENS
00:03:6B	Default	CISCO1
00:12:43	Default	CISCO2
00:0F:E2	Default	H3C
00:60:89	Default	NITSUKO
00:D0:1E	Default	PINTEL
00:E0:75	Default	VERILINK
00:E0:BB	Default	3COM
00:04:0D	Default	AVAYA1
00:1B:4F	Default	AVAYA2

Use the buttons to perform the following tasks:

- To add an OUI, click Add and specify an OUI and its description in the available fields.
- To remove one or more configured OUIs, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 390: OUI Table Summary Fields

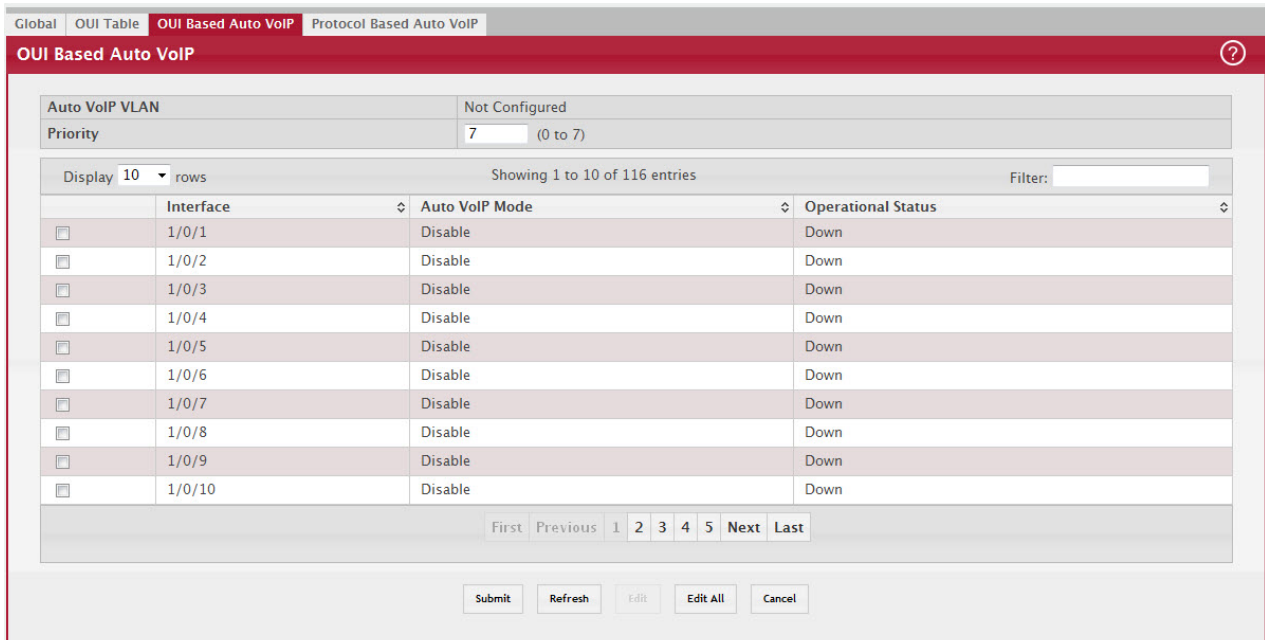
Field	Description
Telephony OUI	The unique OUI that identifies the device manufacturer or vendor. The OUI is specified in three octet values (each octet is represented as two hexadecimal digits) separated by colons.
Status	Identifies whether the OUI is preconfigured on the system (Default) or added by a user (Configured).
Description	Identifies the manufacturer or vendor associated with the OUI.

### 8.2.3 OUI Based Auto VoIP

Use this page to configure the Organizationally Unique Identifier (OUI) based Auto VoIP priority and to enable or disable the Auto VoIP mode on the interfaces.

To display the Auto VoIP OUI Table page, click Quality of Service > Auto VoIP > OUI Based Auto VoIP in the navigation menu.

Figure 413: OUI Based Auto VoIP



Use the buttons to perform the following tasks:

- To configure the settings for one or more interfaces, select each entry to modify and click Edit.
- To apply the same settings to all interfaces, click Edit All.

Table 391: OUI Based Auto VoIP Fields

Field	Description
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic. All VoIP traffic that matches a value in the known OUI list gets assigned to this VoIP VLAN.
Priority	The 802.1p priority used for traffic that matches a value in the known OUI list. If the Auto VoIP mode is enabled and the interface detects an OUI match, the device assigns the traffic in that session to the traffic class mapped to this priority value. Traffic classes with a higher value are generally used for time-sensitive traffic.
Interface	The interface associated with the rest of the data in the row. When editing Auto VoIP settings on one or more interfaces, this field identifies the interface(s) being configured.
Auto VoIP Mode	The administrative mode of OUI-based Auto VoIP on the interface.
Operational Status	The operational status of an interface. To be up, an interface must be administratively enabled and have a link.

## 8.2.4 Protocol Based Auto VoIP

Use this page to configure the protocol-based Auto VoIP priority settings and to enable or disable the protocol-based Auto VoIP mode on the interfaces.

To display the Protocol Based Auto VoIP page, click Quality of Service > Auto VoIP > Protocol Based Auto VoIP in the navigation menu. A portion of the web page is shown below.

Figure 414: Protocol Based Auto VoIP

Interface	Auto VoIP Mode	Operational Status
1/0/1	Disable	Down
1/0/2	Disable	Down
1/0/3	Disable	Down
1/0/4	Disable	Down
1/0/5	Disable	Down
1/0/6	Disable	Down
1/0/7	Disable	Down
1/0/8	Disable	Down
0/1/1	Disable	Down
0/1/2	Disable	Down

Use the buttons to perform the following tasks:

- To configure the settings for one or more interfaces, select each entry to modify and click Edit.
- To apply the same settings to all interfaces, click Edit All.

Table 392: Protocol Based Auto VoIP Fields

Field	Description
Auto VoIP VLAN	The VLAN used to segregate VoIP traffic from other non-voice traffic. All VoIP traffic in a session identified by the call-control protocol gets assigned to this VoIP VLAN.
Prioritization Type	The method used to prioritize VoIP traffic when a call-control protocol is detected, which is one of the following: <ul style="list-style-type: none"> <li>• Remark – Remark the voice traffic with the specified 802.1p priority value at the ingress interface.</li> <li>• Traffic Class – Assign VoIP traffic to the specified traffic class when egressing the interface.</li> </ul>
802.1p Priority	The 802.1p priority used for protocol-based VoIP traffic. This field can be configured if the Prioritization Type is 802.1p Priority. If the Auto VoIP mode is enabled and the interface detects a call-control protocol, the device marks traffic in that session with the specified 802.1p priority value to ensure voice traffic always gets the highest priority throughout the network path. Egress tagging must be administratively enabled on the appropriate uplink port to carry the remarked priority at the egress port.

Table 392: Protocol Based Auto VoIP Fields (Continued)

Field	Description
Traffic Class	The traffic class used for protocol-based VoIP traffic. This field can be configured if the Prioritization Type is Traffic Class. If the Auto VoIP mode is enabled and the interface detects a call-control protocol, the device assigns the traffic in that session to the configured Class of Service (CoS) queue. Traffic classes with a higher value are generally used for time-sensitive traffic. The CoS queue associated with the specified traffic class should be configured with the appropriate bandwidth allocation to allow priority treatment for VoIP traffic.
Interface	The interface associated with the rest of the data in the row. When editing Auto VoIP settings on one or more interfaces, this field identifies the interface(s) being configured.
Auto VoIP Mode	The administrative mode of the Auto VoIP feature on the interface: <ul style="list-style-type: none"> <li>• Enable – The interface scans incoming traffic for the following call-control protocols: <ul style="list-style-type: none"> <li>- Session Initiation Protocol (SIP)</li> <li>- H.323</li> <li>- Skinny Client Control Protocol (SCCP)</li> </ul> </li> <li>• Disable – The interface does not use the Auto VoIP feature to scan for call-control protocols.</li> </ul>
Operational Status	The operational status of an interface. To be up, an interface must be administratively enabled and have a link.

- If you change any of the settings on the page, click Submit to send the updated configuration to the switch. Configuration changes take effect immediately. These changes will not be retained across a power cycle unless a Save is performed.
- Click Refresh to update the page with the most current data from the switch.

## 8.3 Configuring Class of Service

The Class of Service (CoS) queuing feature lets you directly configure certain aspects of switch queuing. This provides the desired QoS behavior for different types of network traffic when the complexities of DiffServ are not required. The priority of a packet arriving at an interface can be used to steer the packet to the appropriate outbound CoS queue through a mapping table. CoS queue characteristics that affect queue mapping, such as minimum guaranteed bandwidth, transmission rate shaping, etc., are user-configurable at the queue (or port) level.

Seven queues per port are supported. Although the hardware supports eight queues, one queue is always reserved for internal use by the stacking subsystem.

### 8.3.1 IP DSCP Mapping Configuration

Use the IP DSCP Mapping Configuration page to map an IP DSCP value to an internal traffic class.

To display the IP DSCP Mapping Configuration page, click QoS > Class of Service > IP DSCP in the navigation menu.

Figure 415: CoS IP DSCP Mapping Configuration

Interface	Global
IP DSCP	Traffic Class
0	<input type="radio"/> 0 <input checked="" type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5 <input type="radio"/> 6
1	<input type="radio"/> 0 <input checked="" type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5 <input type="radio"/> 6
2	<input type="radio"/> 0 <input checked="" type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5 <input type="radio"/> 6
3	<input type="radio"/> 0 <input checked="" type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5 <input type="radio"/> 6
4	<input type="radio"/> 0 <input checked="" type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5 <input type="radio"/> 6
5	<input type="radio"/> 0 <input checked="" type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5 <input type="radio"/> 6
6	<input type="radio"/> 0 <input checked="" type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5 <input type="radio"/> 6
7	<input type="radio"/> 0 <input checked="" type="radio"/> 1 <input type="radio"/> 2 <input type="radio"/> 3 <input type="radio"/> 4 <input type="radio"/> 5 <input type="radio"/> 6

Table 393: IP DSCP Mapping Configuration Fields

Field	Description
Interface	The menu contains all CoS configurable interfaces. The only option is Global, which means that the IP DSCP mapping configuration applies to all interfaces and cannot be applied on a per-interface basis.
IP DSCP Values	Lists the IP DSCP values to which you can map an internal traffic class. The values range from 0-63.
Traffic Class	The traffic class is the hardware queue for a port. Higher traffic class values indicate a higher queue position. Before traffic in a lower queue is sent, it must wait for traffic in higher queues to be sent. Valid range is 0 to 6.

If you make changes to the page, click Submit to apply the changes to the system. Click Restore Defaults to reset all interfaces to the default trust value.

### 8.3.2 Interface Configuration

Use the Interface Configuration page to apply an interface shaping rate to all ports or to a specific port.

To display the Interface Configuration page, click QoS > Class of Service > Interface in the navigation menu.

Figure 416: Interface Configuration

Table 394: Interface Configuration Fields

Field	Description
Interface	Selects the CoS configurable interface to be affected by the Interface Shaping Rate. Select Global to apply a rate to all interfaces. Select an individual port to override the global setting.
Interface Shaping Rate	Sets the limit on how much traffic can leave a port. The limit on maximum transmission bandwidth has the effect of smoothing temporary traffic bursts over time so that the transmitted traffic rate is bounded. The specified value represents a percentage of the maximum negotiated bandwidth. The default value is zero (0). Valid values are 0 to 100, in increments of 1. A value of 0 means the maximum is unlimited.
WRED Decay Exponent	Specifies the decay exponent value used with the WRED average queue length calculation algorithm. Default value is 9. Valid Range is (0 to 15).

If you make changes to the page, click Submit to apply the changes to the system. Click Restore Defaults to reset all interfaces to the default trust value.

### 8.3.3 Interface Queue Configuration

Use the Interface Queue Configuration page to define what a particular queue does by configuring switch egress queues. User-configurable parameters control the amount of bandwidth used by the queue, the queue depth during times of congestion, and the scheduling of packet transmission from the set of all queues on a port. Each port has its own CoS queue-related configuration.

The configuration process is simplified by allowing each CoS queue parameter to be configured globally or per-port. A global configuration change is automatically applied to all ports in the system.

To display the Interface Queue Configuration page, click QoS > Class of Service > Queue in the navigation menu.

Figure 417: Interface Queue Configuration

Queue ID	Minimum Bandwidth (%)	Scheduler Type	Queue Management Type
0	0	Weighted	TailDrop
1	0	Weighted	TailDrop
2	0	Weighted	TailDrop
3	0	Weighted	TailDrop
4	0	Weighted	TailDrop
5	0	Weighted	TailDrop
6	0	Weighted	TailDrop

Table 395: Interface Queue Configuration Fields

Field	Description
Interface	Specifies the interface (physical, LAG, or Global) to configure.
Minimum Bandwidth Allocated	Shows the sum of individual Minimum Bandwidth values for all queues in the interface. The sum cannot exceed the defined maximum of 100. This value is considered while configuring the Minimum Bandwidth for a queue in the selected interface.
Queue ID	Use the menu to select the queue per interface to be configured.
Minimum Bandwidth	Specify the minimum guaranteed bandwidth allocated to the selected queue on the interface. Setting this value higher than its corresponding Maximum Bandwidth automatically increases the maximum to the same value. The default value is 0. The valid range is 0 to 100, in increments of 1. The value zero (0) means no guaranteed minimum. The sum of individual Minimum Bandwidth values for all queues in the selected interface cannot exceed defined maximum 100.
Scheduler Type	Selects the type of queue processing from the drop-down menu. Options are <i>Weighted</i> and <i>Strict</i> . Defining on a per-queue basis allows the user to create the desired service characteristics for different types of traffic. <ul style="list-style-type: none"> <li>Weighted: Weighted round robin associates a weight to each queue. This is the default.</li> <li>Strict: Strict priority services traffic with the highest priority on a queue first</li> </ul>
Queue Management Type	Displays the type of queue depth management techniques used for all queues on this interface. This is only used if the device supports independent settings per-queue. Queue Management Type can only be Taildrop. The default value is Taildrop. All packets on a queue are safe until congestion occurs. At this point, any additional packets queued are dropped.

- If you make changes to the page, click Submit to apply the changes to the system.
- Click Restore Defaults for all Queues to reset the settings for the selected interface.
- To reset the defaults for all interfaces, select Global from the Slot/Port menu before you click the button.

### 8.3.4 CoS Interface Queue Drop Precedence Configuration

Use this page to configure the queue drop precedence on a per-queue, per-interface basis. When an interface is configured with taildrop queue management, all packets on a queue are safe until congestion occurs. If congestion occurs, any additional packets queued are dropped. Weighted Random Early Detection (WRED) drops packets selectively based their drop precedence level.

To display the CoS Interface Queue Configuration page, click QoS > Class of Service > Drop Precedence in the navigation menu.

Figure 418: CoS Interface Queue Drop Precedence Configuration

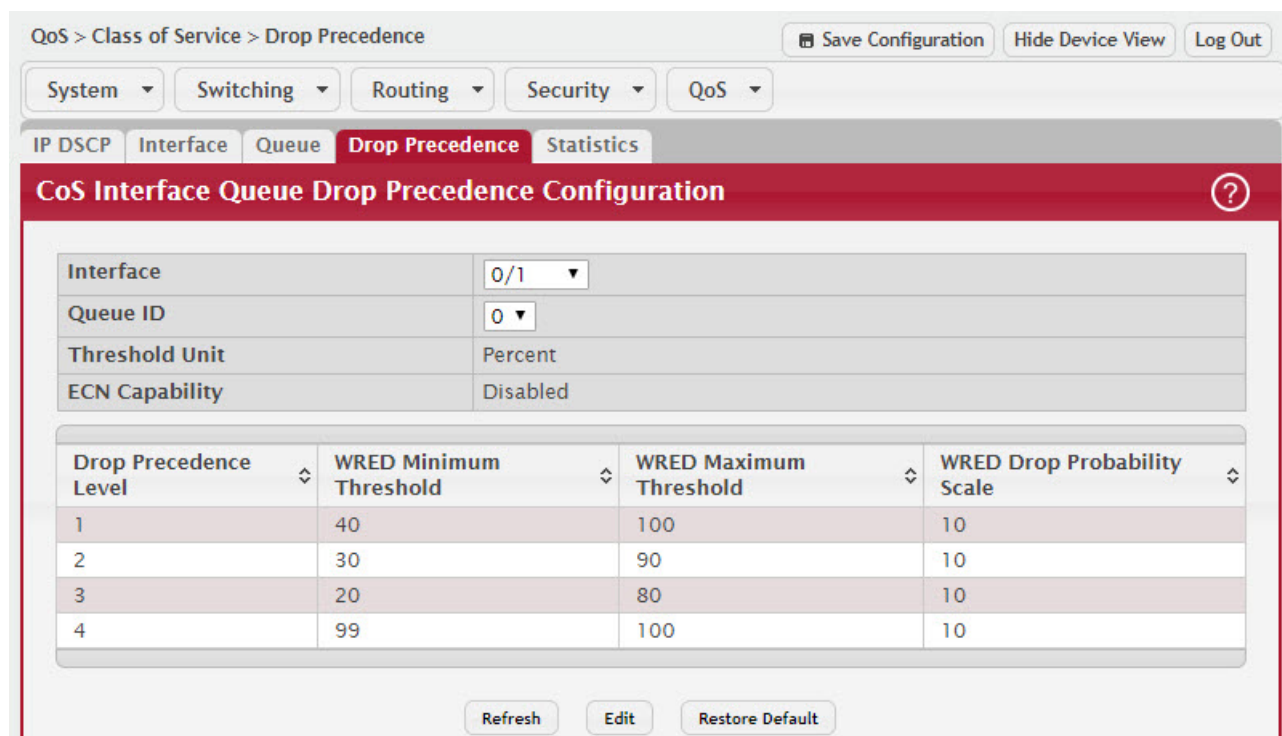


Table 396: CoS Interface Queue Configuration Fields

Field	Description
Interface	The interface on which to configure the queue drop precedence settings. To configure the same settings on all interfaces, select the Global menu option.
Queue ID	The CoS queue on which to configure the drop precedence settings. The higher the queue value, the higher its priority is for sending traffic.
Threshold Unit	The unit for the WRED minimum and maximum threshold values. Threshold unit can be in percentage or Kbytes.
ECN Capability	The Explicit Congestion Notification (ECN) marking on the CoS queue. When ECN capability is enabled, packets marked as ECN capable and exceeding the upper WRED threshold are not dropped. In case of extreme congestion, ECN capable packets may be dropped. <b>Note:</b> ECN statistics is a feature and only available if supported on the platform.



**Table 396: CoS Interface Queue Configuration Fields (Continued)**

Field	Description
Drop Precedence Level	The four drop precedence levels as follows: <ul style="list-style-type: none"> <li>Green–Low drop level 1 for classified TCP packets.</li> <li>Yellow–Medium drop level 2 for classified TCP packets.</li> <li>Red–High drop level 3 for classified TCP packets.</li> <li>Non-TCP–Drop level 4 for non-TCP classified packets.</li> </ul>
WRED Minimum Threshold	The minimum queue threshold below which now packets are dropped for the associated drop precedence level. After the minimum is reached, WRED randomly drops packets based on their priority (DSCP or IP precedence). This setting applies to the interface if it is configured with a WRED queue management type.
WRED Maximum Threshold	The maximum queue threshold above which all packets are dropped for the associated drop precedence level. After the maximum is reached, WRED drops all packets based on their priority (DSCP or IP precedence). This setting applies to the interface if it is configured with a WRED queue management type.
WRED Drop Probability Scale	The packet drop probability for the drop precedence level. This setting applies to the interface if it is configured with a WRED queue management type.

- Click Refresh to refresh the settings displayed on the page.
- Click Restore Defaults to restore all drop precedence settings on the selected interface to the default values. If Global is selected from the Interface menu, all default settings for all interfaces are restored.

### 8.3.4.1 Edit CoS Interface Queue Drop Precedence Configuration

- Click the Edit button to open the Edit CoS Interface Queue Drop Precedence Configuration page. Use this page to configure the per-interface CoS queue settings:
  - Threshold Unit in in percentage or Kbytes.
  - Enable or disable ECN Capability.
  - WRED minimum and maximum thresholds in percentage or Kbytes.
  - WRED drop probability scale.

**Figure 419: Edit CoS Interface Queue Drop Precedence Configuration**

Interface	1/0/1			
Queue ID	0			
Threshold Unit	<input checked="" type="radio"/> Percent <input type="radio"/> Kbytes			
ECN Capability	<input checked="" type="radio"/> Disable <input type="radio"/> Enable			
	Drop Precedence Level			
	1 - Green	2 - Yellow	3 - Red	4 - Non-TCP
WRED Minimum Threshold	40	30	20	99
WRED Maximum Threshold	100	90	80	100
WRED Drop Probability Scale	10	10	10	10

- If you make changes to the page, click Submit to apply the changes to the system.

- Click Cancel to cancel the changes.

### 8.3.5 CoS Statistics

Use the CoS Statistics page to view and clear the CoS statistical information about traffic utilization and color drops for each interface and per CoS queue. Statistics that are not supported in the hardware are displayed as “-” on the web page.

The CoS Statistics page is only available when the ECN feature is supported on the device.

**NOTICE**

To display the CoS Statistics page, click QoS > Class of Service > Statistics in the navigation menu.

Figure 420: CoS Statistics

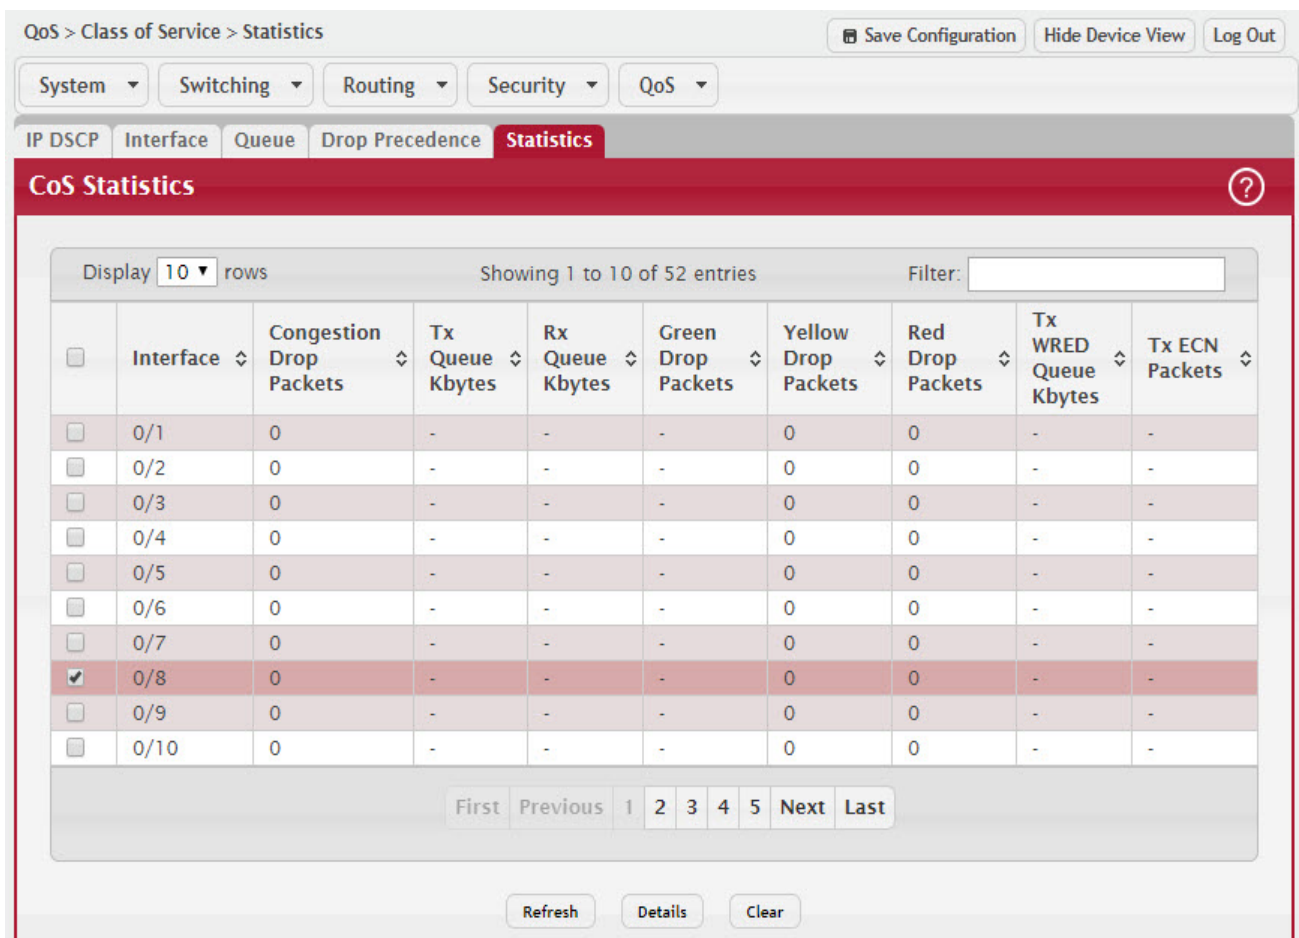


Table 397: CoS Statistics Fields

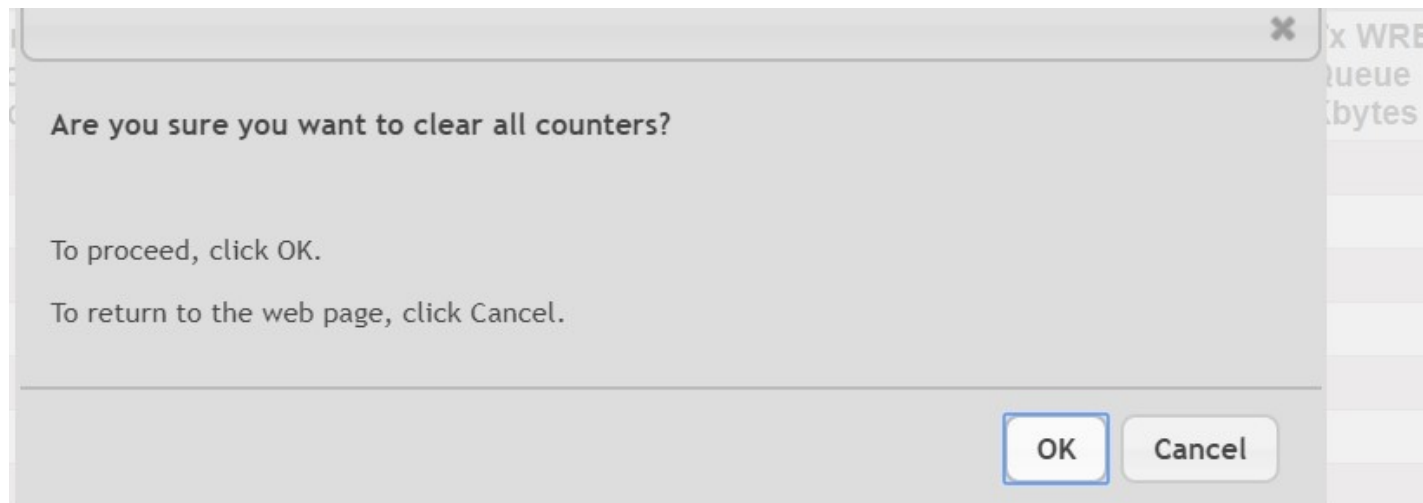
Field	Description
Interface	The interface associated with the rest of the data in the row.
Congestion Drop Packets	The total number of packets dropped on the interface.
Tx Queue Kbytes	The total number of kbytes transmitted on the interface.
Rx Queue Kbytes	The total number of kbytes received on the interface.
Green Drop Packets	The total number of packets dropped on the interface that were colored green.
Yellow Drop Packets	The total number of packets dropped on the interface that were colored yellow.

Table 397: CoS Statistics Fields (Continued)

Field	Description
Red Drop Packets	The total number of packets dropped on the interface that were colored red.
Tx WRED Queue Kbytes	The average queue size transmitted on the interface.
Tx ECN Packets	The total number of ECN packets transmitted on the interface.

Click the Clear button to reset the CoS statistics counters to the default values on one or more selected interfaces. The confirmation page below displays.

Figure 421: Clear CoS Queue Statistics



Click OK to confirm the action to reset the counter values for the selected interfaces.

---

**NOTICE**

The CoS-Queue Peak Kbytes count enqueued to the CoS queues on each interface is not cleared as it is a status value and not a counter.

---

Click Cancel to cancel resetting the counter values for the selected interfaces.

To view additional statistical information per CoS queue for an interface, select the interface and click the Details button.

Figure 422: CoS Queue Statistics Details

Queue ID	Total Drop Packets	Total Kbytes	Peak Kbytes	Current Kbytes	Average Kbytes
0	10	-	-	-	-
1	10	-	-	-	-
2	10	-	-	-	-
3	10	-	-	-	-
4	10	-	-	-	-
5	10	-	-	-	-
6	10	-	-	-	-

Table 398: CoS Queue Statistics Details

Field	Description
Interface	The interface associated with the CoS queue statistic details.
Queue ID	The CoS queue associated with the rest of the data in the row.
Total Drop Packets	The total number of packets dropped for any reason for the associated queue.
Total Kbytes	The total number of kbytes enqueued to the associated queue.
Peak Kbytes	The total number of peak kbytes enqueued to the associated queue. The peak count is not cleared when counters are reset to the default values, as it is a status value and not a counter.
Current Kbytes	The total number of current kbytes enqueued to the associated queue.
Average Kbytes	The total number of average kbytes enqueued to the associated queue.

## 8.4 Configuring DiffServ

Use this page to configure the administrative mode of Differentiated Services (DiffServ) support on the device and to view the current and maximum number of entries in each of the main DiffServ private MIB tables. DiffServ allows traffic to be classified into streams and given certain QoS treatment in accordance with defined per-hop behaviors.

Packets are classified and processed based on defined criteria. The classification criteria is defined by a class. The processing is defined by a policy's attributes. Policy attributes may be defined on a per-class instance basis, and it is these attributes that are applied when a match occurs. A policy can contain multiples classes. When the policy is active, the actions taken depend on which class matches the packet.

### 8.4.1 DiffServ Global Configuration and Status

Use this page to configure the Global DiffServ settings on the device.

To display the DiffServ Global Configuration and Status page, click QoS > DiffServ > Global in the navigation menu.

Figure 423: DiffServ Global Configuration and Status

MIB Table	Current Number / Maximum Number
Class Table	0 / 32
Class Rule Table	0 / 416
Policy Table	0 / 64
Policy Instance Table	0 / 1792
Policy Attribute Table	0 / 5376
Service Table	0 / 576

Table 399: DiffServ Global Configuration and Status Fields

Field	Description
Diffserv Admin Mode	The administrative mode of DiffServ on the device. While disabled, the DiffServ configuration is retained and can be changed, but it is not active. While enabled, Differentiated Services are active.
MIB Table	The information in this table displays the number of entries (rows) that are currently in each of the main DiffServ private MIB tables and the maximum number of rows that can exist in each table.
Class Table	The current and maximum number of classifier entries in the table. DiffServ classifiers differentiate among traffic types.
Class Rule Table	The current and maximum number of class rule entries in the table. Class rules specify the match criteria that belong to a class definition.
Policy Table	The current and maximum number of policy entries in the table. The policy determines the traffic conditioning or service provisioning actions applied to a traffic class.
Policy Instance Table	The current and maximum number of policy-class instance entries in the table. A policy-class instance is a policy that is associated with an existing DiffServ class.
Policy Attribute Table	The current and maximum number of policy attribute entries in the table. A policy attribute entry attaches various policy attributes to a policy-class instance.
Service Table	The current and maximum number of service entries in the table. A service entry associates a DiffServ policy with an interface and inbound or outbound direction.

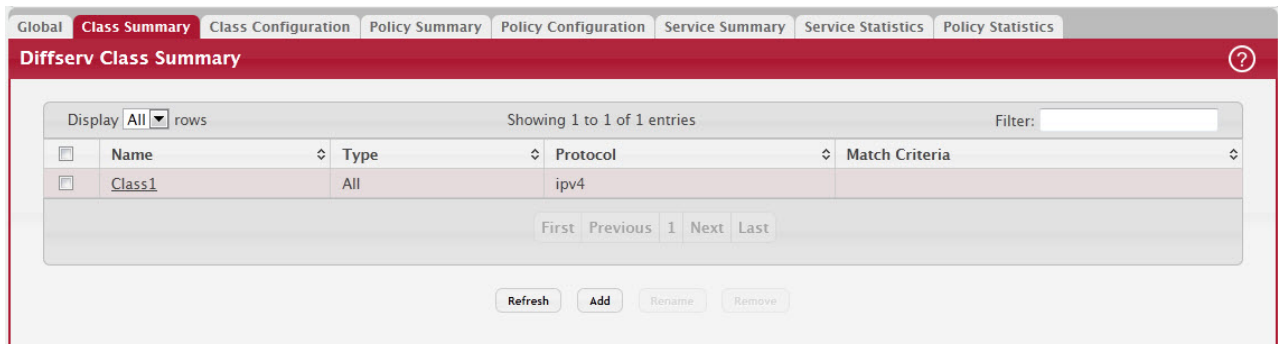
- If you make changes to the page, click Submit to apply the changes to the system. Click Restore Defaults to reset all interfaces to the default trust value.
- Click Refresh to update the page with the most current data from the switch.

## 8.4.2 DiffServ Class Summary

Use this page to create or remove DiffServ classes and to view summary information about the classes that exist on the device. Creating a class is the first step in using DiffServ to provide Quality of Service. After a class is created, you can define the match criteria for the class.

To display the DiffServ Class Summary and Status page, click QoS > DiffServ > Class Summary in the navigation menu.

Figure 424: DiffServ Class Summary



Use the buttons to perform the following tasks:

- To add a DiffServ class, click Add and complete the fields in the Add Class window.
- To change the name of an existing class, select the entry to modify and click Rename.
- To remove one or more configured classes, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 400: DiffServ Class Summary Fields

Field	Description
Name	The name of the DiffServ class. When adding a new class or renaming an existing class, the name of the class is specified in the Class field of the dialog window.
Type	The class type, which is one of the following: <ul style="list-style-type: none"> <li>• All—All the various match criteria defined for the class should be satisfied for a packet match. All signifies the logical AND of all the match criteria.</li> <li>• Any—Any of various match criteria defined for the class can be satisfied for a packet match.</li> </ul>
Protocol	The Layer 3 protocol to use for filtering class types, which is either IPv4 or IPv6.
Match Criteria	The criteria used to match packets.

Click Refresh to update the page with the most current data from the switch.

### 8.4.3 DiffServ Class Configuration

Use this page to define the criteria to associate with a DiffServ class. As packets are received or transmitted, these DiffServ classes are used to classify and prioritize packets. Each class can contain multiple match criteria.

After you select the class to configure from the Class menu, use the buttons to perform the following tasks:

- To define criteria for matching packets within a class, click Add Match Criteria. Once you add a match criteria entry to a class, you cannot edit or remove the entry. However, you can add more match criteria entries to a class until the maximum number of entries has been reached for the class.
- To remove the associated reference class from the selected class, click Remove Reference Class. Note that unless the reference class is the last entry in the list of match criteria, the Reference Class match type remains in the list as a placeholder, but the associated value is N/A, and the previously referenced class is removed.
- To display the DiffServ Class Configuration and Status page, click QoS > DiffServ > Class Configuration in the navigation menu.

Figure 425: DiffServ Class Configuration

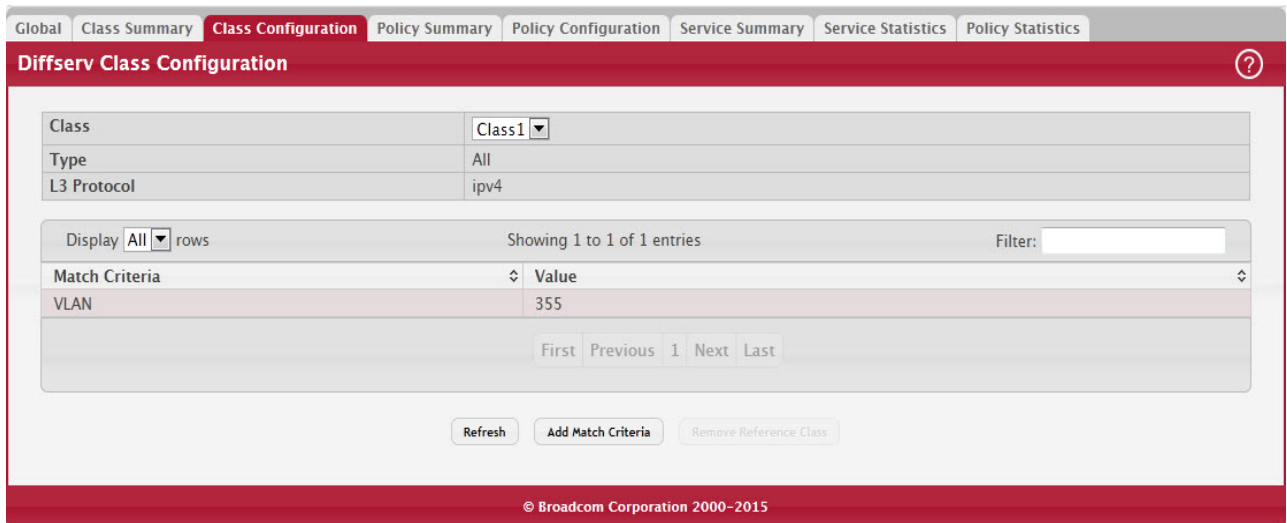


Table 401: DiffServ Class Configuration Fields

Field	Description
Class	The name of the class. To configure match criteria for a class, select its name from the menu.
Type	The class type, which is one of the following: <ul style="list-style-type: none"> <li>All – All the various match criteria defined for the class should be satisfied for a packet match. All signifies the logical AND of all the match criteria.</li> <li>Any – Any of various match criteria defined for the class can be satisfied for a packet match.</li> </ul>
L3 Protocol	The Layer 3 protocol to use for filtering class types, which is either IPv4 or IPv6.
Match Criteria	The type of match criteria defined for the selected class. If the Type is ACL, no information about the match criteria is available on this page.
Value	The configured value of the match criteria that corresponds to the match type.
Any	Select this option to specify that all packets are considered to match the specified class. There is no need to configure additional match criteria if Any is selected because a match will occur on all packets.
Reference ACL	Select this option to require packets to match the criteria defined in the associated ACL. When associating an ACL, the ACL (IP or MAC) options are available only if at least one IP or MAC ACL exists on the device. After you select this option, the <code>ACL Identifier</code> field appears. Use this field to associate an ACL for the match criteria. The <code>ACL Identifier</code> field has the following guidelines: <ul style="list-style-type: none"> <li>The drop-down menu lists the name or number that identifies the ACL for all configured ACLs that are valid for the class type and protocol.</li> <li>Standard and Extended IPv4 ACLs use numbers in the range 1 to 199. All other ACL types use names.</li> <li>If you select an IP ACL, you cannot select the No Protocol option to configure the Class as a non-IP L2 match DiffServ class.</li> </ul>
Reference Class	Select this option to reference another class for criteria. The match criteria defined in the referenced class is as match criteria in addition to the match criteria you define for the selected class. After selecting this option, the classes that can be referenced are displayed. Select the class to reference. A class can reference at most one other class of the same type.

Table 401: DiffServ Class Configuration Fields (Continued)

Field	Description
Class of Service	Select this option to require the Class of Service (CoS) value in an Ethernet frame header to match the specified CoS value.
Secondary Class of Service	Select this option to require the secondary CoS value in an Ethernet frame header to match the specified secondary CoS value.
Ethertype	Select this option to require the EtherType value in the Ethernet frame header to match the specified EtherType value. After you select this option, specify the EtherType value in one of the following two fields: <ul style="list-style-type: none"> <li>Ethertype Keyword – The menu includes several common protocols that are mapped to their EtherType values.</li> <li>Ethertype Value – This field accepts custom EtherType values.</li> </ul>
VLAN	Select this option to require a packet's VLAN ID to match a VLAN ID or a VLAN ID within a continuous range. If you configure a range, a match occurs if a packet's VLAN ID is the same as any VLAN ID within the range. After you select this option, use the following fields to configure the VLAN match criteria: <ul style="list-style-type: none"> <li>VLAN ID Start – The VLAN ID to match or the VLAN ID with the lowest value within a range of VLANs.</li> <li>VLAN ID End – The VLAN ID with the highest value within the range of VLANs. This field is not required if the match criteria is a single VLAN ID.</li> </ul>
Secondary VLAN	Select this option to require a packet's VLAN ID to match a secondary VLAN ID or a secondary VLAN ID within a continuous range. If you configure a range, a match occurs if a packet's secondary VLAN ID is the same as any secondary VLAN ID within the range. After you select this option, use the following fields to configure the secondary VLAN match criteria: <ul style="list-style-type: none"> <li>Secondary VLAN ID Start – The secondary VLAN ID to match or the secondary VLAN ID with the lowest value within a range of VLANs.</li> <li>Secondary VLAN ID End – The secondary VLAN ID with the highest value within the range of VLANs. This field is not required if the match criteria is a single VLAN ID.</li> </ul>
Source MAC Address	Select this option to require a packet's source MAC address to match the specified MAC address. After you select this option, use the following fields to configure the source MAC address match criteria: <ul style="list-style-type: none"> <li>MAC Address – The source MAC address to match.</li> <li>MAC Mask – The MAC mask, which specifies the bits in the source MAC address to compare against an Ethernet frame. Use F's and zeros to configure the MAC mask. An F means that the bit is checked, and a zero in a bit position means that the data is not significant. For example, if the MAC address is aa:bb:cc:dd:ee:ff, and the mask is ff:ff:00:00:00:00, all MAC addresses with aa:bb:xx:xx:xx:xx result in a match (where x is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.</li> </ul>
Destination MAC Address	Select this option to require a packet's destination MAC address to match the specified MAC address. After you select this option, use the following fields to configure the destination MAC address match criteria: <ul style="list-style-type: none"> <li>MAC Address – The destination MAC address to match.</li> <li>MAC Mask – The MAC mask, which specifies the bits in the destination MAC address to compare against an Ethernet frame. Use F's and zeros to configure the MAC mask. An F means that the bit is checked, and a zero in a bit position means that the data is not significant. For example, if the MAC address is aa:bb:cc:dd:ee:ff, and the mask is ff:ff:00:00:00:00, all MAC addresses with aa:bb:xx:xx:xx:xx result in a match (where x is any hexadecimal number). Note that this is not a wildcard mask, which ACLs use.</li> </ul>
Source IPv6 Address	Select this option to require the source IPv6 address in a packet header to match the specified values. After you select this option, use the following fields to configure the source IPv6 address match criteria: <ul style="list-style-type: none"> <li>Source Prefix – The source IPv6 prefix to match.</li> <li>Source Prefix Length – The IPv6 prefix length.</li> </ul>



Table 401: DiffServ Class Configuration Fields (Continued)

Field	Description
Destination IPv6 Address	<p>Select this option to require the destination IPv6 address in a packet header to match the specified values. After you select this option, use the following fields to configure the destination IPv6 address match criteria:</p> <ul style="list-style-type: none"> <li>• Destination Prefix – The destination IPv6 prefix to match.</li> <li>• Destination Prefix Length – The IPv6 prefix length.</li> </ul>
Source L4 Port	<p>Select this option to require a packet's TCP/UDP source port to match the specified port or the port number within a range of port numbers. If you configure a range, a match occurs if a packet's source port number is the same as any source port number within the range. After you select this option, use the following fields to configure a source port keyword, source port number, or source port range for the match criteria:</p> <ul style="list-style-type: none"> <li>• Protocol – Select the desired L4 keyword from the list on which the match is based. If you select a keyword, the other source port configuration fields are not available.</li> <li>• Port End – A user-defined L4 source port number to match or the source port number with the lowest value within a range of ports.</li> <li>• Port Start – The source port with the highest value within the range of ports. This field is not required if the match criteria is a single port.</li> </ul>
Destination L4 Port	<p>Select this option to require a packet's TCP/UDP destination port to match the specified port or the port number within a range of port numbers. If you configure a range, a match occurs if a packet's destination port number is the same as any destination port number within the range. After you select this option, use the following fields to configure a destination port keyword, destination port number, or destination port range for the match criteria:</p> <ul style="list-style-type: none"> <li>• Protocol – Select the desired L4 keyword from the list on which the match is based. If you select a keyword, the other destination port configuration fields are not available.</li> <li>• Port End – A user-defined L4 destination port number to match or the destination port number with the lowest value within a range of ports.</li> <li>• Port Start – The destination port with the highest value within the range of ports. This field is not required if the match criteria is a single port.</li> </ul>
IP DSCP	<p>Select this option to require the packet's IP DiffServ Code Point (DSCP) value to match the specified value. The DSCP value is defined as the high-order six bits of the Service Type octet in the IP header. After you select this option, use one of the following fields to configure the IP DSCP match criteria:</p> <ul style="list-style-type: none"> <li>• IP DSCP Keyword – The IP DSCP keyword code that corresponds to the IP DSCP value to match. If you select a keyword, you cannot configure an IP DSCP Value.</li> <li>• IP DSCP Value – The IP DSCP value to match.</li> </ul>
IP Precedence	<p>Select this option to require the packet's IP Precedence value to match the number configured in the IP Precedence Value field. The IP Precedence field in a packet is defined as the high-order three bits of the Service Type octet in the IP header.</p>
IP TOS	<p>Select this option to require the packet's Type of Service (ToS) bits in the IP header to match the specified value. The IP ToS field in a packet is defined as all eight bits of the Service Type octet in the IP header. After you select this option, use the following fields to configure the ToS match criteria:</p> <ul style="list-style-type: none"> <li>• IP TOS Bits – Enter a two-digit hexadecimal number to match the bits in a packet's ToS field.</li> <li>• IP TOS Mask – Specify the bit positions that are used for comparison against the IP ToS field in a packet.</li> </ul>

Table 401: DiffServ Class Configuration Fields (Continued)

Field	Description
Protocol	<p>Select this option to require a packet header's Layer 4 protocol to match the specified value. After you select this option, use one of the following fields to configure the protocol match criteria:</p> <ul style="list-style-type: none"> <li>No Protocol – A non-IP L2 match DiffServ class. If you select this option, you cannot select a protocol keyword or configure a protocol value.</li> <li>Protocol – The L4 keyword that corresponds to value of the IANA protocol number to match. If you select a keyword, you cannot configure a Protocol Value.</li> <li>Protocol Value – The IANA L4 protocol number value to match.</li> </ul>
Flow Label	<p>Select this option to require an IPv6 packet's flow label to match the configured value. The flow label is a 20-bit number that is unique to an IPv6 packet, used by end stations to signify quality-of-service handling in routers.</p>

- Click Refresh to update the page with the most current data from the switch.

### 8.4.3.1 DiffServ Add Match Criteria

After you click Add Match Criteria, the Add Match Criteria window opens and allows you to define the match criteria for the selected class. The window lists the match criteria that are available for the class.

Figure 426: DiffServ Add Match Criteria

The screenshot shows the 'Add Match Criteria' dialog box. It features a list of match criteria, each with a checkbox. The following criteria are checked: 'Reference ACL' and 'Protocol'. Under 'Reference ACL', the 'ACL Identifier' is set to '< None Available >'. Under 'Protocol', the 'No Protocol' checkbox is unchecked, the 'Protocol' dropdown is set to '< value >', and the 'Protocol Value' text box is empty with '(0 to 255)' next to it. At the bottom right, there are 'Submit' and 'Cancel' buttons.

To add match criteria, select the check box associated with the criteria type. The fields to configure the match values appear after you select the match type. Each match criteria type can be used only once within a class. If a reference class includes the match criteria type, it cannot be used as an additional match type within the class, and the match criteria type cannot be selected or configured.

**NOTICE**

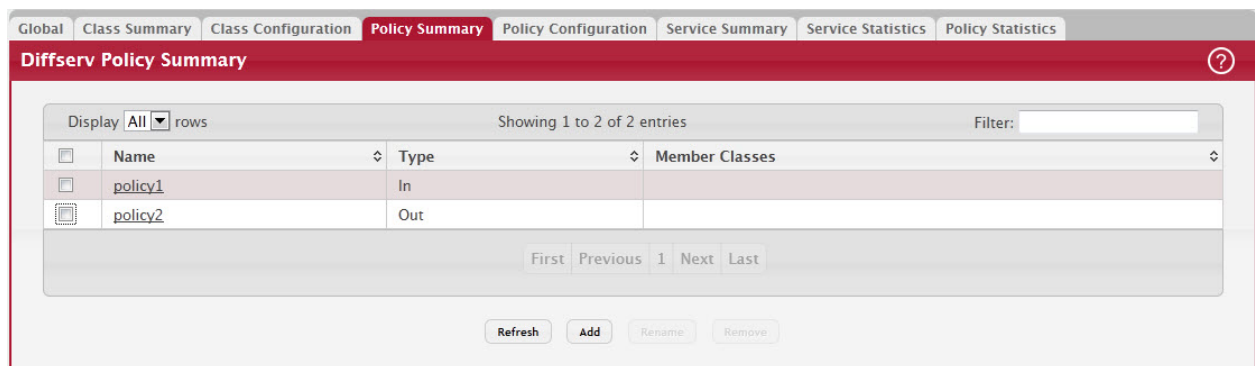
Each match type (other than Reference Class and Reference ACL) includes an option to match any value within the match criteria type except the configured value. This is the Exclude option, which indicates a logical NOT for a match criteria type.

### 8.4.4 DiffServ Policy Summary

Use this page to create or remove DiffServ policies and to view summary information about the policies that exist on the device. A policy defines the QoS attributes for one or more traffic classes. A policy attribute identifies the action taken when a packet matches a class rule. A policy is applied to a packet when a class match within that policy is found.

To display the DiffServ Policy Summary page, click QoS > DiffServ > Policy Summary in the navigation menu.

Figure 427: DiffServ Policy Summary



Use the buttons to perform the following tasks:

- To add a DiffServ policy, click Add.
- To change the name of an existing policy, select the entry to modify and click Rename.
- To remove one or more configured policies, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 402: DiffServ Policy Summary Fields

Field	Description
Name	The name of the DiffServ policy. When adding a new policy or renaming an existing policy, the name of the policy is specified in the Policy field of the dialog window.
Type	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> <li>• In – The policy is specific to inbound traffic.</li> <li>• Out – The policy is specific to outbound traffic direction.</li> </ul>
Member Classes	The DiffServ class or classes that have been added to the policy.

- Click Refresh to update the page with the most current data from the switch.

### 8.4.5 DiffServ Policy Configuration

Use this page to add or remove a DiffServ policy-class association and to configure the policy attributes. The policy attributes identify the action or actions taken when a packet matches a class rule.

After you select the policy to configure from the Policy menu, use the buttons to perform the following tasks:

- To add a class to the policy, click Add Class.

- To add attributes to a policy or to change the policy attributes, select the policy with the attributes to configure and click Add Attribute.
- To remove the most recently associated class from the selected policy, click Remove Last Class.
- To display the DiffServ Policy Configuration page, click QoS > DiffServ > Policy Configuration in the navigation menu.

Figure 428: DiffServ Policy Configuration

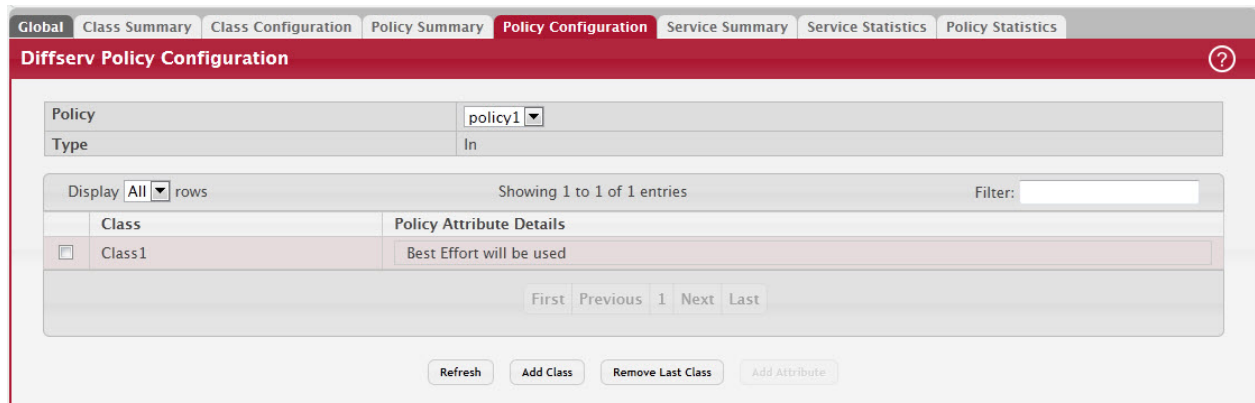


Table 403: DiffServ Policy Configuration Fields

Field	Description
Policy	The name of the policy. To add a class to the policy, remove a class from the policy, or configure the policy attributes, you must first select its name from the menu.
Type	The traffic flow direction to which the policy is applied.
Class	The DiffServ class or classes associated with the policy. The policy is applied to a packet when a class match within that policy-class is found.
Policy Attribute Details	The policy attribute types and their associated values that are configured for the policy.
Assign Queue	Select this option to assign matching packets to a traffic queue. Use the Queue ID Value field to select the queue to which the packets of this policy-class are assigned.
Drop	Select this option to drop packets that match the policy-class.
Mark CoS	Select this option to mark all packets in a traffic stream with the specified Class of Service (CoS) queue value. Use the Class of Service field to select the CoS value to mark in the priority field of the 802.1p header (the only tag in a single tagged packet or the first or outer 802.1Q tag of a double VLAN tagged packet). If the packet does not already contain this header, one is inserted.
Mark Secondary CoS	Select this option to mark all packets in a traffic stream with the specified secondary CoS queue number. Use the Class of Service field to select the CoS value to mark in the priority field of the 802.1p header in the secondary (inner) 802.1Q tag of a double VLAN tagged packet. If the packet does not already contain this header, one is inserted.
Mark IP DSCP	Select this option to mark all packets in the associated traffic stream with the specified IP DSCP value. After you select this option, use one of the following fields to configure the IP DSCP value to mark in packets that match the policy-class: <ul style="list-style-type: none"> <li>• IP DSCP Keyword – The IP DSCP keyword code that corresponds to the IP DSCP value. If you select a keyword, you cannot configure an IP DSCP Value.</li> <li>• IP DSCP Value – The IP DSCP value.</li> </ul>
Mark IP Precedence	Select this option to mark all packets in the associated traffic stream with the specified IP Precedence value. After you select this option, use the IP Precedence Value field to select the IP Precedence value to mark in packets that match the policy-class.

Table 403: DiffServ Policy Configuration Fields (Continued)

Field	Description
Mirror Interface	Select this option to copy the traffic stream to a specified egress port (physical or LAG) without bypassing normal packet forwarding. This action can occur in addition to any marking or policing action. It may also be specified along with a QoS queue assignment. Use the Interface menu to select the interface to which traffic is mirrored.
Police Simple	Select this option to enable the simple traffic policing style for the policy-class. The simple form of the police attribute uses a single data rate and burst size, resulting in two outcomes (conform and violate). After you select this option, configure the following policing criteria: <ul style="list-style-type: none"> <li>• Color Mode – The type of color policing used in DiffServ traffic conditioning.</li> <li>• Color Conform Class – For color-aware policing, packets in this class are metered against both the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: CoS, IP DSCP, IP Precedence, or Secondary COS.</li> <li>• Committed Rate (Kbps) – The maximum allowed arrival rate of incoming packets for this class.</li> <li>• Committed Burst Size (Kbytes) – The amount of conforming traffic allowed in a burst.</li> <li>• Conform Action – The action taken on packets that are considered conforming (below the police rate).</li> <li>• Violate Action – The action taken on packets that are considered non-conforming (above the police rate).</li> </ul>
Police Single Rate	Select this option to enable the single-rate traffic policing style for the policy-class. The single-rate form of the police attribute uses a single data rate and two burst sizes, resulting in three outcomes (conform, exceed, and violate). After you select this option, configure the following policing criteria: <ul style="list-style-type: none"> <li>• Color Mode – The type of color policing used in DiffServ traffic conditioning.</li> <li>• Color Conform Class – For color-aware policing, packets are metered against the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: CoS, IP DSCP, IP Precedence, or Secondary COS. This field is available only if one or more classes that meets the color-awareness criteria exist.</li> <li>• Color Exceed Class – For color-aware policing, packets are metered against the PIR only.</li> <li>• Committed Rate (Kbps) – The maximum allowed arrival rate of incoming packets for this class.</li> <li>• Committed Burst Size (Kbytes) – The amount of conforming traffic allowed in a burst.</li> <li>• Excess Burst Size (Kbytes) – The amount of conforming traffic allowed to accumulate beyond the Committed Burst Size (Kbytes) value during longer-than-normal idle times. This value allows for occasional bursting.</li> <li>• Conform Action – The action taken on packets that are considered conforming (below the police rate).</li> <li>• Exceed Action – The action taken on packets that are considered to exceed the committed burst size but are within the excessive burst size.</li> <li>• Violate Action – The action taken on packets that are considered non-conforming (above the police rate).</li> </ul>

Table 403: DiffServ Policy Configuration Fields (Continued)

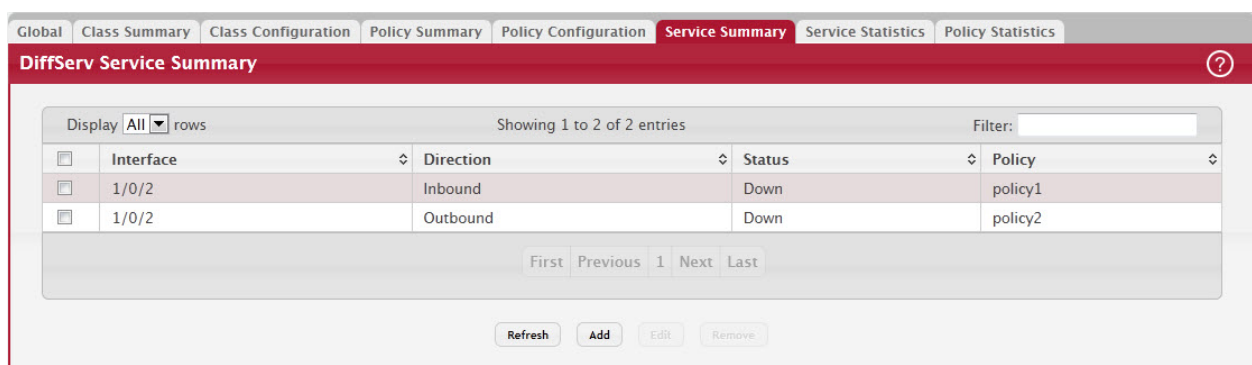
Field	Description
Police Two Rate	<p>Select this option to enable the two-rate traffic policing style for the policy-class. The two-rate form of the police attribute uses two data rates and two burst sizes. Only the smaller of the two data rates is intended to be guaranteed. After you select this option, configure the following policing criteria:</p> <ul style="list-style-type: none"> <li>• Color Mode – The type of color policing used in DiffServ traffic conditioning.</li> <li>• Color Conform Class – For color-aware policing, packets are metered against the committed information rate (CIR) and the peak information rate (PIR). The class definition used for policing color awareness is only allowed to contain a single, non-excluded class match condition identifying one of the supported comparison fields: CoS, IP DSCP, IP Precedence, or Secondary COS. This field is available only if one or more classes that meets the color-awareness criteria exist.</li> <li>• Color Exceed Class – For color-aware policing, packets are metered against the PIR.</li> <li>• Committed Rate (Kbps) – The maximum allowed arrival rate of incoming packets for this class.</li> <li>• Committed Burst Size (Kbytes) – The amount of conforming traffic allowed in a burst.</li> <li>• Peak Rate (Kbps) – The maximum peak information rate for the arrival of incoming packets for this class.</li> <li>• Excess Burst Size (Kbytes) – The maximum size of the packet burst that can be accepted to maintain the Peak Rate (Kbps).</li> <li>• Conform Action – The action taken on packets that are considered conforming (below the police rate).</li> <li>• Exceed Action – The action taken on packets that are considered to exceed the committed burst size but are within the excessive burst size.</li> <li>• Violate Action – The action taken on packets that are considered non-conforming (above the police rate).</li> </ul>
Redirect Interface	<p>Select this option to force a classified traffic stream to the specified egress port (physical port or LAG). Use the Interface field to select the interface to which traffic is redirected.</p>

- Click Refresh to update the page with the most current data from the switch.

### 8.4.6 DiffServ Service Summary

Use this page to add DiffServ policies to interfaces, remove policies from interfaces, and edit policy-interface mappings. To display the DiffServ Service Summary page, click QoS > DiffServ > Service Summary in the navigation menu.

Figure 429: DiffServ Service Summary



Use the buttons to perform the following tasks:

- To add a policy to an interface, click Add.
- To edit a configured interface-policy association, select the entry to modify and click Edit.
- To remove one or more configured interface-policy associations, select each entry to delete and click Remove. You must confirm the action before the entry is deleted.

Table 404: DiffServ Service Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. Only interfaces that have an associated policy are listed in the table.
Direction	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> <li>Inbound – The policy is applied to traffic as it enters the interface.</li> <li>Outbound – The policy is applied to traffic as it exits the interface.</li> </ul>
Status	The status of the policy on the interface. A policy is Up if DiffServ is globally enabled, and if the interface is administratively enabled and has a link. Otherwise, the status is Down.
Policy	The DiffServ policy associated with the interface.
When you click Add or Edit, the Configure Service window opens and allows you to configure DiffServ interface policies. Specifying 'None' for a policy has no effect when adding or editing interface policies. To remove an interface policy mapping, use the Remove button on the parent page. The following information describes the fields in this window.	
Interface	Select an interface to associate with a policy.
Policy In	The menu lists all policies configured with a type of In. Select the policy to apply to traffic as it enters the interface.
Policy Out	The menu lists all policies configured with a type of Out. Select the policy to apply to traffic as it exits the interface.

- Click Refresh to update the page with the most current data from the switch.

### 8.4.7 DiffServ Service Statistics

This page displays service-level statistical information for all interfaces in the system to which a DiffServ policy has been attached.

To display the DiffServ Service Statistics page, click QoS > DiffServ > Service Statistics in the navigation menu.

Figure 430: DiffServ Service Performance Statistics

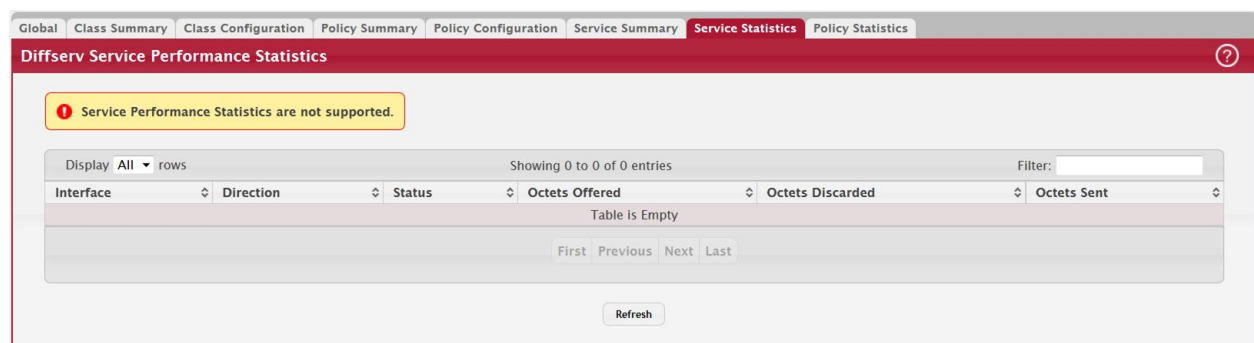


Table 405: DiffServ Service Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. The table displays all interfaces that have a DiffServ policy currently attached in a traffic flow direction.
Direction	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> <li>In – The policy is applied to traffic as it enters the interface.</li> <li>Out – The policy is applied to traffic as it exits the interface.</li> </ul>
Status	The operational status of this service interface, either Up or Down.

Table 405: DiffServ Service Statistics Fields (Continued)

Field	Description
Octets Offered	The total number of octets offered to all class instances in this service policy before their defined DiffServ treatment is applied. This is the overall count per-interface, per-direction.
Octets Discarded	The total number of octets discarded for all class instances in this service policy for any reason due to DiffServ treatment. This is the overall count per-interface, per-direction.
Octets Sent	The total number of octets forwarded for all class instances in this service policy after their defined DiffServ treatments were applied. In this case, forwarding means the traffic stream was passed to the next functional element in the data path, such as the switching or routing function of an outbound link transmission element. This is the overall count per-interface, per-direction.

- Click Refresh to update the page with the most current data from the switch.

### 8.4.8 DiffServ Service Policy Statistics

This page displays class-oriented statistical information for the policy, which is specified by the interface and direction. To display the DiffServ Service Policy Statistics page, click QoS > DiffServ > Policy Statistics in the navigation menu.

Figure 431: DiffServ Service Policy Statistics

Table 406: DiffServ Service Policy Statistics Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. The table displays all interfaces that have a DiffServ policy currently attached in a traffic flow direction.
Direction	The traffic flow direction to which the policy is applied: <ul style="list-style-type: none"> <li>In – The policy is applied to traffic as it enters the interface.</li> <li>Out – The policy is applied to traffic as it exits the interface.</li> </ul>
Policy	The name of the policy currently attached to the interface.
Status	The operational status of the policy currently attached to the interface.
Class	The DiffServ class currently defined for the attached policy.
Octets Offered	The total number of octets offered to all class instances in this service policy before their defined DiffServ treatment is applied. This is the overall count per-interface, per-direction.
Octets Discarded	The total number of octets discarded for all class instances in this service policy for any reason due to DiffServ treatment. This is the overall count per-interface, per-direction.



**Table 406: DiffServ Service Policy Statistics Fields (Continued)**

Field	Description
Packets Offered	The total number of packets offered to all class instances in this service policy before their defined DiffServ treatment is applied. This is the overall count per-interface, per-direction.
Packets Discarded	The total number of packets discarded for all class instances in this service policy for any reason due to DiffServ treatment. This is the overall count per-interface, per-direction.

Click Refresh to update the page with the most current data from the switch.

## 9/ Appendix: Configuration Examples

This appendix contains examples of how to configure selected features available in the FASTPATH software. Each example contains procedures on how to configure the feature by using the Web interface, and/or CLI, and/or SNMP.

Each configuration example starts from a factory-default configuration unless otherwise noted.


### NOTICE

### 9.1 Configuring VLANs

The diagram in this section shows a switch with four ports configured to handle the traffic for two VLANs. Port 1/0/2 handles traffic for both VLANs, while port 1/0/1 is a member of VLAN 2 only, and ports 1/0/3 and 1/0/4 are members of VLAN 3 only.

The following examples show how to create VLANs, assign ports to the VLANs, and assign a VLAN as the default VLAN to a port.

Figure 432: VLAN Example Network Diagram



#### 9.1.1 Using the Web Interface to Configure VLANs

1. Access the Switching > VLAN > Status page.
2. Click Add to create a new VLAN.
3. Type 2-3 in the VLAN ID-Individual/Range field.

Add VLAN
✕

VLAN ID or Range	<input style="width: 90%;" type="text"/> <small>Enter VLAN ID in the range 2 to 4093. Use '-' to specify a range and ',' to separate VLAN IDs or VLAN ranges in the list.</small>
------------------	--

4. Click Submit.

5. From the Port Configuration Page, Select VLAN 2 from the VLAN ID List.
6. From the Participation column in the interface table, select Include for ports 1/0/1 and 1/0/2 to specify that these ports are members of VLAN 2.
7. Select the interface check box and click Edit. Select the Tagging All check box to specify that frames will always be transmitted tagged from ports that are members of VLAN 2.
8. Click Submit.
9. Select VLAN 3 from the VLAN ID and Name List.
10. Select the Participate option in the VLAN field.
11. For ports 1/0/2, 1/0/3 and 1/0/4, select Include from the Participation menu to specify that these ports are members of VLAN 3.
12. Click Submit.
13. Go to the Switching > VLAN > Port Configuration page.
14. From the Interface menu, select 1/0/1.
15. In the Acceptable Frame Types field, select AdmitTaggedOnly to specify that untagged frames will be rejected on receipt.
16. Click Submit.
17. From the Interface menu, select 1/0/2.
18. In the Port VLAN ID field, enter 3 to assign VLAN 3 as the default VLAN for the port.
19. In the Acceptable Frame Types field, select AdmitTaggedOnly to specify that untagged frames will be rejected on receipt.

The screenshot displays the 'VLAN Port Configuration' page. At the top, there are tabs for 'Status', 'Port Configuration', 'Port Summary', 'Switchport Summary', 'Reset', and 'RSPAN'. The main title is 'VLAN Port Configuration'. Below the title, there is a 'VLAN ID' dropdown menu set to '1'. A table shows the configuration for 10 interfaces. The table has columns for 'Interface', 'Status', 'Participation', and 'Tagging'. All interfaces are listed with 'Include' status and 'Include' participation, and 'Untagged' tagging. Below the table, there are navigation buttons: 'First', 'Previous', '1', '2', '3', 'Next', and 'Last'. At the bottom of the page, there are buttons for 'Refresh', 'Edit', and 'Edit All'. The footer of the page reads '© Broadcom Corporation 2000-2015'.

Interface	Status	Participation	Tagging
1	Include	Include	Untagged
2	Include	Include	Untagged
3	Include	Include	Untagged
4	Include	Include	Untagged
5	Include	Include	Untagged
6	Include	Include	Untagged
7	Include	Include	Untagged
8	Include	Include	Untagged
9	Include	Include	Untagged
10	Include	Include	Untagged

20. Click Submit.

## 9.1.2 Using the CLI to Configure VLANs

### 1. Create VLAN 2 and VLAN 3.

```
(Broadcom FASTPATH Routing) #vlan database
vlan 2
vlan 3
exit
```

### 2. Assign ports 1/0/1 and 1/0/2 to VLAN2 and specify that untagged frames will be rejected on receipt.

```
(Broadcom FASTPATH Routing) #Config
interface 1/0/1
vlan participation include 2
vlan acceptframe vlanonly
exit
interface 1/0/2
vlan participation include 2
vlan acceptframe vlanonly
```

### 3. While in interface config mode for port 1/0/2, assign VLAN3 as the default VLAN.

```
(Broadcom FASTPATH Routing) (Interface 1/0/2)#vlan pvid 3
exit
```

### 4. Specify that frames will always be transmitted tagged from ports that are members of VLAN 2.

```
(Broadcom FASTPATH Routing) (Config)#vlan port tagging all 2
exit
```

### 5. Assign the ports that will belong to VLAN 3.

---

Port 1/0/2 belongs to both VLANs, and port 1/0/1 can never belong to VLAN 3.

#### **NOTICE**

```
(Broadcom FASTPATH Routing) #Config
interface 1/0/2
vlan participation include 3
exit
interface 1/0/3
vlan participation include 3
exit
interface 1/0/4
vlan participation include 3
exit
exit
```

### 6. Specify that untagged frames will be accepted on port 1/0/4.

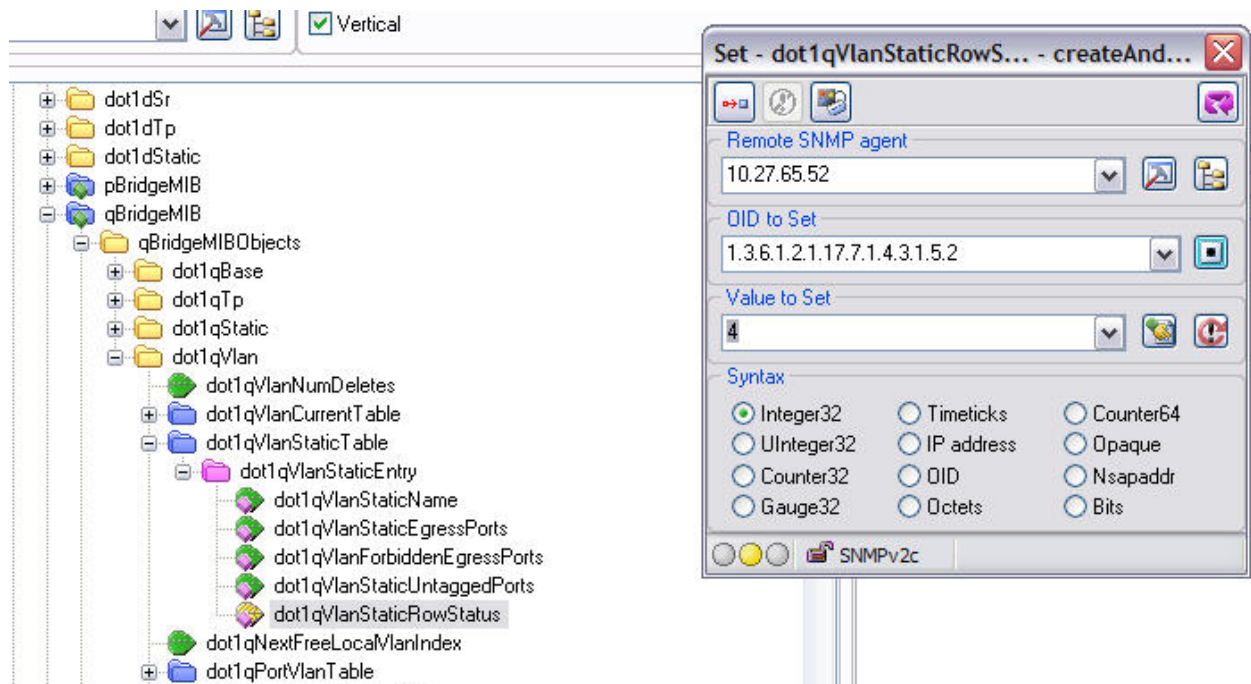
```
(Broadcom FASTPATH Routing) #Config
interface 1/0/4
vlan acceptframe all
exit
exit
```

## 9.1.3 Using the SNMP to Configure VLANs

### 1. Use the objects in dot1qVlanStaticTable (in dot1qVlan in the QBRIDGE-MIB module) to create VLANs 2 and 3.

Set the dot1qVlanStaticRowStatus object to 'CreateandGo (4)' to create a VLAN. If the other parameters are not specified, simply specifying the dot1qVlanIndex and dot1qVlanStaticRowStatus is sufficient to create the VLAN.

The full path to the object is iso(1).org(3).dod(6).internet(1).mgmt(2).mib-2(1).dot1dBridge(17).qBridgeMIB(7).qBridgeMIBObjects(1).dot1qVlan(4).dot1qVlanStaticTable(3).dot1qVlanStaticEntry(1).dot1qVlanStaticRowStatus(5).



- To assign ports 1/0/1 and 1/0/2 to VLAN2, retrieve the current `dot1qStaticEgressPorts` mask and append interfaces 1/0/1 and 1/0/2 to this mask by setting the first octet to 0xC0.

The `dot1qVlanStaticEgressPorts` bit mask can be constructed according to the following rules:

- Each octet within this value specifies a set of eight ports, with the first octet specifying ports (1-8), the second octet specifying ports (9-16), and so on.
- Within each octet, the most significant bit represents the lowest numbered port, and the least significant bit represents the highest numbered port. Thus, each port of the bridge is represented by a single bit within the value of this object. If that bit has a value of (1), then that port is included in the set of ports. The port is not included if its bit has a value of (0).

For example if the switch has 12 ports and we want to add ports 1 and 4 in the VLAN and exclude all other ports, then the bit mask in hex will be 0x50 0x00.

- To specify that frames will always be transmitted tagged from ports that are members of VLAN 2, use the `dot1qVlanStaticUntaggedPorts` object and set the value of the appropriate number of octets to 0.

Each octet represents eight ports, so for a 48-port switch, the first six octets would be zero.

- To specify that ports 1/0/1 and 1/0/2 will only accept tagged frames and will reject untagged frames on receipt, set the `dot1qPortAcceptableFrameTypes` object to `admitOnlyVlanTagged(2)`.

The object is in `dot1qPortVlanEntry` in the `dot1qPortVlanTable`.

- To assign VLAN3 as the default VLAN for interface 1/0/2, set the value of `dot1qPvid` for 1/0/2 (instance 2) to 3.
- To assign ports 1/0/2, 1/0/3, and 1/0/4 to VLAN3, retrieve the current `dot1qStaticEgressPorts` mask and append the interfaces to this mask by setting the first octet to 0x70.

## 9.2 Configuring Multiple Spanning Tree Protocol

This example shows how to enable IEEE 802.1s Multiple Spanning Tree (MST) protocol on the switch and all of the ports and to set the bridge priority.

To make multiple switches be part of the same MSTP region, make sure the Force Protocol Version setting for all switches is IEEE 802.1s. Also, make sure the configuration name, digest key, and revision level are the same for all switches in the region.

### NOTICE

The digest key is generated based on the association of VLANs to different instances. To ensure the digest key is same, the mapping of VLAN to instance must be the same on each switch in the region. For example, if VLAN 10 is associated with instance 10 on one switch, you must associate VLAN 10 and instance 10 on the other switches.

### 9.2.1 Using the Web UI to Configure MSTP

1. Create VLANs 10 and 20.
  - a. Access the Switching > VLAN > Status page.
  - b. Click Add to create a VLAN.
  - c. Select the VLAN ID-Individual option and enter 10.
  - d. Click Submit.
  - e. Repeat the steps to add VLAN 20.
2. Enable MSTP (IEEE 802.1s) on the switch and change the configuration name.

Changing the configuration name allows all the bridges that want to be part of the same region to join.

- a. Go to the Switching > Spanning Tree > Switch page.
- b. From the Spanning Tree Admin Mode menu, select Enable.
- c. In the Configuration Name field, enter `broadcom`.
- d. Click Submit.

3. Create two MST instances.
  - a. Go to the Switching > Spanning Tree > MST page.
  - b. From the MST page, click Add.
  - c. In the MST ID field, enter 10.
  - d. Associate MST ID 10 with VLAN 10 and assign a bridge priority of 16384.
  - e. Click Submit.
  - f. Repeat the steps to create an MST instance with an ID of 20.

4. Use similar procedures to associate MST instance 20 to VLAN 20 and assign it a bridge priority value of 61440. By using a lower priority for MST 20, MST 10 becomes the root bridge.
5. Force port 1/0/2 to be the root port for MST 20, which is the non-root bridge.
  - a. Go to the Switching > Spanning Tree > MST page.
  - b. From the MST ID menu, select 20.
  - c. From the Interface menu, select 1/0/2.
  - d. In the Port Priority field, enter 64.
  - e. Click Submit.

## 9.2.2 Using the CLI to Configure MSTP

1. Create VLAN 10 and VLAN 20.
 

```
(Broadcom FASTPATH Routing) #vlan database
vlan 10
vlan 20
exit
```
2. Enable spanning tree Globally
 

```
(Broadcom FASTPATH Routing) #config
spanning-tree
```
3. Create MST instances 10 and 20.
 

```
spanning-tree mst instance 10
spanning-tree mst instance 20
```
4. Associate MST instance 10 to VLAN 10 and MST instance 20 to VLAN 20
 

```
spanning-tree mst vlan 10 10
spanning tree mst vlan 20 20
```
5. Change the name so that all the bridges that want to be part of the same region can form the region.
 

```
spanning-tree configuration name broadcom
```
6. Make the MST ID 10 bridge the root bridge by lowering the priority.
 

```
spanning-tree mst priority 10 16384
```
7. Change the priority of MST ID 20 to ensure the other bridge is the root bridge.
 

```
spanning-tree mst priority 20 61440
```
8. Enable STP on interface 1/0/1
 

```
interface 1/0/1
spanning-tree port mode
exit
```
9. Enable STP on interface 1/0/2
 

```
interface 1/0/2
spanning-tree port mode
```
10. On the non-root bridge change the priority to force port 1/0/2 to be the root port.

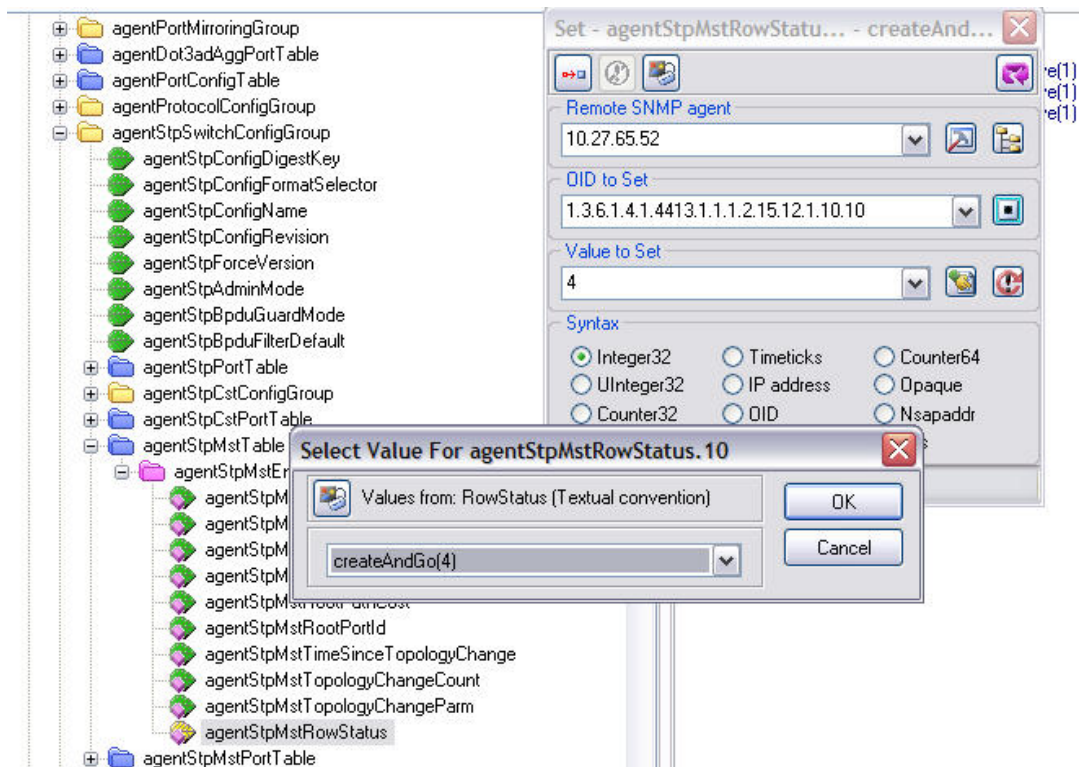
```
spanning-tree mst 20 port-priority 64
exit
```

### 9.2.3 Using SNMP to Configure MSTP

1. Use the objects in dot1qVlanStaticTable (in dot1qVlan in the QBRIDGE-MIB module) to create VLANs 10 and 20.
2. To enable spanning tree globally, set the agentStpAdminMode object in the FASTPATH-SWITCHING-MIB module to enable (2).

The full path to the object is iso(1).org(3).dod(6).internet(1).private(4).enterprises(1).broadcom(4413).broadcomProducts(1).fastPath(1).fastPathSwitching(1).agentConfigGroup(2).agentStpSwitchConfigGroup(15).agentStpAdminMode(6).

3. Use the agentStpConfigName object in the agentStpSwitchConfigGroup to change the name so that all the bridges that want to be part of the same region can form the region.
4. Use the agentStpMstRowStatus object in the agentStpMstTable to create MST instances 10 and 20.



5. Use the agentStpMstBridgePriority object to set the bridge priorities for MST 10 and MST 20:
  - For MST ID 10, set the value to 16384 to make it the root bridge.
  - For MST ID 20, set the value to 61440 to ensure the other bridge is the root bridge.
6. Use the agentStpMstVlanRowStatusAssociate object in the agentStpMstVlanTable to associate MST instance 10 to VLAN 10 and MST instance 20 to VLAN 20.
  - For MST ID 20, the OID to set is 1.3.6.1.4.1.4413.1.1.1.2.15.14.1.1.10.10 (the final .10 is the VLAN ID)
  - For MST ID 20, the OID to set is 1.3.6.1.4.1.4413.1.1.1.2.15.14.1.1.20.20

Set the value to CreateAndGo (4)

7. Use the agentStpPortState in agentStpPortTable under agentStpSwitchConfigGroup to enable STP on interface 1/0/1 and interface 1/0/2.

For instance 1 and 2, set the value to enable (1).

8. Use the agentStpMstPortPriority object in agentStpMstPortTable to change the port priority on interface 1/0/2 to force the port to be the root port on the non-root bridge.

For instance 2, set the value to 64.




## 9.3 Configuring VLAN Routing

This section provides an example of how to configure FASTPATH software to support VLAN routing. The configuration of the VLAN router port is similar to that of a physical port. The main difference is that, after the VLAN has been created, you must use the `show ip vlan` command to determine the VLAN's interface ID so that you can use it in the router configuration commands.

The diagram in this section shows a Layer 3 switch configured for port routing. It connects two VLANs, with two ports participating in one VLAN, and one port in the other. The script shows the commands you would use to configure FASTPATH software to provide the VLAN routing support shown in the diagram.

Figure 433: VLAN Routing Example Network Diagram



### 9.3.1 Using the CLI to Configure VLAN Routing

1. Create VLAN 10 and VLAN 20.

```
(Broadcom FASTPATH Routing) #vlan database
vlan 10
vlan 20
exit
```

2. Configure ports 1/0/1, 1/0/2 as members of VLAN 10 and specify that untagged frames received on these ports will be assigned to VLAN 10.

```
config
interface 1/0/1
vlan participation include 10
vlan pvid 10
exit
interface 1/0/2
vlan participation include 10
vlan pvid 10
exit
```

3. Configure port 1/0/3 as a member of VLAN 20 and specify that untagged frames received on these ports will be assigned to VLAN 20

```
interface 1/0/3
vlan participation include 20
vlan pvid 20
exit
exit
```

4. Specify that all frames transmitted for VLANs 10 and 20 will be tagged.

```
config
  vlan port tagging all 10
  vlan port tagging all 20
  exit
```

5. Enable routing for the VLANs:

```
(Broadcom FASTPATH Routing) #vlan database
  vlan routing 10
  vlan routing 20
  exit
```

6. View the logical interface IDs assigned to the VLAN routing interfaces.

```
(Broadcom FASTPATH Routing) #show ip vlan
```

MAC Address used by Routing VLANs: 00:00:AA:12:65:12

VLAN ID	Logical Interface	IP Address	Subnet Mask
10	0/4/1	0.0.0.0	0.0.0.0
20	0/4/2	0.0.0.0	0.0.0.0

As the output shows, VLAN 10 is assigned ID 0/4/1 and VLAN 20 is assigned ID 0/4/2

7. Enable routing for the switch:

```
config
  ip routing
  exit
```

8. Configure the IP addresses and subnet masks for the virtual router ports.

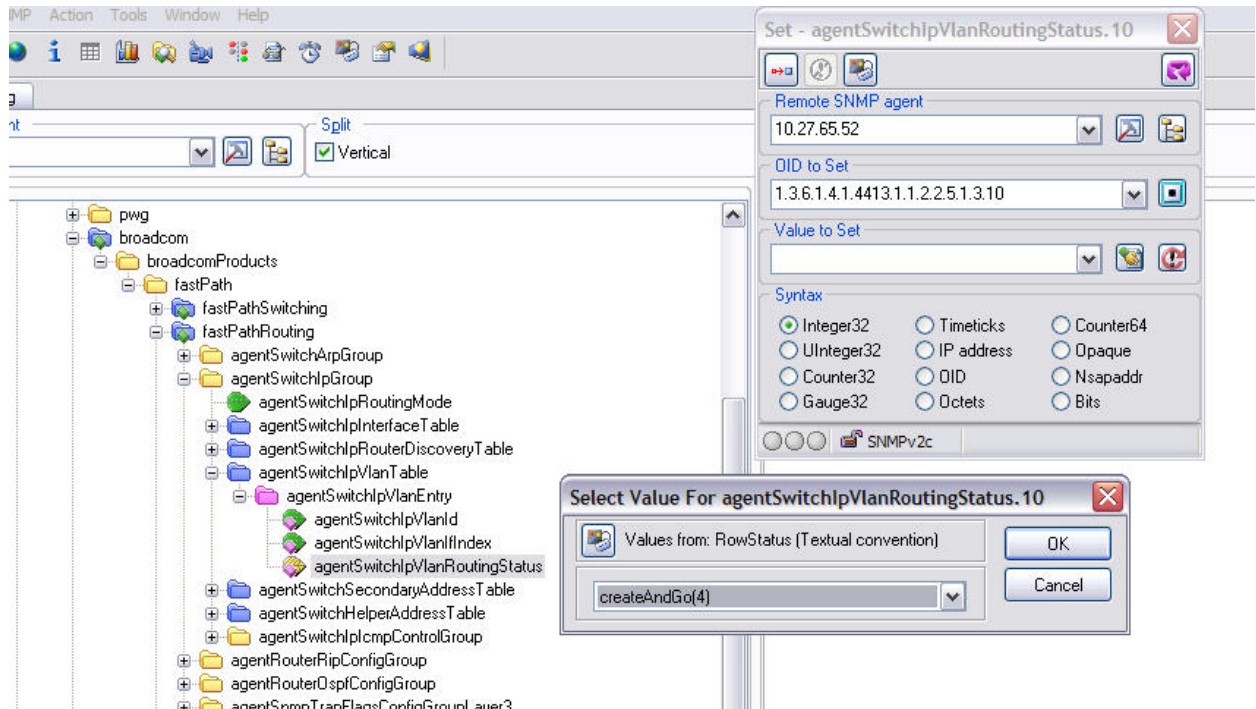
```
config
  interface 0/4/1
    ip address 192.150.3.1 255.255.255.0
  exit
  interface 0/4/2
    ip address 192.150.4.1 255.255.255.0
  exit
  exit
```

### 9.3.2 Using SNMP to Configure VLAN Routing

- Use the `dot1qVlanStaticRowStatus` object in the `dot1qVlanStaticTable` to create VLAN 10 and VLAN 20.
- To configure VLAN membership, retrieve the current `dot1qStaticEgressPorts` mask and append the desired interfaces to the mask.
  - VLAN 10: 1/0/1 and 1/0/2
  - VLAN 20: 1/0/3
- To assign the PVID for an interface, use the `dot1qPvid` object.
  - 1/0/1: PVID 10
  - 1/0/2: PVID 10
  - 1/0/3: PVID 20
- To specify that all frames transmitted for VLANs 10 and 20 will be tagged, use the `dot1qVlanStaticUntaggedPorts` object and set the value of the appropriate number of octets to 0.

Each octet represents eight ports, so for a 48-port switch, the first six octets would be zero.

- To enable routing for the VLANs, use the `agentSwitchIpVlanRoutingStatus` object in the `agentSwitchIpVlanTable` under `agentSwitchIpGroup` in `fastPathRouting` to set the value for VLAN 10 and VLAN 20 to `CreateAndGo (4)`.



6. Walk the agentSwitchIpVlanIndex object to view the logical interface IDs assigned to the VLAN routing interfaces.
7. Set the agentSwitchIpRoutingMode object to enable (1) to enable routing for the switch:
8. Use the agentSwitchIpInterfaceIpAddress and agentSwitchIpInterfaceIpMask objects in the agentSwitchIpInterfaceTable to configure the IP addresses and subnet mask for the virtual router ports.

### NOTICE

While setting the ip address for the VLAN interface, the agentSwitchIpInterfaceIpAddress and agentSwitchIpInterfaceNetMask should be set together.

- VLAN index 482 (VLAN 10): 192.150.3.1 255.255.255.0
- VLAN index 483 (VLAN 20): 192.150.4.1 255.255.255.0

## 9.4 Configuring Policy Based Routing

In contemporary networks, network administrators who manage organizations should be provided with a choice for implementing packet forwarding/routing according to the organization's policies. Policy Based Routing (PBR) is a feature that fits this purpose. PBR provides a flexible mechanism to implement solutions in cases where organizational constraints dictate that traffic be routed through specific network paths.

Configuring PBR involves configuring a route-map with match and set commands and then applying the corresponding route-map to the interface.

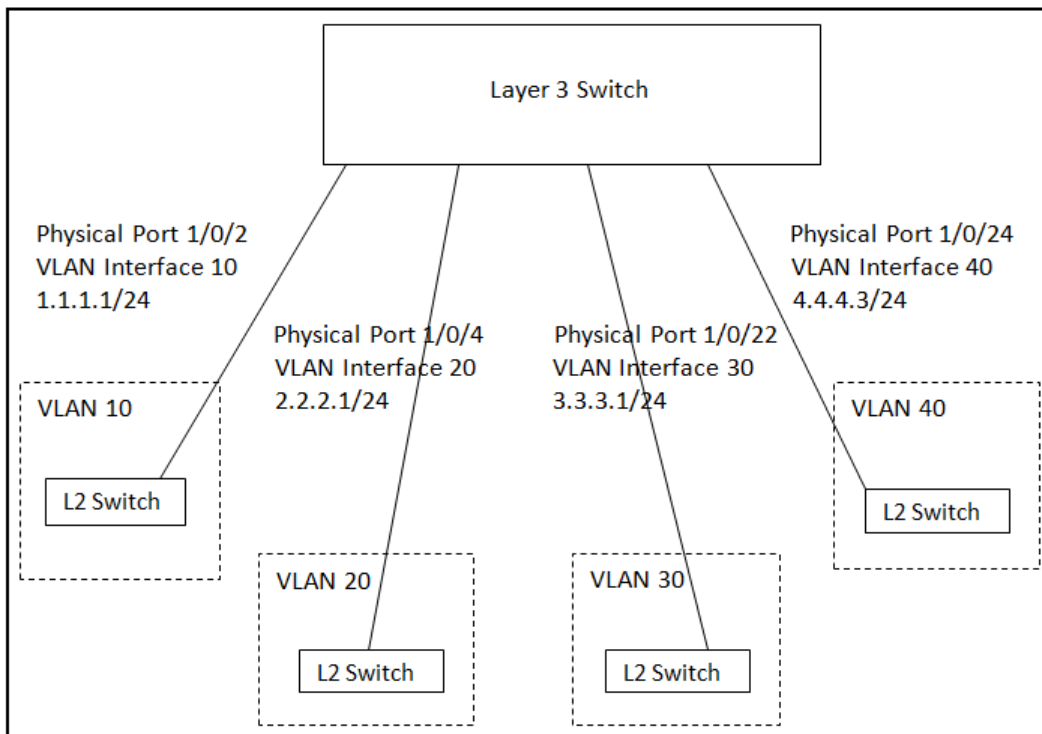
Policy Based Routing is applied to inbound traffic on physical routing/VLAN routing interfaces. Enabling the feature causes the router to analyze all packets incoming on the interface using a route-map configured for that purpose. One interface can only have one route-map tag, but an administrator can have multiple route-map entries with different sequence numbers. These entries are evaluated in sequence number order until the first match. If there is no match, the packets are routed as usual.

## 9.4.1 Configuring Policy Based Routing Using the CLI

In the following configuration example, we have a Layer3 Switch/Router with four VLAN routing interfaces – VLAN 10, VLAN 20, VLAN 30, and VLAN 40. Each of these interfaces is connected to an L2 network.

- Physical Interface 1/0/2 – Member of VLAN 10
- Physical Interface 1/0/4 – Member of VLAN 20
- Physical Interface 1/0/22 – Member of VLAN 30
- Physical Interface 1/0/24 – Member of VLAN 40

Figure 434: Policy Based Routing Example



In this example, the procedure to configure policy route traffic from VLAN routing interface 10 to VLAN routing interface 30 is shown in [Figure 434: "Policy Based Routing Example," on page 502](#). Traffic sent to VLAN Interface 10 is destined for VLAN Interface 20. In order to override the traditional destination routing and send the same traffic to VLAN Interface 30, use the following procedure.

1. Create VLANs 10, 20, 30, 40, and enable routing on these VLANs.

```
(Broadcom FASTPATH Routing) #vlan database
vlan 10,20,30,40
vlan routing 10 1
vlan routing 20 2
vlan routing 30 3
vlan routing 40 4
exit
```

2. Add physical ports to the VLANs and configure PVID on the corresponding interfaces.

```
config
interface 1/0/2
vlan pvid 10
vlan participation exclude 1
vlan participation include 10
exit
interface 1/0/4
vlan pvid 20
```

```

        vlan participation exclude 1
        vlan participation include 20
    exit
interface 1/0/22
    vlan pvid 30
    vlan participation exclude 1
    vlan participation include 30
    exit
interface 1/0/24
    vlan pvid 40
    vlan participation exclude 1
    vlan participation include 40
    exit
exit

```

3. Enable routing on each VLAN interface and assign an IP address.

```

config
    interface vlan 10
        routing
        ip address 1.1.1.1 255.255.255.0
    exit
interface vlan 20
    routing
    ip address 2.2.2.1 255.255.255.0
    exit
interface vlan 30
    routing
    ip address 3.3.3.1 255.255.255.0
    exit
interface vlan 40
    routing
    ip address 4.4.4.3 255.255.255.0
    exit

```

4. Enable IP Routing (Global configuration).

```

config
    ip routing
    exit

```

After this step, if traffic with the following characteristics is sent, it will be routed from VLAN routing interface 10 to VLAN routing interface 20.

Source IP: 1.1.1.2  
Destination IP: 2.2.2.2

In order to policy route such traffic to VLAN routing interface 30, continue with the following steps:

5. Create an access-list matching incoming traffic.

```

config
    access-list 1 permit 1.1.1.2 0.0.0.255
    exit

```

6. Create a route-map and add match/set terms to the route-map.

```

configure
    route-map pbr_test permit 10
    match ip address 1
    set ip next-hop 3.3.3.3
    exit
exit

```

7. Assign a route-map to VLAN routing interface 10.

```

config

```

```

interface vlan 10
ip policy pbr_test
exit
exit

```

After this step, traffic mentioned in [Step 5](#) is policy routed to VLAN interface 30. Counters are incremented in the show route-map command indicating that traffic is being policy routed.

#### 8. Run the show command.

```

(Broadcom FASTPATH Routing) #show route-map pbr_test
route-map pbr_test permit 10

```

Match clauses:

```

ip address (access-lists) : 1

```

Set clauses:

```

ip next-hop 3.3.3.3

```

Policy routing matches: 19922869 packets, 1275063872 bytes

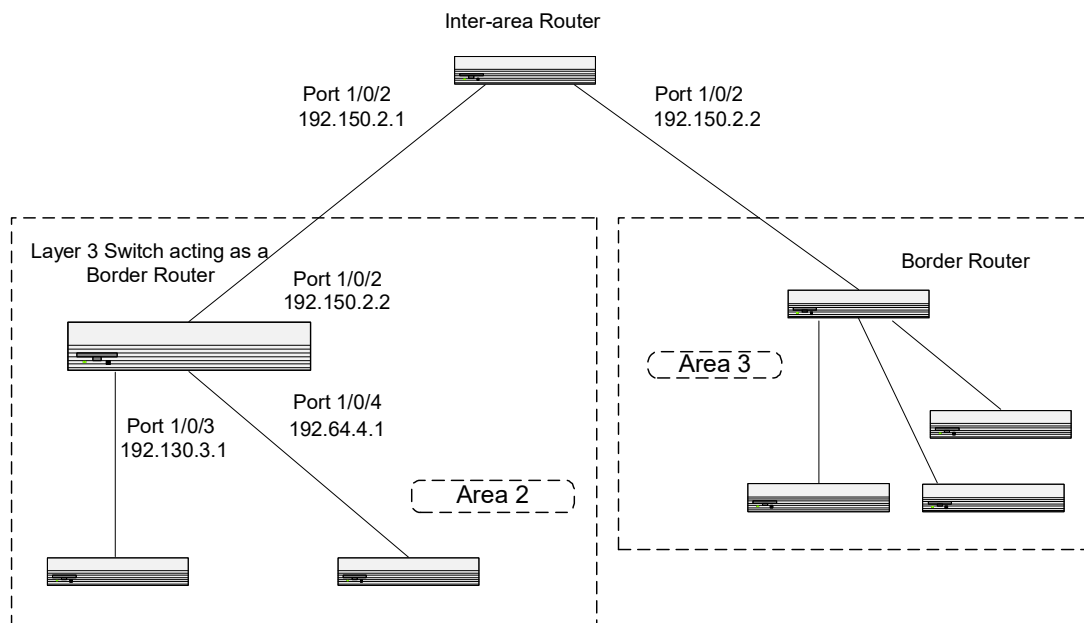
## 9.5 Configuring OSPF

This section contains two OSPF configuration examples.

### Example 1: Configuring an OSPF Border Router and Setting Interface Costs

The following example shows you how to configure an OSPF border router areas and interfaces in FASTPATH software.

Figure 435: OSPF Example Network Diagram: Border Router



### 9.5.1 Using the CLI to Configure OSPF

#### 1. Enable routing on the switch.

```

(Broadcom FASTPATH Routing) #config
ip routing
exit

```

#### 2. For ports 1/0/2, 1/0/3, and 1/0/4, enable routing and assign IP addresses.

```

config
interface 1/0/2

```

```

    routing
    ip address 192.150.2.2 255.255.255.0
    exit
interface 1/0/3
    routing
    ip address 192.130.3.1 255.255.255.0
    exit
interface 1/0/4
    routing
    ip address 192.64.4.1 255.255.255.0
    exit
exit

```

3. Specify a router ID and disable 1583 compatibility to prevent a routing loop (IPv4-only).

```

config
router ospf
    router-id 192.150.9.9
    no 1583compatibility
    exit
exit

```

4. Configure the OSPF area ID, priority, and cost for each interface.

OSPF is globally enabled by default. To make it operational on the router, you configure OSPF for particular interfaces and identify which area the interface is associated with. The following commands also sets the priority and cost for the ports:

```

config
interface 1/0/2
    ip ospf area 0.0.0.0
    ip ospf priority 128
    ip ospf cost 32
    exit
interface 1/0/3
    ip ospf area 0.0.0.2
    ip ospf priority 255
    ip ospf cost 64
    exit
interface 1/0/4
    ip ospf area 0.0.0.2
    ip ospf priority 255
    ip ospf cost 64
    exit
exit

```

---

### **NOTICE**

In OSPFv2, you can also enable OSPF on an interface in global configuration mode by associating a network interface, identified by a network IP address and wildcard mask, with an area. The following example is equivalent to defining interface 1/0/4 in area 2, as in the previous example:

```

(Broadcom FASTPATH Routing) #config
router ospf
network 192.164.4.0 0.0.0.255 area 2

```

---

### Example 2: Configuring Stub and NSSA Areas

In this example, Area 0 connects directly to two other areas: Area 1 is defined as a stub area and Area 2 is defined as an NSSA area.

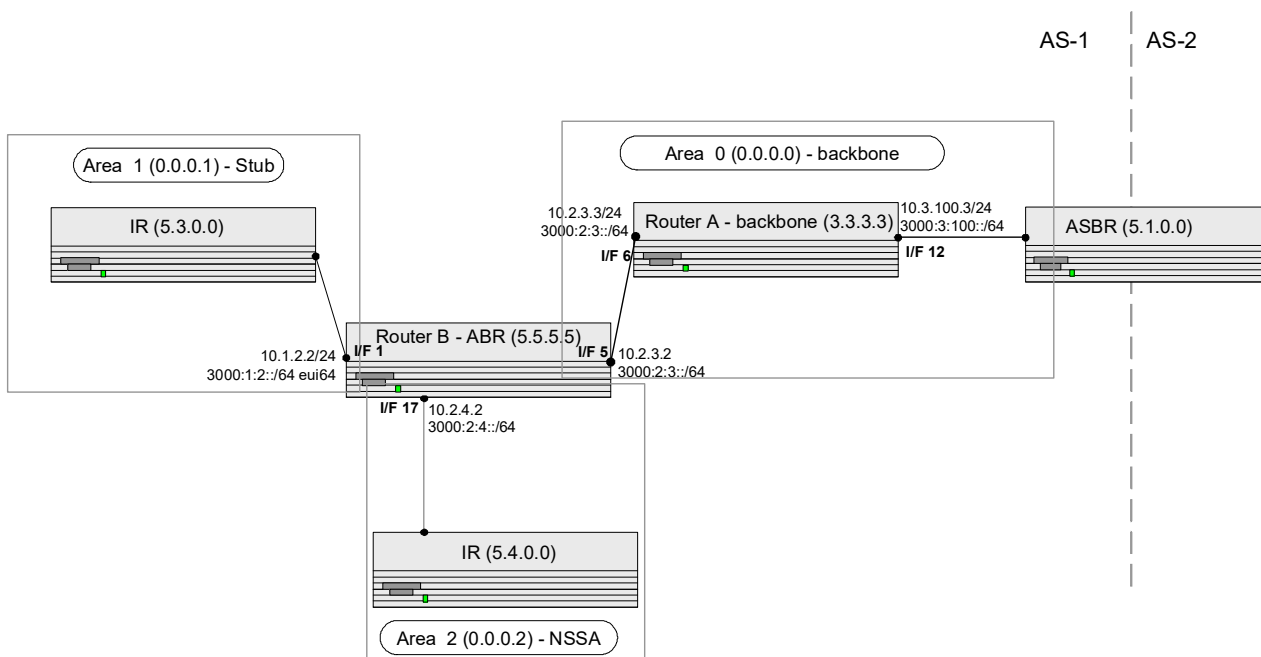
---

### **NOTICE**

OSPFv2 and OSPFv3 can operate concurrently on a network and on the same interfaces (although they do not interact). This example configures both protocols simultaneously.

---

Figure 436: OSPF Configuration—Stub Area and NSSA Area



## 9.5.2 Using the CLI to Configure OSPF Areas

Configure Router A: Router A is a backbone router. It links to an ASBR (not defined here) that routes traffic outside the AS.

### 1. Globally enable IPv6 and IPv4 routing:

```
Broadcom FASTPATH Routing) #configure
  ipv6 unicast-routing
  ip routing
```

### 2. Configure IP address and enable OSPF on interfaces 6 and 12 and enable IPv6 OSPF on the interfaces. (OSPF is enabled on the IPv4 interface in the next code group.)

```
interface 1/0/6
  routing
  ipv6 enable
  ip address 10.2.3.3 255.255.255.0
  ipv6 address 3000:2:3::/64 eui64
  ipv6 ospf
  exit

interface 1/0/12
  routing
  ip address 10.3.100.3 255.255.255.0
  ipv6 address 3000:3:100::/64 eui64
  ipv6 enable
  ipv6 ospf
  exit
```

### 3. Define an OSPF router. Enable OSPF for IPv4 on the two interfaces by globally defining the range of IP addresses associated with each interface, and then associating those ranges with Area 0:

```
ipv6 router ospf
  router-id 3.3.3.3
  exit
router ospf
  router-id 3.3.3.3
  network 10.2.3.0 0.0.0.255 area 0.0.0.0
  network 10.3.100.0 0.0.0.255 area 0.0.0.0
  exit
```



```
exit
```

Configure Router B: Router B is a ABR that connects Area 0 to Areas 1 and 2.

1. **Configure IPv6 and IPv4 routing.** The static routes are included for illustration only: Redistributed static routes, like routes distributed from other protocols, are not injected into stub areas such as Area 1:

```
Broadcom FASTPATH Routing) #configure
ipv6 unicast-routing
ipv6 route 3000:44:44::/64 3000:2:3::210:18ff:fe82:c14
ip route 10.23.67.0 255.255.255.0 10.2.3.3
```

2. **On interfaces 1, 5, and 17, configure IPv4 and IPv6 addresses and enable OSPF on the interfaces.** For IPv6, associate interface 1 with Area 1 and interface 17 with Area 2. (OSPF is enabled on the IPv4 interface in the next code group.)

```
interface 1/0/1
  routing
  ip address 10.1.2.2 255.255.255.0
  ipv6 address 3000:1:2::/64 eui64
  ipv6 ospf
  ipv6 ospf areaid 1
  exit
interface 1/0/5
  routing
  ip address 10.2.3.2 255.255.255.0
  ipv6 address 3000:2:3::/64 eui64
  ipv6 ospf
  exit
interface 1/0/17
  routing
  ip address 10.2.4.2 255.255.255.0
  ipv6 address 3000:2:4::/64 eui64
  ipv6 ospf
  ipv6 ospf areaid 2
  exit
```

3. **For IPv4: Define an OSPF router.** Define Area 1 as a stub. Enable OSPF for IPv4 on interfaces 1, 5, and 17 by globally defining the range of IP addresses associated with each interface, and then associating those ranges with Areas 1, 0, and 17, respectively. Then, configure a metric cost to associate with static routes when they are redistributed via OSPF:

```
router ospf
  router-id 2.2.2.2
  area 0.0.0.1 stub
  area 0.0.0.2 nssa
  network 10.1.2.0 0.0.0.255 area 0.0.0.1
  network 10.2.3.0 0.0.0.255 area 0.0.0.0
  network 10.2.4.0 0.0.0.255 area 0.0.0.2
  redistribute static metric 1 subnets
  exit
```

4. **For IPv6: Define an OSPF router.** Define Area 1 as a stub and area 2 as a Not-So-Stubby-Area (NSSA). Configure a metric cost to associate with static routes when they are redistributed via OSPF:

```
ipv6 router ospf
  router-id 2.2.2.2
  area 0.0.0.1 stub
  area 0.0.0.2 nssa
  redistribute static metric 105 metric-type 1
  exit
exit
```

### 9.5.3 Using the CLI to Configure OSPFv3 Enhancements

To configure OSPFv3 enhancements using the CLI:

1. Enable the IPv6 router admin mode on the switch.

```
(Broadcom FASTPATH Routing) #config
ipv6 unicast-routing
exit
```

2. On port 1/0/1:

- Enable routing
- Assign an IPv6 address
- Enable OSPFv3 and specify its area
- Set the OSPFv3 interface type
- Enable Link LSA Suppression.

```
config
interface 1/0/1
routing
ipv6 address 2000::1/64
ipv6 ospf area 0
ipv6 ospf network point-to-point
ipv6 ospf link-lsa-suppression
exit
exit
```


3. Specify a router ID, configure the stub router mode LSA metrics, and OSPF timers.

```
config
ipv6 router ospf
router-id 1.1.1.1
max-metric router-lsa on-startup 1000 summary-lsa
timers throttle spf 50 2000 5000
timers pacing lsa-group 120
exit
exit
```

## 9.6 Configuring 802.1X Network Access Control

This example configures a single RADIUS server used for authentication and accounting at 10.10.10.10. The shared secret is configured to be secret. The switch is configured to require that the 802.1X access method is through a RADIUS server. IEEE 802.1X port-based access control is enabled for the system, and interface 1/0/1 is configured to be in force-authorized mode because this is where the RADIUS server and protected network resources are located.

Figure 437: Switch with 802.1x Network Access Control



If a user, or supplicant, attempts to communicate via the switch on any interface except interface 1/0/1, the system challenges the supplicant for login credentials. The system encrypts the provided information and transmits it to the RADIUS server. If the RADIUS server grants access, the system sets the 802.1X port state of the interface to authorized, and the supplicant is able to access network resources.

## 9.6.1 Using the CLI to configure 802.1X Port-Based Access Control

1. Configure the RADIUS authentication server IP address.

```
(Broadcom FASTPATH Routing) #config
radius server host auth 10.10.10.10
```

2. Configure the RADIUS authentication server secret.

```
radius server key auth 10.10.10.10
secret
secret
```

3. Configure the RADIUS accounting server IP address.

```
radius server host acct 10.10.10.10
```

4. Configure the RADIUS accounting server secret.

```
radius server key acct 10.10.10.10
secret
secret
```

5. Enable RADIUS accounting mode.

```
radius accounting mode
```

6. Set IEEE 802.1X to use RADIUS as the AAA method.

```
aaa authentication dot1x default radius
```

7. Enable 802.1X authentication on the switch.

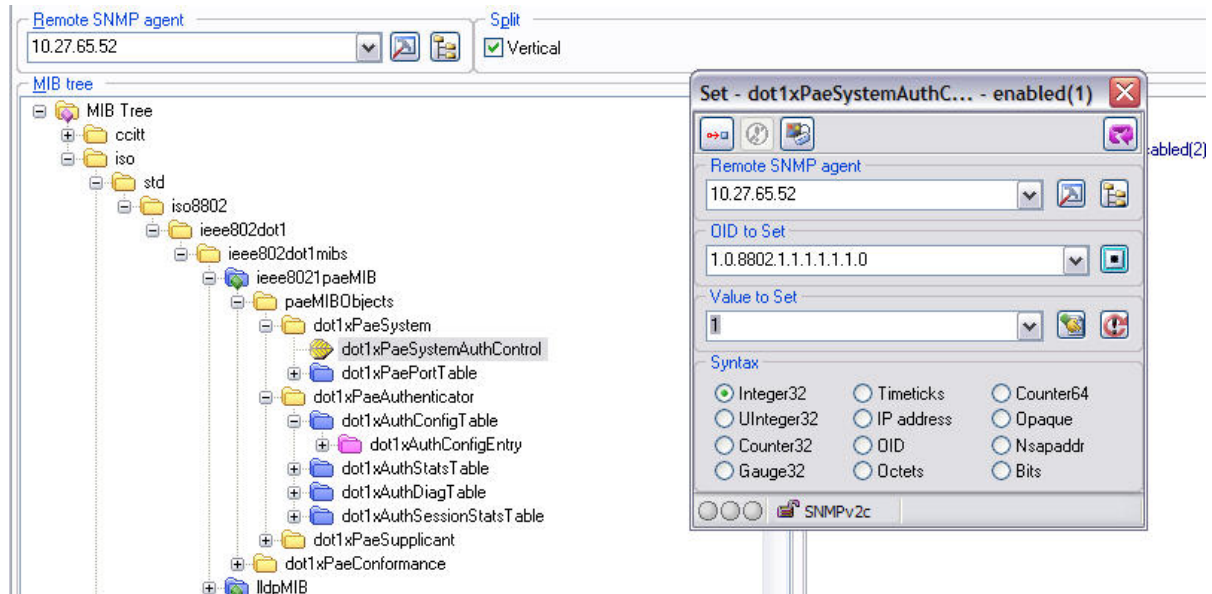
```
dot1x system-auth-control
```

8. Set the 802.1X mode for port 1/0/1 to Force Authorized.

```
interface 1/0/1
dot1x port-control force-authorized
exit
```

## 9.6.2 Using SNMP to configure 802.1X Port-Based Access Control

1. Use the `agentRadiusServerStatus` in the `agentRadiusServerConfigTable` under the `FASTPATH-RADIUS-AUTH-CLIENT-MIB` to create a new RADIUS server entry.
2. Use the `agentRadiusServerAddress` object to configure the RADIUS authentication server IP address as 10.10.10.10.
3. Use the `agentRadiusServerSecret` object to configure the RADIUS authentication server secret.
4. Use the `agentRadiusAccountingStatus` object in the `agentRadiusAccountingConfigTable` to create a RADIUS accounting server.
5. Use the `agentRadiusAccountingServerAddress` object to configure the RADIUS accounting server IP address. as 10.10.10.10.
6. Use the `agentRadiusAccountingSecret` object to configure the RADIUS accounting server secret.
7. Use the `agentRadiusAccountingStatus` object to enable RADIUS accounting mode.
8. Use the `agentUserConfigDefaultAuthenticationList` object in `agentAuthenticationGroup` in the `FASTPATH-SWITCHING` module to set RADIUS as the default login list for dot1x.
9. To enable 802.1X authentication on the switch, set the `dot1xPaeSystemAuthControl` object in the `IEEE8021-PAE-MIB` module to enable (1).



10. To set the 802.1X mode for port 1/0/1 to Force Authorized, use the `agentDot1xPortControlMode` object in the `agentDot1xPortConfigTable`, which is in FASTPATH-DOT1X-ADVANCED-FEATURES-MIB.

## 9.7 Configuring Authentication Tiering

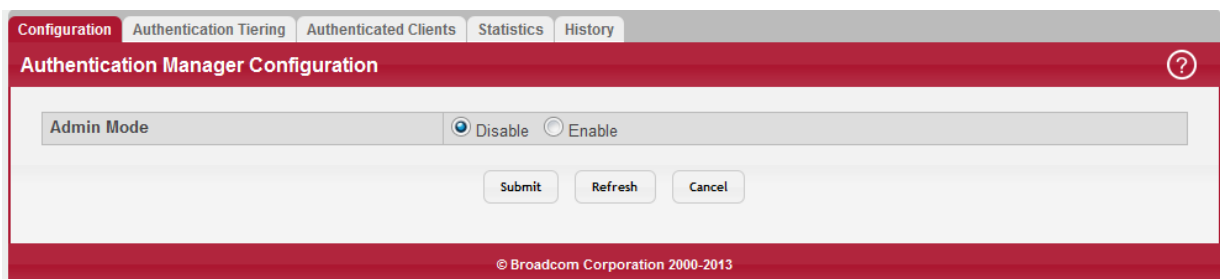
Authentication Tiering can be configured through either the web interface or the CLI.

### 9.7.1 Configuring Authentication Tiering Using the Web Interface

To configure Authentication Tiering through the web interface:

1. Access the Security > Authentication Manager > Configuration page.

Figure 438: Authentication Manager Configuration Page



2. Select Enable.
3. Click Submit.
4. Access the Security > Authentication Manager > Authentication Tiering page.

Figure 439: Authentication Tiering Page

<input type="checkbox"/>	Interface	Configured Order	Enabled Order	Configured Priority	Enabled Priority	Authenticated Clients	Re-Authentication Timer
<input type="checkbox"/>	0/1		⏻		⏻	0	300
<input type="checkbox"/>	0/2		⏻		⏻	0	300
<input checked="" type="checkbox"/>	0/3		⏻		⏻	0	300
<input type="checkbox"/>	0/4		⏻		⏻	0	300
<input type="checkbox"/>	0/5		⏻		⏻	0	300
<input type="checkbox"/>	0/6		⏻		⏻	0	300
<input type="checkbox"/>	0/7		⏻		⏻	0	300
<input type="checkbox"/>	0/8		⏻		⏻	0	300
<input type="checkbox"/>	0/9		⏻		⏻	0	300
<input type="checkbox"/>	0/10		⏻		⏻	0	300

© Broadcom Corporation 2000-2013

5. Select interface 1/0/3 check box and click Edit.

The Edit Authentication Tiering page displays (see [Figure 440: "Edit Authentication Tiering Page,"](#) on page 511).

Figure 440: Edit Authentication Tiering Page

Interface: 0/3

Re-Authentication Timer (Seconds): 300 (300 to 65535)

Configured Method Order

Methods	Order
Dot1x	
MAB	
Captive Portal	

Configured Method Priority

Methods	Priority
Dot1x	
MAB	
Captive Portal	

Submit Cancel

6. Type 10000 in the Re-Authentication Timer field.
7. In the Configured Method Order box, move Dot1x, MAB, and Captive Portal to the Order box by selecting the method and clicking the > button.
8. In the Configured Method Priority box, move Dot1x and Captive Portal to the Priority box by selecting the method and clicking the > button.
9. Click Submit.

## 9.7.2 Configuring Authentication Tiering Using the CLI

To Configure Authentication Tiering Using the CLI:

1. Enable Authentication Tiering globally.

```
config
  authentication enable
  exit
```

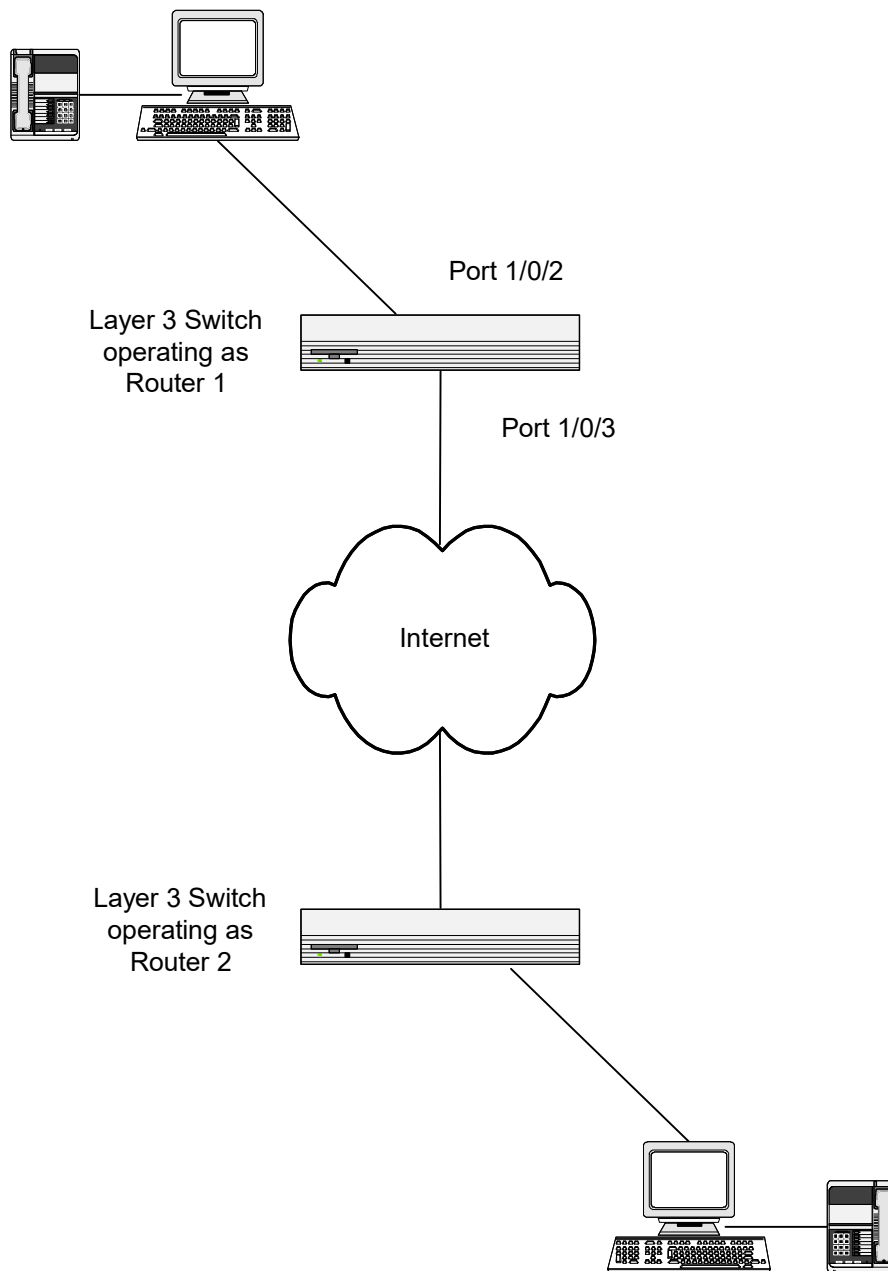
2. Configure the authentication order, priority, and restart timer on interface 1/0/3.

```
config
  interface 1/0/3
  authentication order dot1x mab captive-portal
  authentication priority captive-portal dot1x
  authentication restart 10000
  exit
exit
```

## 9.8 Configuring Differentiated Services for VoIP

One of the most valuable uses of DiffServ is to support Voice over IP (VoIP). VoIP traffic is inherently time-sensitive: for a network to provide acceptable service, a guaranteed transmission rate is vital. This example shows one way to provide the necessary quality of service: how to set up a class for UDP traffic, have that traffic marked on the inbound side, and then expedite the traffic on the outbound side. The configuration script is for Router 1 in the accompanying diagram: a similar script should be applied to Router 2.

Figure 441: DiffServ VoIP Example Network Diagram



### 9.8.1 Using the CLI to Configure DiffServ VoIP Support

1. Enter Global Config mode. Set queue 5 on all ports to use strict priority mode. This queue shall be used for all VoIP packets. Activate DiffServ for the switch.  
(Broadcom FASTPATH Routing) #config  
cos-queue strict 5  
diffserv
2. Create a DiffServ classifier named `class_voip` and define a single match criterion to detect UDP packets. The class type `match-all` indicates that all match criteria defined for the class must be satisfied in order for a packet to be considered a match.  
class-map match-all class\_voip  
match protocol udp  
exit

3. Create a second DiffServ classifier named `class_ef` and define a single match criterion to detect a DiffServ code point (DSCP) of 'EF' (expedited forwarding). This handles incoming traffic that was previously marked as expedited elsewhere in the network.

```
class-map match-all class_ef
  match ip dscp ef
  exit
```

4. Create a DiffServ policy for inbound traffic named `pol_voip`, then add the previously created classes `class_ef` and `class_voip` as instances within this policy.

This policy handles incoming packets already marked with a DSCP value of EF (per `class_ef` definition), or marks UDP packets per the `class_voip` definition) with a DSCP value of EF. In each case, the matching packets are assigned internally to use queue 5 of the egress port to which they are forwarded.

```
policy-map pol_voip in
  class class_ef
    assign-queue 5
  exit
  class class_voip
    mark ip-dscp ef
    assign-queue 5
  exit
exit
```

Attach the defined policy to an inbound service interface.

```
interface 1/0/2
  service-policy in pol_voip
  exit
exit
```

## 9.8.2 Using SNMP to Configure DiffServ VoIP Support

1. Use the `agentDiffServGenStatusAdminMode` object in `agentDiffServGenStatusGroup` under `fastPathQOSDiffServPrivate` in the FASTPATH-QOS-DIFFSERV-PRIVATE-MIB module to activate DiffServ for the switch.
2. To set queue 5 on all ports to use strict priority mode, use the `agentCosQueueSchedulerType` in the `agentCosQueueTable` in the FASTPATH-QOS-COS-MIB module. This queue is used for all VoIP packets.
3. Use the `agentDiffServClassRowStatus` object in the `agentDiffServClassTable` to create two new DiffServ instances. Set the value to `CreateAndGo` (4).
4. Use the `agentDiffServClassName` in the `agentDiffServClassTable` to name the first DiffServ classifier `class_voip` and the second classifier `class_ef`.
5. Use the `agentDiffServClassType` in the `agentDiffServClassTable` to set the class type for each classifier to `All` (1).
6. Use the `agentDiffServClassRuleMatchEntryType` in the `agentDiffServClassRuleTable` to set `class_voip` to match a protocol (9) and `class_ef` to match an IP DSCP value (6).
7. For `class_voip`, define a single match criterion to detect UDP packets by setting the `agentDiffServClassRuleMatchProtocolNum` in the `agentDiffServClassRuleTable` to 17.
8. Use the `agentDiffServClassRuleMatchIpDscp` object in the `agentDiffServClassRuleTable` to define a single match criterion to detect a DSCP of EF (46). This handles incoming traffic that was previously marked as expedited elsewhere in the network.
9. Use the `agentDiffServPolicyRowStatus` object in the `agentDiffServPolicyTable` to create a DiffServ policy. Set the value to `CreateAndGo` (4).
10. Use the `agentDiffServPolicyType` object to set the policy direction so that it applies to inbound (1) traffic.
11. Use the `agentDiffServPolicyName` object to name the new DiffServ instance `pol_voip`.
12. Use the `agentDiffServPolicyInstRowStatus` object in the `agentDiffServPolicyInstTable` to create new instances that will be associated with the previously created classes (`class_ef` and `class_voip`).
13. Use the `agentDiffServPolicyInstClassIndex` object to associate `class_ef` and `class_voip` with the policy instances.
14. Use the `agentDiffServPolicyAttrRowStatus` object in the `agentDiffServPolicyAttrTable` to create three instances.



15. Use the `agentDiffServPolicyAttrStmntAssignQueueId` to set the queue value for instances 1.1.1 and 1.2.2 to 5, so that matching packets are assigned internally to use queue 5 of the egress port to which they are forwarded.
16. Use the `agentDiffServPolicyAttrStmntMarkIpDscpVal` object to set the value of instance 1.2.1 to 46, which marks UDP packets (per the `class_voip` definition) with a DSCP value of EF.
17. Create an instance for the interface that will have the policy attached by using the `agentDiffServServiceRowStatus` object in the `agentDiffServServiceTable`. For example, to create an instance for interface 1/0/2, set 2.1 to `CreateAndGO (4)`.
18. Attach the policy to the interface instance by using the `agentDiffServServicePolicyIndex` object. Set the value of the instance to 1.

## 9.9 Configuring PIM

Protocol Independent Multicast (PIM) is a standard multicast routing protocol that provides scalable inter-domain multicast routing across the Internet, independent of the mechanisms provided by any particular unicast routing protocol.

PIM-SM is used to efficiently route multicast traffic to multicast groups that may span wide area networks where bandwidth is a constraint. PIM-SM is defined in RFC 4601.

The following example configures PIM-SM for IPv4 on a router.

### 9.9.1 Using the CLI to Configure PIM-SMv4

The following example configures PIM-SM for IPv4 in a FASTPATH router:

1. Configure IP routing, IP multicast, IGMP, and PIM-SM in the global configuration mode with the following commands:

```
(Broadcom FASTPATH Routing) #configure
    ip routing
    ip multicast
    ip igmp
    ip pim sparse
exit
```

2. Configure a PIM-SM rendezvous point with an IP address and group range. The RP IP address will serve as an RP for the range of potential multicast groups specified in the group range. Use the following command:

```
ip pim rp-address 1.1.1.1 224.0.0.0 255.0.0.0
```

3. Enable routing, IGMP, PIM-SM on one or more interfaces with the following commands:

```
interface 1/0/1
    routing
    ip address 1.1.1.1 255.255.255.0
    ip igmp
    ip pim
exit
interface 1/0/2
    routing
    ip address 2.2.2.2 255.255.255.0
    ip igmp
    ip pim
exit
```

The above configuration example enabled PIM-SM on the router.

## 9.9.2 Using SNMP to Configure PIM-SMv4

- Use the following objects to configure an OSPF router and globally enable IP routing, multicast, IGMP, and PIM-SM.
  - Enable OSPF: `ospfAdminStat` under `ospfGeneralGroup` in the OSPF-MIB module
  - Set OSPF router ID: `ospfRouterId` under `ospfGeneralGroup` in the OSPF-MIB module
  - Enable routing: `agentSwitchIpRoutingMode` object in `agentSwitchIpGroup` under `fastPathRouting`
  - Enable multicast: `agentMulticastRoutingAdminMode` under `agentMulticastRoutingConfigGroup` in the FASTPATH-MULTICAST-MIB module
  - Enable IGMP: `agentMulticastIGMPAdminMode` under `agentMulticastIGMPConfigGroup`
  - Enable PIM-SM: `agentMulticastPIMSMAdminMode` under `agentMulticastPIMSMConfigGroup`
- Use the `pimSmStaticRPIPAAddress` object in the `agentMulticastPIMSMStaticRPTTable` under `agentMulticastPIMSMConfigGroup` to configure a PIM-SM rendezvous point with an IP address (1.1.1.1) and group range 224.0.0.0 to 240.0.0.0. The IP address will serve as an RP for the range of potential multicast groups specified in the group range.
- Use the following objects to enable routing, IGMP, PIM-SM, and OSPF on one or more interfaces:
  - Enable routing on the interface: `agentSwitchIpInterfaceRoutingMode` in the `agentSwitchIpInterfaceTable` under the FASTPATH-ROUTING-MIB module.
  - Enable IGMP on the interface: `mgmdRouterInterfaceStatus` in the `mgmdRouterInterfaceTable` under the MGMD-STD-MIB module.
  - Enable PIM-SM on an interface: `pimSmInterfaceStatus` in the `pimSmInterfaceTable` under the PIM-STD-MIB module.
  - Enable OSPF on an interface: `ospfIfStatus` in the `ospfIfTable` in the OSPF-MIB module.
  - Use the `agentSwitchIpInterfaceIpAddress` and `agentSwitchIpInterfaceNetMask` objects in the `agentSwitchIpInterfaceTable` under FASTPATH-ROUTING-MIB to assign an IP address and subnet mask to each interface.

## 9.9.3 Configuring IP Multicast Routing with PIM Sparse Mode

Protocol Independent Multicast (PIM) is a standard multicast routing protocol that provides scalable inter-domain multicast routing across the Internet, independent of the mechanisms provided by any particular unicast routing protocol.

PIM Sparse mode (PIM-SM) is used to efficiently route multicast traffic to multicast groups that may span wide area networks where bandwidth is a constraint. PIM-SM is defined in RFC 4601.

### 9.9.3.1 Configuring PIM-SM

The following example configures PIM-SM for IPv4 in a FASTPATH router:

- Configure IP routing, IP multicast, IGMP, and PIM-SM in the global configuration mode with the following commands:

```
(Broadcom FASTPATH Routing) #configure
    ip routing
    ip multicast
    ip igmp
    ip pim sparse
exit
```

- Configure a PIM-SM rendezvous point with an IP address and group range. The RP IP address will serve as an RP for the range of potential multicast groups specified in the group range. Use the following command:

```
ip pim rp-address 1.1.1.1 224.0.0.0 255.0.0.0
```

- Enable routing, IGMP, PIM-SM on one or more interfaces with the following commands:

```
interface 1/0/1
    routing
    ip address 1.1.1.1 255.255.255.0
    ip igmp
    ip pim
exit
```

```

interface 1/0/2
  routing
  ip address 2.2.2.2 255.255.255.0
  ip igmp
  ip pim
exit

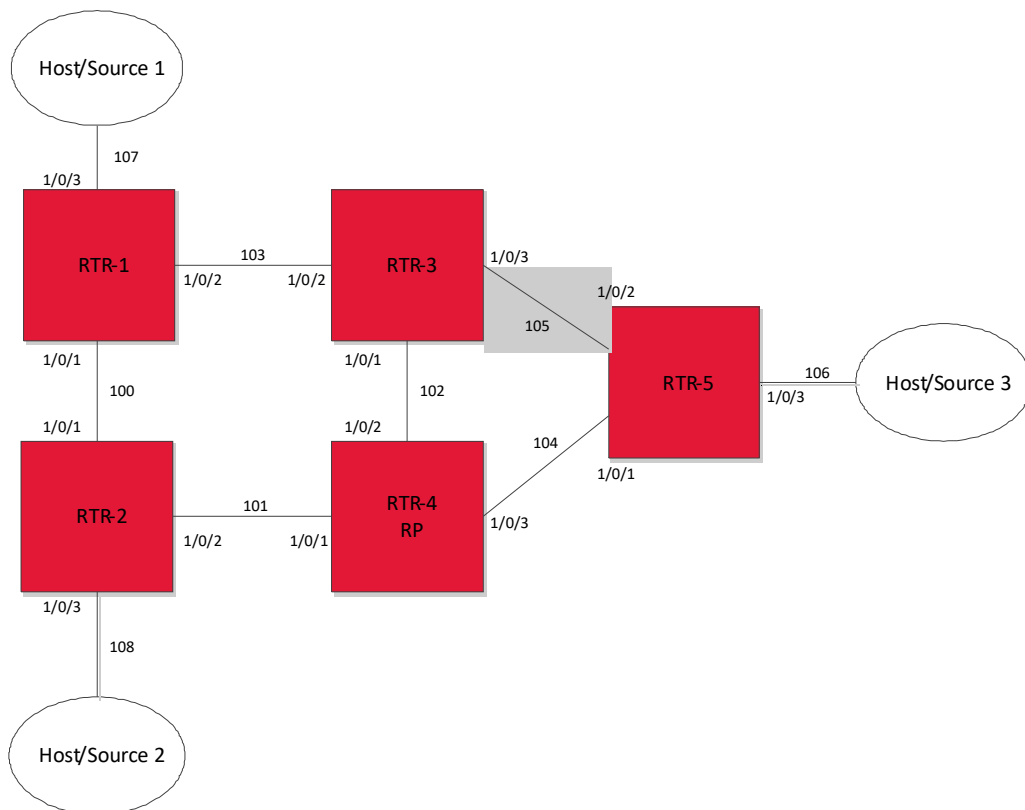
```

The above configuration example enabled PIM-SM on the router.

### 9.9.3.2 Configuring Multicast Data Forwarding with PIM RP as an Intermediate Router

The following example configures and explains how PIM-SM works in a topology where a PIM RP is configured to be more than one hop away from the multicast source.

Figure 442: Multicast Data Forwarding Example



Configure PIM-SM on all the routers in the above topology in a similar way as was explained in [Section 9.9.3.1: "Configuring PIM-SM"](#). The exact configuration of each of the routers can be found at the end of this example.

1. Configure RTR-4 as a PIM Rendezvous point as shown below:

```

ip pim rp-address 192.168.102.4 24.0.0.0 255.0.0.0

```

In the topology shown in [Figure 442: "Multicast Data Forwarding Example," on page 517](#), RTR-4 is the RP. Source-1 (192.168.107.10) is the multicast source transmitting data for the multicast group 224.1.2.1. Host 3 announces interest in receiving data for the multicast group 224.1.2.1 by sending an IGMPv2 Join message to RTR-5. RTR-5 initiates creation of the shared tree by sending a PIM (\*,G) Join towards the RP and creates a (\*,G) entry in their route table.

The multicast route table at RTR-5 looks like the following example:

```

(FASTPATH-RTR5)#show ip mcast mroute summary

```

```

Multicast Route Table Summary
Incoming Outgoing

```

```
Source IP   Group IP   Protocol Interface Interface List
-----
```

```
*    224.1.2.1   PIMSM   vlan104  vlan106
```

After Source 1 starts sending data, RTR-1 (DR adjacent the source) encapsulates the multicast data in PIM Register messages and unicasts it to RP (RTR-4). RTR-4 decapsulates the data packets and forwards them natively to RTR-5. FASTPATH registers encapsulation and decapsulation in the software. If the source transmits data at higher rates, this can be taxing on the CPU and has to be avoided. To address this, on receiving the first Register message, RTR-4 (RP) initiates a switch-over to the source tree by sending a (S,G) Join towards the source. This example assumes the register rate-limit is set to the default value of 0. The ip pim register-rate-limit command configures the rate limit value. After RTR-4 starts receiving multicast data natively on the source tree, RTR-4 sends a Register-Stop message to stop RTR-1 from sending further encapsulated Register messages.

After the RTR-4 sends a Register-Stop message, the multicast route table at RTR-4 (RP) looks like the following example:

```
(FASTPATH-RTR4)#show ip mcast mroute summary
```

```

          Multicast Route Table Summary
Source IP      Group IP      Protocol   Incoming   Outgoing
              Interface Interface List
-----
*             224.1.2.1   PIMSM          vlan104
192.168.107.10 224.1.2.1   PIMSM          vlan101  vlan104
```

The multicast route table at RTR-5 looks like the following example:

```
(FASTPATH-RTR5)#show ip mcast mroute summary
```

```

          Multicast Route Table Summary
Source IP      Group IP      Protocol   Incoming   Outgoing
              Interface Interface List
-----
*             224.1.2.1   PIMSM          vlan104  vlan106
192.168.107.10 224.1.2.1   PIMSM          vlan104  vlan106
```

The unicast route table at RTR-5 shows that the best path towards Source-1 is through RTR-3, but RTR-5 receives data through RTR-4, which is a sub-optimal (longer) path. At the last hop router, FASTPATH PIM-SM, on receiving the first data packet initiates a switch-over to the source tree by sending a (S,G) Join towards the source and prunes the RP tree by sending a (\*,G) Prune towards the RP.

After the switch-over, the multicast route table at RTR-5 is as follows:

```
(FASTPATH-RTR5)#show ip mcast mroute summary
```

```

          Multicast Route Table Summary
Source IP      Group IP      Protocol   Incoming   Outgoing
              Interface Interface List
-----
*             224.1.2.1   PIMSM          vlan104
192.168.107.10 224.1.2.1   PIMSM          vlan105  vlan106
```

Host-3 now receives multicast traffic from the source on the optimal path. Now consider a case where a network re-convergence takes place. Assume that RTR-3 goes down, and as a result the multicast source tree that was built before (RTR-5 to RTR-1) is broken and Host-3 stops receiving multicast data. Note that the shared tree (RTR-5 to RTR-4) is intact as the network did not fail in this path.

RTR-3 is the Incoming interface for the (S,G) stream that RTR-5 receives in the shortest path. Therefore, when an Incoming interface is down, FASTPATH PIM-SM deactivates all the (S,G) entries with a source that is reachable through that Incoming interface. In this case, as RTR-3 is the Incoming interface for RTR-5's (S,G) entry, RTR-5 deactivates the (S,G) entry.

At this point in time, the multicast route table at RTR-5 looks like the following example:

```
(FASTPATH-RTR5)#show ip mcast mroute summary
```

```

          Multicast Route Table Summary
Source IP      Group IP      Protocol   Incoming   Outgoing
              Interface Interface List
-----
```

```
*          224.1.2.1      PIMSM      vlan104
```

After the unicast routing table reconverges after the failover, the underlying unicast routing table manager informs PIM-SM about changes to the existing paths. RTR-5, with the help of the deactivated (S,G) entry, then detects a change in the path to the source, with the new path being via RTR-4. RTR-5 attempts to rebuild the source tree by sending a (S,G) Join towards the source to RTR-4 and eventually receives multicast data from the source via RTR-4 on the shortest path tree.

After the network re-convergence, the multicast route table at RTR-5 shows the following information:

```
(FASTPATH-RTR5)#show ip mcast mroute summary
```

```

          Multicast Route Table Summary
Source IP      Group IP      Protocol    Incoming   Outgoing
-----
          *          224.1.2.1    PIMSM      vlan104   vlan106
192.168.107.10 224.1.2.1    PIMSM      vlan104   vlan106

```

This example demonstrated how multicast routing reconverges and how multicast traffic resumes after a network failover.

For troubleshooting purposes, after a network re-convergence, if multicast data is not resumed from the source to Host 3, the multicast route table at RTR-5 looks similar to the following example. You may notice such behavior if the Outgoing Interface List in the following command output is empty.

```
(FASTPATH-RTR5)#show ip mcast mroute summary
```

```

          Multicast Route Table Summary
Source IP      Group IP      Protocol    Incoming   Outgoing
-----
          *          224.1.2.1    PIMSM      vlan104
192.168.107.10 224.1.2.1    PIMSM      vlan104

```

## 9.10 Configuring FIP Snooping

FIP snooping is a frame inspection method used by the FASTPATH FIP Snooping Bridge to monitor FIP frames and apply policies based on the L2 header information in those frames, following the recommendations in Annex C of FC\_BB\_5 Rev 2.00.

FIP Snooping enables the following features:

- Auto-configuration of Ethernet ACLs based on information in the Ethernet headers of FIP frames.
- Emulation of fibre channel (FC) point-to-point links within the DCB Ethernet network.
- Enhanced FCoE security/robustness by preventing FCoE MAC spoofing.

The FIP Snooping Bridge solution in FASTPATH is intended for use only at the edge or perimeter of the switched network and not on an interior switch.

### 9.10.1 Using the CLI to Configure FIP snooping

Web-based configuration pages for FIP snooping are not available.

To configure FIP snooping:

1. For ports connected to CNAs/ENodes, enable LLDP and DCBX and configure them as DCBX auto-down ports. In this example, the ports connected to the CNAs/ENodes are ports 9 and 10.

```

config
  interface 1/0/9-1/0/10
    lldp transmit
    lldp receive
    lldp dcbx port-role auto-down
  exit

```

2. For ports connected to the FCF (Cisco Nexus 5010/5548), enable LLDP and DCBX and configure these ports as DCBX auto-up ports. In this example, the port connected to the FCF is port 11.

```

interface 1/0/11
  lldp transmit
  lldp receive
  lldp dcbx port-role auto-up

```

```
exit
```

3. In Global Config mode, configure one-to-one global dot1p mapping.

```
classofservice dot1p-mapping 0 0
classofservice dot1p-mapping 1 1
classofservice dot1p-mapping 2 2
classofservice dot1p-mapping 3 3
classofservice dot1p-mapping 4 4
classofservice dot1p-mapping 5 5
classofservice dot1p-mapping 6 6
exit
```

4. Create the FCoE VLAN. In this example, FCoE VLAN ID is 1000.

```
vlan database
vlan 1000
exit
```

5. Add VLAN 1000 membership to the ports connected to CNAs and FCF. Enable VLAN tagging on these ports for FCoE VLAN using the interface command below.

```
config
interface 1/0/9-1/0/11
vlan participation include 1000
vlan tagging 1000
exit
exit
```

6. Enable FIP snooping in FCoE VLAN 1000. Also enable FIP snooping in VLAN 1 to allow FIP VLAN discovery to happen in untagged mode.

```
configure
feature fip-snooping
vlan 1,1000
fip-snooping enable
exit
exit
```

7. Configure FCF facing ports using the interface command below. By default, FIP snooping ports are configured as host/ENode mode.

```
configure
interface 1/0/11
fip-snooping port-mode fcf
exit
exit
```

The following code sample shows the configuration script for the FIP snooping switch configured in the example. Two interfaces (1/0/9 and 1/0/10) are connected to CNAs, and 1/0/11 is connected to CISCO Nexus 5010 FCF.

```
vlan database
vlan 1000
exit
```

```
configure
```

```
feature fip-snooping
vlan 1,1000
fip-snooping enable
exit
```

```
classofservice dot1p-mapping 0 0
classofservice dot1p-mapping 1 1
classofservice dot1p-mapping 2 2
classofservice dot1p-mapping 3 3
classofservice dot1p-mapping 4 4
classofservice dot1p-mapping 5 5
classofservice dot1p-mapping 6 6
```

```
interface 1/0/9
```

```

description 'Brocade CNA'
vlan participation include 1000
vlan tagging 1000
vlan priority 3
lldp transmit
lldp receive
lldp dcbx port-role auto-down
exit

interface 1/0/10
description 'Broadcom CNA'
vlan participation include 1000
vlan tagging 1000
vlan priority 3
lldp transmit
lldp receive
lldp dcbx port-role auto-down
exit

interface 1/0/11
description 'CISCO Nx5010-FCF Facing'
vlan participation include 1000
vlan tagging 1000
vlan priority 3
fip-snooping port-mode fcf
lldp transmit
lldp receive
lldp dcbx port-role auto-up
exit

exit

```

## 9.11 Configuring OpenFlow

The purpose of the OpenFlow feature is to demonstrate Broadcom hardware and software capabilities and to provide a platform on which to build custom networking features, such as the data center tenant networking feature.

OpenFlow 1.0 mode enables the switch to interoperate with standard OpenFlow controllers, such as NOX. The NOX controller has several built-in OpenFlow controller applications that can be used with FASTPATH switches.

The only difference between the FASTPATH tenant networking mode and the OpenFlow 1.0 mode is the switch management paradigm. In tenant networking mode, the FASTPATH switch communicates with the OpenFlow manager to obtain the configuration for OpenFlow controllers, CAPWAP tunnels, and rate limiters. In the OpenFlow 1.0 mode, these configuration parameters are defined through the switch user interface.

The underlying FASTPATH OpenFlow implementation inherently supports multiple hardware tables. Even when operating in the OpenFlow 1.0 mode, the OpenFlow controller can access the different hardware tables by slightly modifying the OFPT\_FLOW\_MOD message. A pure OpenFlow 1.0 standards-compliant controller can only access one hardware table. The administrator can configure the default hardware table accessed by the OpenFlow 1.0 protocol.

### 9.11.1 Enabling and Disabling OpenFlow

The OpenFlow feature can be enabled or disabled by the network administrator. Although this feature is administratively enabled, it is not operational until the switch has an IP address. A separate operational state indicates whether the OpenFlow feature is operational. If the feature is not operational, then another state indicates the reason for the feature to be disabled.

After administratively disabling the feature, the network administrator must wait until the OpenFlow Feature is operationally disabled before re-enabling the feature. The OpenFlow feature can be administratively disabled at any time.

The administrator can allow the switch to automatically assign an IP address to the OpenFlow feature or they can manually select the address. When the address is assigned automatically and the interface with the assigned address goes offline, the switch will select another active interface if one is available. The OpenFlow feature becomes operationally disabled and re-enabled when a new IP address is selected. If the address is assigned statically, the OpenFlow feature becomes operational only when a switch interface with the matching IP address becomes active.

The automatic IP addresses selection is done in the following order of preference.

1. **The loopback interfaces.**
2. **The routing interfaces.**
3. **The network interface.**

The service port cannot be used as an IP interface for the OpenFlow feature. If routing is enabled, the Network interface cannot be used as the OpenFlow interface.

Once the IP address is selected, it is used until the interface goes offline, the feature is disabled, or a more preferred interface becomes available. The selected IP address is used as the end-point of SSL connections and the end-point of CAPWAP tunnels. When the OpenFlow feature is operationally disabled the switch drops connections with the OpenFlow managers and controllers, the switch also purges all flows and configuration programmed by the managers and controllers.

If the administrator changes the OpenFlow variant while the OpenFlow feature is enabled, the switch automatically disables and re-enables the OpenFlow feature causing all flows to be deleted and connections to the controllers to be dropped.

If the administrator changes the default hardware table for OpenFlow 1.0 and if the switch is currently operating in OpenFlow 1.0 variant, the OpenFlow feature is automatically disabled and re-enabled.

### 9.11.2 Interacting with the OpenFlow Manager

The OpenFlow manager is a device that uses the Open vSwitch management protocol to send commands and retrieve status from the switch.

The FASTPATH OpenFlow feature supports the OpenFlow manager only when the DCTENANT\_NET component is selected in CCHelper. If the DCTENANT\_NET component is not selected, the code for interacting with the OpenFlow manager is excluded from the file system whenever practical, and conditionally compiled out from common files. If the DCTENANT\_NET component is selected, but the OpenFlow variant is not configured to be Tenant Networking then the communications with the OpenFlow manager is not supported.

In order to interact with the OpenFlow manager, the OpenFlow feature must be administratively enabled. The administrator must also configure IP addresses of the OpenFlow managers using the switch UI. The OpenFlow manager interaction is handled by the Open vSwitch module called OVSDB.

### 9.11.3 Deploying OpenFlow

The OpenFlow manager uses the management protocol to tell the switch how to communicate with the OpenFlow controllers and the IP addresses of switches in which CAPWAP tunnels must be set up.

If the administrator selects the OpenFlow 1.0 variant of the OpenFlow protocol, the Controller IP addresses are manually assigned through the switch user interface and the CAPWAP tunnel destination IP addresses are also manually assigned.

### 9.11.4 OpenFlow Scenarios

The OpenFlow feature is mainly used in a data center network where devices are located in different parts of the network and require layer-2 connectivity.

The tunneling feature enables the devices to communicate over a layer-3 infrastructure. The flow management feature enables customers to avoid scaling problems and loops associated with the layer-2 network.

The OpenFlow feature can also be used in a research environment, but there are two limitations that make the research use case less attractive. First, only one OpenFlow controller can manage the switch at a time, meaning that concurrent experiments are not supported unless concurrency is handled at the controller level. Second, the OpenFlow controller has complete access to all ports and VLANs, meaning that using the switch for mixed production and experimental traffic is not advisable.

### 9.11.5 OpenFlow Interaction with Other Functions

The OpenFlow component interacts with multiple FASTPATH components by either communicating with these components or sharing common resources with the components.



## 9.11.6 OpenFlow Variants

### 9.11.6.1 OpenFlow 1.0

In OpenFlow 1.0 mode, the switch is a hybrid OpenFlow switch and supports the OpenFlow 1.0 standard. Hybrid OpenFlow switch means OpenFlow acts as a protocol in conjunction with existing switch functionality. OpenFlow 1.0 mode enables the switch to inter-operate with the standard OpenFlow controllers such as NOX, Beacon, and Big Switch. If COTS versions of these controllers are not available, testing is limited to verification via the OVS\_VXCTL tool.


### 9.11.6.2 Data Center Tenant Networking

In Tenant Networking mode, the FASTPATH switch communicates with the OpenFlow manager to obtain the configuration for the OpenFlow controllers, CAPWAP tunnels, and rate limiters. In OpenFlow 1.0 mode, these configuration parameters are defined through the switch user interface.

## 9.11.7 Configuring OpenFlow

The following example uses the network interface's IP address. All switches shown in [Figure 443: "OpenFlow Network Example," on page 523](#) have the same OpenFlow configuration.

Figure 443: OpenFlow Network Example



Use the following commands to configure an OpenFlow network:

1. Configure the network protocol as DHCP with the following command:

```
(Routing) #network protocol dhcp
```

2. Since the controller IP address in this example is configured from the Switch CLI, set the OpenFlow variant mode to openflow1.0 with the following command:

```
Routing) (Config)# openflow variant openflow10
```

3. Set the controller IP address with the following command:

```
(Routing) (Config)#openflow controller 192.168.1.1 6633 tcp
```

4. To insert the flow into the OpenFlow 1.0 match table which can match on all OpenFlow 1.0 fields, set the OpenFlow default flow table to Full-Match with the following command:

```
(Routing) (Config)# openflow default-table full-match
```

5. Enable OpenFlow on the switch with the following command:

```
(Routing) (Config)# openflow enable
```

6. Verify the OpenFlow configuration with the following command:

```
(Routing) #show openflow
```

```
Administrative Mode..... Disable
Operational Status..... Disabled
Disable Reason..... Admin-Disabled
IP Address..... None
Static IP Mode..... Disable
Static IP Address..... 0.0.0.0
Network MTU..... 1518
OpenFlow Variant..... OpenFlow 1.0
Default Table..... full-match
```

```
OpenFlow Manager IP:port Addresses
```

```
-----
```

```
(Routing) #show openflow configured controller
```

IP Address	IP Port	Connection Mode
192.168.1.1	6633	tcp

7. The controller installs rules in the switches. In this example, the following rules have been installed:

#### Switch 1

- Forward any traffic with ingress port 0/1 with Source MAC 00:00:00:00:00:11 to egress port 0/4
- Forward any traffic with ingress port 0/1 with Source MAC 00:00:00:00:00:33 to egress port 0/4
- Forward any traffic with ingress port 0/1 with Source MAC 00:00:00:00:00:55 to egress port 0/4

#### Switch 2

- Forward any traffic with ingress port 0/1 with Source MAC 00:00:00:00:00:11 to egress port 0/2
- Forward any traffic with ingress port 0/1 with Source MAC 00:00:00:00:00:33 to egress port 0/2
- Forward any traffic with ingress port 0/1 with Source MAC 00:00:00:00:00:55 to egress port 0/2

#### Switch 3

- Forward any traffic with ingress port 0/1 with Source MAC 00:00:00:00:00:11 to egress port 0/2
- Forward any traffic with ingress port 0/1 with Source MAC 00:00:00:00:00:33 to egress port 0/2
- Forward any traffic with ingress port 0/1 with Source MAC 00:00:00:00:00:55 to egress port 0/2

8. To verify the installed flows for Switch 1, use the following command:

```
(Routing) #show openflow installed flows
```

Flow 0C9E0D00 type "1DOT0"

Match criteria:

```
Flow table          24 : Priority          32768
Ingress port       0/1 : Src MAC  00:00:00:00:00:11
Actions:
Egress port        0/4
Status:
Duration           7 : Idle                5 : installed in hardware  1
```

Flow F6880900 type "1DOT0"

Match criteria:

```
Flow table          24 : Priority          32768
Ingress port       0/2 : Src MAC  00:00:00:00:00:33
Actions:
Egress port        0/4
Status:
Duration           11 : Idle               9 : installed in hardware  1
```

Flow 36370100 type "1DOT0"

Match criteria:

```
Flow table          24 : Priority          32768
Ingress port       0/3 : Src MAC  00:00:00:00:00:55
Actions:
Egress port        0/4
Status:
Duration           1121 : Idle             1119 : installed in hardware  1
```

9. To verify the installed flows for Switch 2, use the following command:

(Routing) #show openflow installed flows

Flow 0C9E0D00 type "1DOT0"

Match criteria:

```
Flow table          24 : Priority          32768
Ingress port       0/1 : Src MAC  00:00:00:00:00:11
Actions:
Egress port        0/2
Status:
Duration           7 : Idle                5 : installed in hardware  1
```

Flow F6880900 type "1DOT0"

Match criteria:

```
Flow table          24 : Priority          32768
Ingress port       0/1 : Src MAC  00:00:00:00:00:33
Actions:
Egress port        0/2
Status:
Duration           11 : Idle               9 : installed in hardware  1
```

Flow 36370100 type "1DOT0"

Match criteria:

```
Flow table          24 : Priority          32768
Ingress port       0/1 : Src MAC  00:00:00:00:00:55
Actions:
Egress port        0/2
Status:
Duration           1121 : Idle             1119 : installed in hardware  1
```

10. To verify the installed flows for Switch 3, use the following command:

```
(Routing) #show openflow installed flows

Flow 0C9E0D00 type "1DOT0"

Match criteria:
Flow table          24 : Priority          32768
Ingress port       0/1 : Src MAC  00:00:00:00:00:11
Actions:
Egress port        0/2
Status:
Duration           7 : Idle                    5 : installed in hardware    1

Flow F6880900 type "1DOT0"

Match criteria:
Flow table          24 : Priority          32768
Ingress port       0/1 : Src MAC  00:00:00:00:00:33
Actions:
Egress port        0/2
Status:
Duration           11 : Idle                   9 : installed in hardware    1

Flow 36370100 type "1DOT0"

Match criteria:
Flow table          24 : Priority          32768
Ingress port       0/1 : Src MAC  00:00:00:00:00:55
Actions:
Egress port        0/2
Status:
Duration           1121 : Idle                 1119 : installed in hardware  1
```

## 9.12 IGMP and MLD Snooping Switches

### 9.12.1 Snooping Functionality on a FASTPATH Switch

A snooping switch can be configured to receive IGMP packets in a subnet and identify ports on which interested IP multicast listeners are present. It also identifies the ports on which multicast routers are attached, and these are the likely ports on which IP multicast sources are present. This section describes how the snooping switch handles IGMP messages, addresses considerations for IGMP packet and IP multicast traffic forwarding, and provides information about support for IGMP and MLD versions.

#### 9.12.1.1 Query Processing

When a port receives an IGMP query, the snooping switch identifies the receiving port as a multicast router-attached port. The switch maintains the list of multicast router-attached ports on a per-VLAN basis. If the port does not receive periodic IGMP queries, the learned entries maintained in the list expire after a configured interval. The snooping switch stores the last received query on a VLAN because the switch uses this to calculate the max response time value during IGMP Leave message processing.

The snooping switch treats Distance Vector Multicast Routing Protocol (DVMRP) probe messages and Protocol Independent Multicast (PIM) versions 1 and 2 hello messages that it receives similar to IGMP queries by adding the interfaces that receive these messages to the list of multicast router-attached ports.

#### 9.12.1.2 Group Registration

Multicast listeners can register to an IP multicast group by sending an IGMP Report message in response to a general query from a multicast router or by sending an unsolicited IGMP Report message. When the snooping switch processes an IGMP Report message, it creates an entry in the Layer 2 multicast forwarding table for the requested multicast group. Each entry contains a unique VLAN and multicast group combination along with a list of ports on which the IGMP Report was received. Multicast router-attached ports discovered during query processing on the incoming VLAN are automatically added to the newly created Layer 2 multicast forwarding entry.

The created entries expire if no additional IGMP Report messages are received for that multicast group, VLAN, and received port combination. The snooping switch administrator can configure the group expiry on a per-VLAN basis. If all the host registrations expire for a Layer 2 multicast forwarding entry, the entry is removed from the table.

### 9.12.1.3 Group Leave

Multicast listeners can opt to voluntarily leave a group by sending an IGMP Leave message or by not responding to the periodic IGMP queries sent by the multicast router. Upon receiving an IGMP Leave message, the snooping switch sends a group specific query on the received port to solicit IGMP Reports from other interested hosts on the same network segment. The snooping switch waits for the interval specified by the last received query packet (max response time) to receive a response for the Leave query. If there is no response, the port is removed from the Layer 2 multicast forwarding entry. If no querier information is available, a configured value is used. If an IGMP Report is received, the entry remains the same.

Alternatively the administrator can configure the snooping switch to remove the interface that received the IGMP Leave message from the Layer 2 multicast forwarding entry immediately upon processing the message. No IGMP Leave query is sent in this scenario. Configuring the immediate leave is useful in situations where instantaneous control of group registrations is required, which results in better bandwidth control.

### 9.12.1.4 IGMP Packet Forwarding Considerations

The snooping switch forwards received IGMP Report and Leave messages only to multicast router-attached ports in that VLAN. IGMP queries are forwarded to all member interfaces of the VLAN.

The snooping switch is aware of link-layer changes caused by spanning tree operations. When a port is enabled or disabled by spanning tree, a general query is generated by the root bridge. This Topology Change Notification query is sent to all non-multicast router-attached ports of the root bridge, which aids in updating L2 multicast forwarding entries faster so that network disruptions are felt only momentary.

The snooping switch processes all IGMP messages and drops invalid IGMP and MLD messages. Any unrecognized IGMP or MLD messages are forwarded in the VLAN. When the snooping switch originates an IGMP query (leave processing or TCN), it does not alter the version number or fields. The snooping switch leaves this information the same as the query information it received most recently on that VLAN.

### 9.12.1.5 IP Multicast Data Forwarding Considerations

When processing a packet whose destination MAC address is a multicast address, an IEEE standard bridge forwards a copy of the packet to each of the remaining network interfaces that are members of the same VLAN.

By default, unregistered multicast data packets are flooded to all ports in the VLAN.

By creating static L2 multicast forwarding entries, multicast groups can be registered, and data can be forwarded only to selected ports.

### 9.12.1.6 Version Compatibility

The following table shows the IGMP/MLD versions that the FASTPATH snooping switch supports.

Table 407: IGMP/MLD Version Support

Protocol Version	Support
IGMPv1	Yes
IGMPv2	Yes
IGMPv3	Yes
MLDv1	Yes
MLDv2	Yes

## 9.12.2 Snooping Switch Restrictions

This section describes the IGMP and MLD Snooping implementation on a FASTPATH-based snooping switch.

### 9.12.2.1 IGMPv3 and MLDv2 Support

The IGMPv3 and MLDv2 protocols allow multicast listeners to specify the list of hosts from which they want to receive the traffic. FASTPATH snooping switches support the following record types:

- `MODE_IS_INCLUDE`
- `MODE_IS_EXCLUDE`
- `CHANGE_TO_INCLUDE_MODE`
- `CHANGE_TO_EXCLUDE_MODE`
- `ALLOW_NEW_SOURCES`
- `BLOCK_OLD_SOURCES`

The forwarding database built using IGMPv3 reports is based on the Source IP address, Multicast Group address, and VLAN.

When a switch receives an older version (IGMPv2 or IGMPv1) report on a given interface, and on a given VLAN, all the previously gathered source filtering information for that group on the given interface and on the given VLAN is ignored. All IGMPv3 membership reports received on the given interface for that group and on the given VLAN are ignored until IGMPv2/IGMPv1 group times out. This is not in compliance to RFC 3376 section 7.3.2.

### 9.12.2.2 MAC Address-Based Multicast Group

The L2 multicast forwarding table (built using IGMPv2/V1 reports) consists of the IP Multicast group MAC address. For IPv4 multicast groups, 16 IP multicast group addresses map to the same multicast MAC address. For example, 224.1.1.1 and 225.1.1.1 map to the MAC address 01:00:5E:01:01:01, and IP addresses in the range [224-239].3.3.3 map to 01:00:5E:03:03:03. As a result, if a host requests 225.1.1.1 using IGMPv2 or IGMPv1, then it might receive multicast traffic of group 226.1.1.1 as well.

### 9.12.2.3 IGMP Snooping in a Multicast Router

IGMP snooping is a Layer 2 feature and is achieved by using the L2 multicast forwarding table. However, when multicast routing is enabled on a FASTPATH switch, L2 multicast forwarding entries do not affect multicast data forwarding. Instead, the corresponding IP multicast table entries must be created to achieve similar behavior.

On a multicast router, for IGMP snooping to be functional, any multicast routing protocol needs to be operationally enabled on the routing interface. IGMP snooping also needs to be enabled on the VLAN corresponding to the routing interface. Note that IGMP snooping behavior will not be functional on VLANs that are not enabled for VLAN routing.

## 9.12.3 Configuring IGMP and MLD Snooping

### 9.12.3.1 Configuration Commands

The FASTPATH Command Line Interface (CLI) includes several commands that are used to configure the IGMP and MLD snooping features. For more information about each command, and for information about commands that are not described in this section, refer to the FASTPATH CLI Command Reference.

### 9.12.3.2 Enabling IGMP Snooping

To globally enable IGMP snooping on the switch enter Global Configuration mode and use the `set igmp` command, for example:

```
console(config) #set igmp
```

To enable IGMP snooping on an interface, enter Interface Configuration mode and use the `set igmp` command, for example:

```
console(config) #interface 1/0/1
console(config-if-1/0/1) #set igmp
```

To enable IGMP snooping on a VLAN, enter VLAN Config mode and use the `set igmp vlan_id` command. The following example enables IGMP snooping on VLAN 10:

```
console #vlan database
console(config-vlan) #ip igmp 10
```

### 9.12.3.3 Configuring IGMP Snooping Parameters

The following example shows how to configure the group membership interval on an interface (Interface Config mode):

```
console(Interface 1/0/1) #set igmp groupmembership-interval 250
```

The following example shows how to configure the group membership interval on VLAN 10 (VLAN Config mode):

```
console(Vlan) #set igmp groupmembership-interval 10 250
```

The following example shows how to configure the max response interval on an interface (Interface Config mode):

```
console(Interface 1/0/1) #set igmp maxresponse 10
```

The following example shows how to configure the max response interval on VLAN 10 (VLAN Config mode):

```
console(Vlan) #set igmp maxresponse 10 10
```

The following example shows how to enable fast leave mode on VLAN 10 (VLAN Config mode):

```
console(Vlan) #set igmp fast-leave 10
```

The following example shows how to configure the multicast router attached ports expiry interval on an interface (Interface Config mode):

```
console(Interface 1/0/1) #set igmp mcrtrexpiretime 60
```

The following example shows how to configure the multicast router attached ports expiry interval on VLAN 10 (VLAN Config mode):

```
console(config-vlan) #set igmp mcrtrexpiretime 10 60
```

### 9.12.3.4 Display IGMP Snooping Information

The following example shows how to display the IGMP snooping groups:

```
console#show mac-address-table igmpsnooping
```

VLAN ID	MAC Address	Type	Description	Interfaces
1	01:00:5E:01:02:03	Dynamic	Network Assist	Fwd: 1/0/2

The following command shows the forwarding database built by snooping IGMPv3 reports:

```
console#show igmpsnooping ssm entries
```

VLAN ID	Group	Source Ip	Filter Mode	Interfaces
1	232.10.11.12	1.1.1.1	include	1/0/5

The following command displays the IGMPv3 group learned by the snooping switch:

```
console#show igmpsnooping ssm groups
```

VLAN ID	Group	Interface	Reporter	Filter Mode	Source Address List
1	224.10.11.12	1/0/5	192.168.1.1	include	1.1.1.1

### 9.12.3.5 Configuring Static Multicast Forwarding Entries

The following example shows how to create a static multicast forwarding entry for VLAN 1 and multicast MAC address 01:00:5E:11:22:33, associate it with the destination port 1/0/2 and the source port 1/0/4.

```
console(config)#macfilter 01:00:5e:11:22:33 1
```

```
console(config)#interface 1/0/2
```

```
console(Interface 1/0/2)#macfilter adddest 01:00:5e:11:22:33 1
```

```
console(Interface 1/0/2)#exit
```

```
console(config)#interface 1/0/4
```

```
console(Interface 1/0/4)#macfilter addsrc 01:00:5e:11:22:33 1
```

```
console(Interface 1/0/4)#exit
```

```
console#show mac-address-table multicast
```

```

VLAN ID MAC Address          Source  Type   Description      Interface  Fwd
-----
1         01:00:5E:11:22:33 Filter  Static Mgmt Config      Fwd:      Fwd:
                                           1/0/2    1/0/2

```

```
console#show mac-address-table static all
```

```

MAC Address      VLAN ID      Source      Destination
-----
01:00:5E:11:22:33  1           1/0/4      1/0/2

```

```
console#show mac-address-table multicast 01:00:5e:11:22:33 1
```

```

VLAN ID MAC Address          Source  Type   Description      Interface  Fwd
-----
1         01:00:5E:11:22:33 Filter  Static Mgmt Config      Fwd:      Fwd:
                                           1/0/2    1/0/2

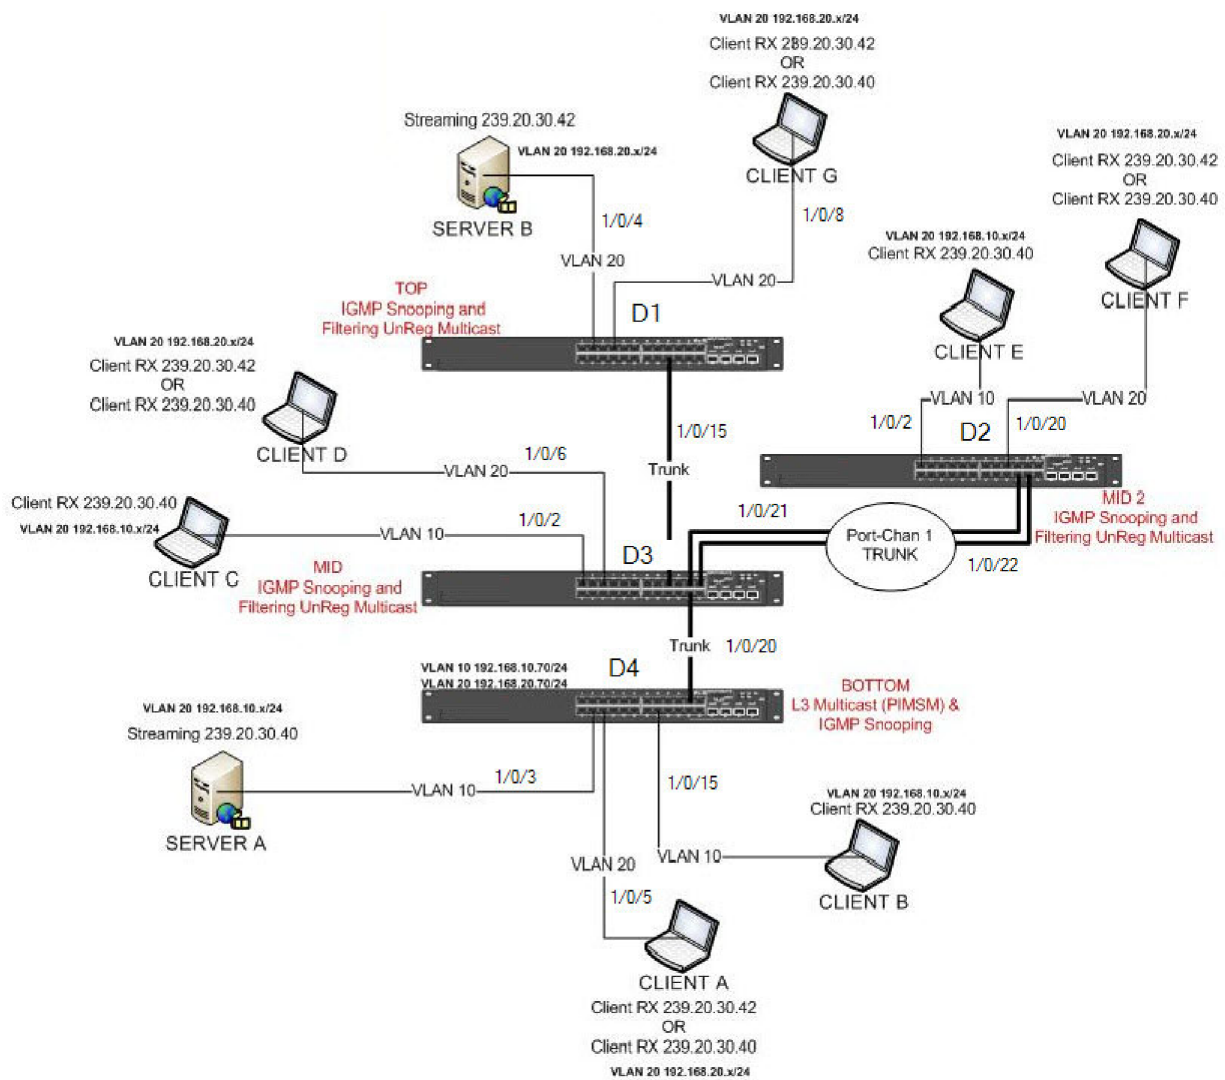
```

### 9.13 Multicast Snooping Example (with IP Multicast Routing)

The examples in this section use the network topology shown in [Figure 444: "Network Topology for Multicast Snooping with IP Multicast Routing,"](#) on page 531.



Figure 444: Network Topology for Multicast Snooping with IP Multicast Routing



The above network topology includes the following elements:

- Snooping Switches: D1, D2, D3 with snooping on VLAN 10, 20
- Multicast Router: D4 with PIM-SM and snooping on VLAN10, 20
- Multicast Listeners: Client A–G
- Multicast Sources: Server A—239.20.30.40, Server B—239.20.30.42
- Subnets: VLAN 10 - 192.168.10.70/24, VLAN 20—192.168.20.70/24
- Router attached ports: D3—1/0/20, D2—PortChannel1, D1-1/0/15

## 9.13.1 Snooping Within a Subnet

In the example network topology, the multicast source and listeners are in the same subnet: VLAN20—192.168.20.70/24. D4 sends periodic queries on VLAN 10, and these queries are forwarded to D1, D2, and D3 via trunk links. Snooping switches D1, D2, and D3 forward these queries to clients G, F, and D respectively.

### 9.13.1.1 Directly Connected Snooping Switch

In this scenario, the multicast source and listener are directly connected to a snooping switch. The following steps show what happens when Client G requests a multicast stream that Server B provides.

1. Client G sends a report for 239.20.30.42.
2. The report is forwarded to multicast router D4 via D1—1/0/15 and D3—1/0/20.
3. A forwarding entry is created by D1 for VLAN20, 239.20.30.42—1/0/8, 1/0/15.
4. Client G receives the multicast stream from Server B.
5. D3 receives the multicast stream and is forwarded to D4 via mrouter port D3-1/0/20.
6. Client D sends a report for 239.20.30.42.
7. The report is forwarded to multicast router D4 via D3—1/0/20.
8. A forwarding entry is created by D3 for VLAN20, 239.20.30.42—1/0/6, 1/0/20.
9. Client D receives the multicast stream from Server B.
10. Client F does not receive the multicast stream because it did not respond to queries from D4.

### 9.13.1.2 Intermediate Snooping Switch

In this scenario, the multicast source and listener are connected by intermediate snooping switches. The following steps show what happens when Client D requests a multicast stream that Server B provides.

1. Client D sends a report for 239.20.30.42.
2. The report is forwarded to multicast router D4 via D3—1/0/20.
3. A forwarding entry is created by D3 for VLAN20, 239.20.30.42—1/0/6, 1/0/20.
4. Client D receives a multicast stream from server B via D1-1/0/15 and D3-1/0/6. D1 forwards an unregistered multicast data stream (239.20.30.42 is unregistered on D1) to mrouter port (D1-1/0/15).
5. Client G will not receive the Server B multicast stream because it did not request it.
6. Client F does not receive the multicast stream because it did not respond to queries from D4.

## 9.13.2 Snooping on a Multicast Router

In the example network topology, consider Client B and Server A. Both are in the same subnet VLAN10—192.168.10.70/24. Server A is a source for multicast stream 239.20.30.40. D4 sends periodic queries on VLAN 10 and VLAN 20, and these queries reach D1, D2, and D3 via trunk links, which in turn forward them in VLAN 10 and VLAN 20 to reach their respective attached clients.

### 9.13.2.1 Multicast Source and Listener on the Same Routing VLAN

In this scenario, the multicast source and listener are directly connected to the multicast router on the same routing VLAN. The following steps show what happens when Server A floods a multicast stream on the routing VLAN that includes Client B.

1. As multicast routing and snooping is enabled on D4 VLAN 10, an IP multicast table entry is created with an empty L2 forwarding list. As a result, multicast traffic is not flooded in VLAN 10.
2. Client B sends a report for 239.20.30.40.
3. The IP multicast table entry is modified to include only D4—1/0/15 as the L2 forwarding list member.
4. Client B receives multicast data.
5. The multicast stream is not forwarded to D3 on trunk link 1/0/20 because no other clients requested this data.

### 9.13.2.2 Multicast Source Connected to Multicast Router and Listener Connected to Snooping Switch (Different Routing VLANs)

In this scenario the multicast source is directly connected to multicast router, and the multicast listener is connected to a different routing VLAN via intermediate snooping switches. The following steps show what happens when Client F requests the multicast stream that Server A provides. Clients A, D and F are in the same subnet: VLAN20—192.168.20.70/24. Server A is in a different subnet: VLAN10—192.168.10.70/24.

1. Client F sends a report for 239.20.30.40.
2. A multicast forwarding entry is created on D2 VLAN20, 239.20.30.40—1/0/20, PortChannel1.
3. The Client F report message is forwarded to D3—PortChannel1 (multicast router attached port).
4. A multicast forwarding entry is created on D3 VLAN 20, 239.20.30.40—PortChannel1, 1/0/20.
5. The Client F report message is forwarded to D4 via D3—1/0/20 (multicast router attached port).
6. An IP multicast routing entry is created on D4 VLAN 10 —VLAN 20 with the L3 outgoing port list as VLAN 20—1/0/20.
7. The multicast stream is routed to D3.
8. The multicast stream is forwarded to listener Client F using forwarding entries created on D3 and D2.
9. Clients A and D do not receive the Server A multicast stream because they did not send a report.

### 9.13.2.3 Multicast Source Connected to Snooping Switch and Listener Connected to Multicast Router (Different Routing VLANs)

In this scenario, the multicast source is connected to a multicast router via intermediate snooping switches, and the listener is directly connected to the multicast router on a different routing interface. The following steps show what happens when Client B requests the multicast stream that Server B provides. Server A and Clients B, C, and E are on the same subnet VLAN10—192.168.10.70/24. Server B is in a different subnet VLAN20 —192.168.20.70/24.

1. Client B sends a report for 239.20.30.42.
2. Multicast Router D4 learns group 239.20.30.42.
3. The multicast stream from Server B reaches D4 via trunk links D1-1/0/15 and D3-1/0/20 as there are mrouter ports and the snooping switch forwards unregistered multicast data to mrouter ports.
4. An IP multicast routing entry is created on D4 VLAN20 - VLAN 10 with the L3 outgoing port list as VLAN10 - 1/0/15.
5. Client B receives multicast data from Server B.
6. Server A and Clients C and E do not receive Server B data because no report messages were sent requesting Server B traffic.

### 9.13.2.4 Multicast Source and Listener Connected to Snooping Switches (Different Routing VLANs)

In this scenario, the multicast source and listener are connected to the multicast router via intermediate snooping switches and are part of different routing VLANs. The following steps show what happens when Client E requests the multicast stream that Server B provides. Clients E, B, and C are on the same subnet VLAN10—192.168.10.70/24. Server B is in a different subnet VLAN20—192.168.20.70/24.

1. Client E sends a report for 239.20.30.42.
2. A multicast forwarding entry is created on D2 VLAN10, 239.20.30.42—1/0/2, PortChannel1.
3. The report from Client E is forwarded to D3 via D2—PortChannel1.
4. A multicast forwarding entry is created on D3 VLAN10, 239.20.30.42—PortChannel1, 1/0/20.
5. The report from Client E is forwarded to D4 via D3—1/0/20.
6. Multicast Router D4 learns group 239.20.30.42.

7. The multicast stream from Server B reaches D4 via trunk links D1-1/0/15 and D3-1/0/20 as there are mrouter ports and a snooping switch forwards unregistered multicast data to mrouter ports.
8. An IP multicast routing entry is created on D4 VLAN20 - VLAN 10 with the L3 outgoing port list as VLAN10—1/0/20.
9. Client E receives multicast data from Server B. Clients B and C do not receive Server B data because no report messages were sent requesting Server B traffic.

## 9.14 Configuring Port Mirroring

Port mirroring is used to monitor the network traffic that a port sends and receives. The Port Mirroring feature creates a copy of the traffic that the source port handles and sends it to a destination port. The source port is the port that is being monitored. The destination port is monitoring the source port. The destination port is where you would connect a network protocol analyzer to learn more about the traffic that is handled by the source port.

A port monitoring session includes one or more source ports that mirror traffic to a single destination port. The FAST-PATH software supports a single port monitoring session. LAGs (port channels) cannot be used as source or destination ports.

For each source port, you can specify whether to mirror ingress traffic (traffic the port receives, or RX), egress traffic (traffic the port sends, or TX), or both ingress and egress traffic.

The packet that is copied to the destination port is in the same format as the original packet on the wire. This means that if the mirror is copying a received packet, the copied packet is VLAN tagged or untagged as it was received on the source port. If the mirror is copying a transmitted packet, the copied packet is VLAN tagged or untagged as it is being transmitted on the source port.

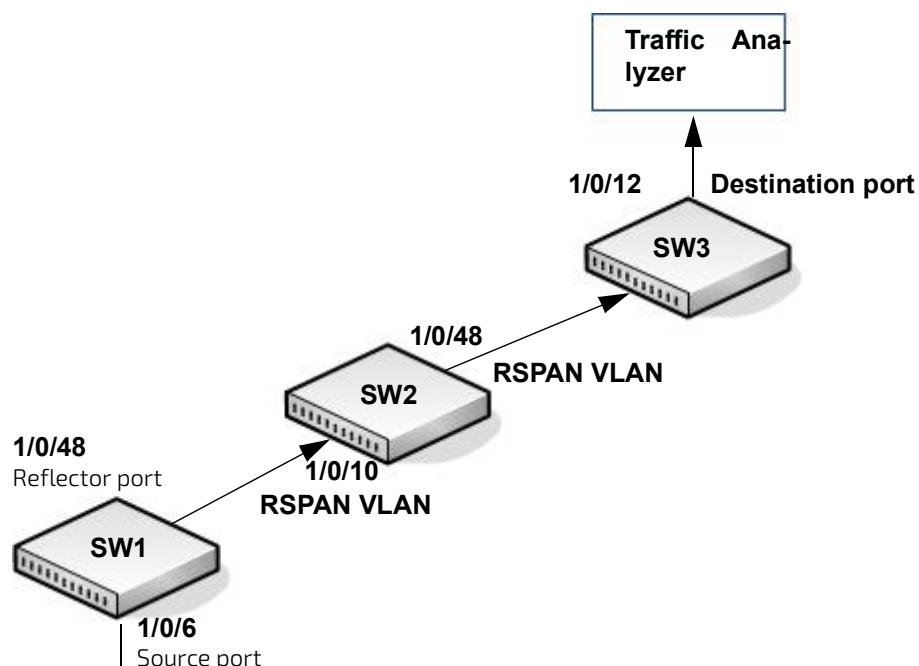
After you configure the port mirroring session, you can enable or disable the administrative mode of the session to start or stop the probe port from receiving mirrored traffic.

## 9.15 Configuring RSPAN

This example mirrors traffic from port 6 on a source switch (SW1) to a probe port on a remote switch (port 12 on SW3). The mirrored traffic is carried in the RSPAN VLAN and VLAN 100, which traverses an intermediate switch (SW2). The steps in this example show how to configure port mirroring on the source, intermediate, and destination switches.

[Figure 445: "RSPAN Configuration Example," on page 534](#) provides a visual overview of the RSPAN configuration example.

Figure 445: RSPAN Configuration Example



## 9.15.1 Configuring RSPAN Using the Web Interface

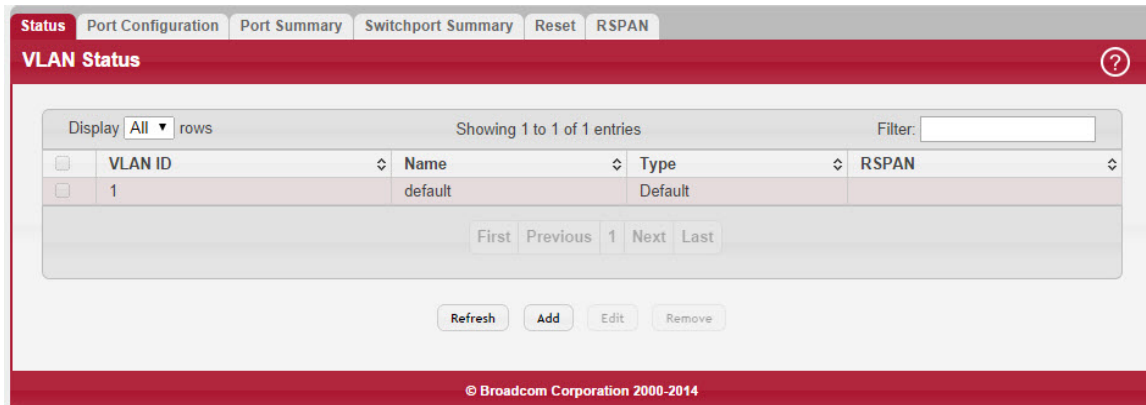
Refer to the following sections to configure RSPAN using the web interface.

### 9.15.1.1 Configuration on the Source Switch (SW1)

To configure the source switch:

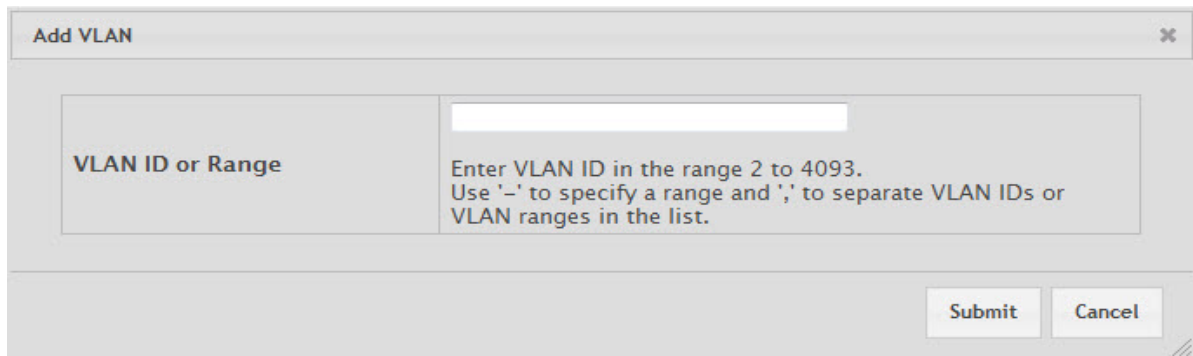
1. Create VLAN100.
  - a. Access the Switching > VLAN > Status page.  
The **VLAN Status** page displays.

Figure 446: VLAN Status



- b. Click Add.  
The **Add VLAN** page displays.

Figure 447: Add VLAN



- c. In the VLAN ID or Range field, type 100.
  - d. Click Submit.
2. Configure VLAN 100 as the RSPAN VLAN.
  - a. Access the Switching > VLAN > RSPAN page.  
The **RSPAN Configuration** page displays.

Figure 448: Configure RSPAN

- b. From the RSPAN VLAN menu, select VLAN 100.
  - c. Click Submit.
3. Configure the RSPAN VLAN as the destination port and the reflector port as port 1/0/48.
    - a. Access the System > Port > Mirroring page.  
The **Multiple Port Mirroring** page displays.

Figure 449: Multiple Port Mirroring

- b. In the Destination field, click the Edit icon.
  - c. From the Port menu, select 1/0/48.
  - d. Click Submit.
4. Configure the source interface port as port 1/0/6.
    - a. Click Configure Source.  
The **Configure Source** page displays.

Figure 450: Configure Source

The screenshot shows the 'Multiple Port Mirroring' configuration page. At the top, there are tabs for 'Summary', 'Description', and 'Mirroring'. Below the tabs, the session details are displayed: Session ID is 1, Mode is Disabled, and Destination is None. A table below shows 'Display All rows' and 'Showing 0 to 0 of 0 entries'. The table has columns for 'Source' and 'Direction', but it is empty. At the bottom, there are buttons for 'Refresh', 'Configure Session', 'Configure Source' (circled in red), and 'Remove Source'. The footer indicates '© Broadcom Corporation 2000-2013'.

- b. From Type, select Interface.
  - c. In Available Source Port, select 1/0/6.
  - d. Click Submit.
5. Enable the port mirroring session.
    - a. Click Configure Session.  
The **Configure Session** page displays.

Figure 451: Configure Session

The screenshot shows the 'Multiple Port Mirroring' configuration page. At the top, there are tabs for 'Summary', 'Description', and 'Mirroring'. Below the tabs, the session details are displayed: Session ID is 1, Mode is Disabled, and Destination is None. A table below shows 'Display All rows' and 'Showing 0 to 0 of 0 entries'. The table has columns for 'Source' and 'Direction', but it is empty. At the bottom, there are buttons for 'Refresh', 'Configure Session' (circled in red), 'Configure Source', and 'Remove Source'. The footer indicates '© Broadcom Corporation 2000-2013'.

- b. In Mode, select Enable.
- c. Click Submit.

### 9.15.1.2 Configuration on the Intermediate Switch (SW2)

To configure the intermediate switch:

1. Create VLAN 100.
  - a. Access the Switching > VLAN > Status page.  
The **VLAN Status** page displays (see [Figure 446: "VLAN Status,"](#) on page 535).
  - b. Click Add.  
The **Add VLAN** page displays (see [Figure 447: "Add VLAN,"](#) on page 535).
  - c. In the VLAN ID or Range field, type 100.
  - d. Click Submit.

2. Configure VLAN 100 as the RSPAN VLAN.
  - a. Access the Switching > VLAN > RSPAN page.  
The **RSPAN Configuration** page displays (see [Figure 448: "Configure RSPAN," on page 536](#)).
  - b. From the RSPAN VLAN menu, select VLAN 100.
  - c. Click Submit.
3. Configure ports 1/0/10 and 1/0/48 as members of VLAN 100, and enable tagging so that frames transmitted in this VLAN will include the VLAN 100 tag in the Ethernet header.
  - a. Access the Switching > VLAN > Port Configuration page.  
The **VLAN Port Configuration** page displays.

Figure 452: VLAN Port Configuration

The screenshot shows the 'VLAN Port Configuration' page for VLAN ID 100. The page has a navigation bar with tabs: Status, Port Configuration (selected), Port Summary, Internal Usage, Reset, and RSPAN. Below the navigation bar, the 'VLAN ID' is set to 100. The main content area displays a table with 10 rows, showing 1 to 10 of 116 entries. The table has columns for Interface, Status, Participation, and Tagging. The interface 1/0/10 is selected, and its status is 'Include' with 'Include' participation and 'Tagged' tagging. Below the table are navigation buttons: First, Previous, 1, 2, 3, 4, 5, Next, Last. At the bottom, there are buttons for Refresh, Edit, and Edit All.

Interface	Status	Participation	Tagging
1/0/1	Exclude	Auto Detect	Untagged
1/0/2	Exclude	Auto Detect	Untagged
1/0/3	Exclude	Auto Detect	Untagged
1/0/4	Exclude	nil	Untagged
1/0/5	Include	Include	Tagged
1/0/6	Exclude	nil	Untagged
1/0/7	Exclude	Auto Detect	Untagged
1/0/8	Exclude	nil	Untagged
1/0/9	Exclude	Auto Detect	Untagged
1/0/10	Include	Include	Tagged

- b. From VLAN ID, select 100.
- c. Select interfaces 1/0/10 and 1/0/48.
- d. Click Edit.  
The **Edit VLAN Port Configuration** page displays.



Figure 453: Edit VLAN Port Configuration

**NOTE: Tagging will only be enabled for VLAN member ports.**

VLAN ID	100
Interface	1/0/48
Participation	<input checked="" type="radio"/> Include <input type="radio"/> Auto Detect <input type="radio"/> Exclude
Tagging	<input type="radio"/> Untagged <input checked="" type="radio"/> Tagged

Submit Cancel

- e. In the Participation field, select Include.
- f. In Tagging, select Tagged.

### 9.15.1.3 Configuration on the Destination Switch (SW3)

To configure the destination switch:

1. Create VLAN 100.
  - a. Access the Switching > VLAN > Status page.  
The **VLAN Status** page displays (see [Figure 446: "VLAN Status," on page 535](#)).
  - b. Click Add.  
The **Add VLAN** page displays (see [Figure 447: "Add VLAN," on page 535](#)).
  - c. In the VLAN ID or Range field, type 100.
  - d. Click Submit.
2. Configure VLAN 100 as the RSPAN VLAN.
  - a. Access the Switching > VLAN > RSPAN page.  
The **RSPAN Configuration** page displays (see [Figure 448: "Configure RSPAN," on page 536](#)).
  - b. From the RSPAN VLAN menu, select VLAN 100.
  - c. Click Submit.
3. Configure 0/12 as the destination (probe) port.
 

The **Multiple Port Mirroring** page displays (see [Figure 449: "Multiple Port Mirroring," on page 536](#)).

  - a. In the Destination field, click the Edit icon.
  - b. In the Type field, select Interface.
  - c. From the Port menu, select 0/12
  - d. Click Submit.
4. Configure the source interface port as port 0/6.
  - a. Click Configure Source.  
The **Configure Source** page displays (see [Figure 450: "Configure Source," on page 537](#)).
  - b. From Type, select Remote VLAN.
  - c. Click Submit.

5. Enable the port mirroring session.

a. Click Configure Session.

The **Configure Session** page displays (see [Figure 451: "Configure Session," on page 537](#)).

b. In Mode, select Enable

c. Click Submit.

## 9.15.2 Configuring RSPAN Using the CLI

Refer to the following sections to configure RSPAN using the CLI interface.

### 9.15.2.1 Configuration on the Source Switch (SW1)

To configure the source switch:

1. Access the VLAN configuration mode and create VLAN 100, which will be the RSPAN VLAN.

```
(Routing) #vlan database
(Routing) (Vlan) #vlan 100
(Routing) (Vlan) #exit
```

2. Configure VLAN 100 as the RSPAN VLAN.

```
(Routing) #configure
(Routing) (Config) #vlan 100
(Routing) (Config) (vlan 100) #remote-span
(Routing) (Config) (vlan 100) #exit
```

3. Configure the RSPAN VLAN as the destination port and the reflector port as port 0/48.

```
(Routing) (Config) #monitor session 1 destination remote vlan 100 reflector-port 0/48
```

4. Configure the source interface port as port 0/6.

```
(Routing) (Config) #monitor session 1 source interface 0/6
```

5. Enable the port mirroring session on the switch.

```
(Routing) (Config) #monitor session 1 mode
(Routing) #exit
```

Configuration on the Intermediate Switch (SW2)

To configure the intermediate switch (SW2):

1. Access the VLAN configuration mode and create VLAN 100.

```
(Routing) #vlan database
(Routing) (Vlan) #vlan 100
(Routing) (Vlan) #exit
```

2. Enable RSPAN on vlan 100.

```
(Routing) #configure
(Routing) (Config) #vlan 100
(Routing) (Config) (vlan 100) #remote-span
(Routing) (Config) (vlan 100) #exit
```

3. Configure VLAN participation so that the interface is always a member of the VLAN.

```
(Routing) (Config) #vlan participation include 100
(Routing) (Config) #interface 0/10
```

4. Enable VLAN tagging on the interface.

```
(Routing) (Config) #vlan tagging 100
(Routing) (Config) #exit
```

5. Configure VLAN participation so the interface is always a member of the VLAN.(Routing) (Config) #vlan participation

```
include 100
(Routing) (Config) #interface 0/48
(Routing) (Config) #exit
```

### 9.15.2.2 Configuration on the Destination Switch (SW2)

To configure the destination switch (SW3):

1. Access the VLAN configuration mode and create VLAN 100.

```
(Routing) #vlan database
(Routing) (Vlan) #vlan 100
(Routing) (Vlan) #exit
```

2. Enable RSPAN on vlan 100.

```
(Routing) #configure
(Routing) (Config) #vlan 100
(Routing) (Config) (vlan 100) #remote-span
(Routing) (Config) (vlan 100) #exit
```

3. Configure the RSPAN VLAN as the source interface for the port mirroring session.

```
(Routing) #configure
(Routing) (Config) #monitor session 1 source remote vlan 100
```

4. Configure the destination port as port 0/12. This is the probe port that is attached to a network traffic analyzer.

```
(Routing) (Config) #monitor session 1 destination interface 0/12
```

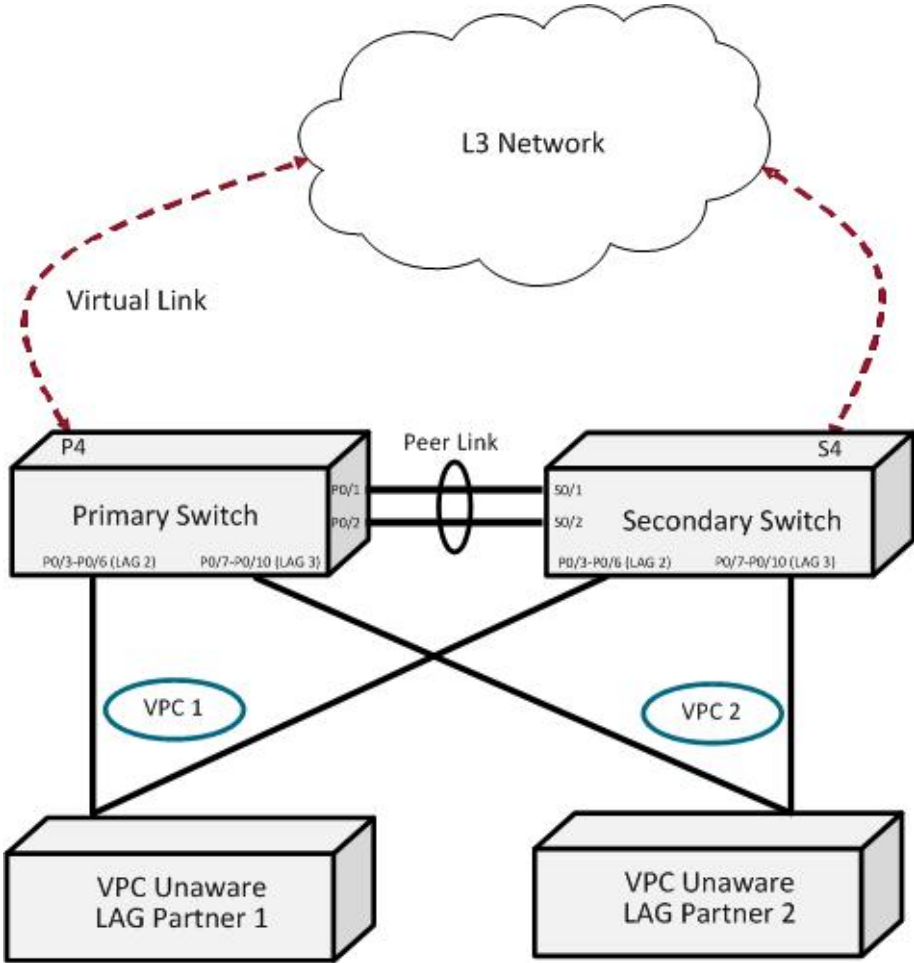
5. Enable the port mirroring session on the switch.

```
(Routing) (Config) #monitor session 1 mode
(Routing) (Config) #exit
```

## 9.16 Configuring VPC

See [Figure 454: "VPC Configuration Diagram," on page 542](#) for a visual overview of the VPC configuration steps.

Figure 454: VPC Configuration Diagram



### 9.16.1 Configuring VPC Using the Web Interface

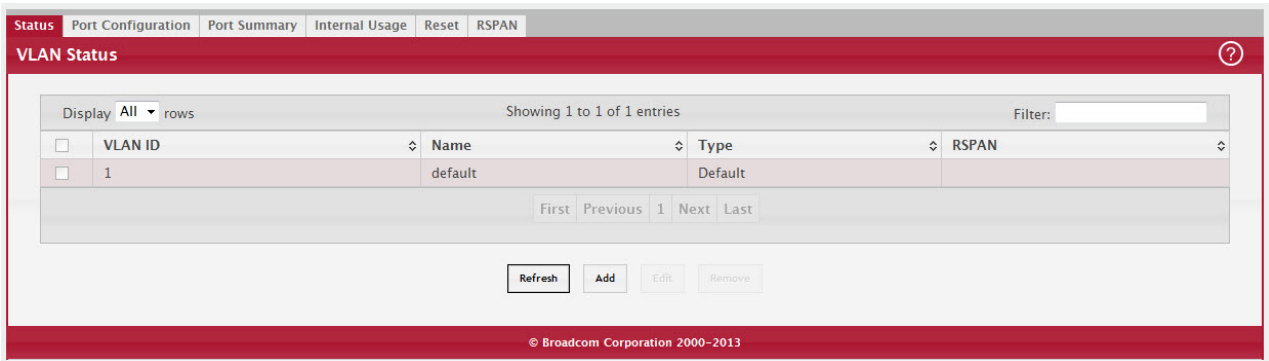
Refer to the following sections to configure VPC using the web interface.

#### 9.16.1.1 Configuring the VLANs and Port Channels

Before you configure the VPC global settings, you must first configure the system VLANs and port channels.

1. To create the VPC VLANs, from the web interface, click Switching > VLAN > Status in the navigation menu.

Figure 455: VLAN Status



2. Click Add to create the VLANs.

Figure 456: VLAN Status

3. In the VLAN ID or Range field, enter 10-17 and click Submit.
4. Click Add to create the VLAN routing interface that will be used for the Dual Control Plane detection Protocol.
5. In the VLAN ID or Range field, enter 100 and click Submit.
6. To modify Port Channels 1, 2, and 3, click Switching > Port Channel> Summary in the navigation menu.

Figure 457: Port Channel Summary

Name	Type	Admin Mode	STP Mode	Link State	Link Trap	Members	Active Ports	Load Balance
ch1	Static	Enable	Disable	Down	Disable	0/1, 0/2		Source/Destination MAC, VLAN, EtherType, Incoming Port
ch2	Static	Enable	Enable	Down	Disable	0/3, 0/4, 0/5, 0/6		Source/Destination MAC, VLAN, EtherType, Incoming Port
ch3	Static	Enable	Enable	Down	Disable	0/7, 0/8, 0/9, 0/10		Source/Destination MAC, VLAN, EtherType, Incoming Port
ch4	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch5	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch6	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch7	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch8	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch9	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port
ch10	Static	Enable	Enable	Down	Disable			Source/Destination MAC, VLAN, EtherType, Incoming Port

- a. Edit CH1 to include 0/1 and 0/2.
- b. Edit CH1 to include 0/3, 0/4, 0/5, and 0/6.
- c. Edit CH1 to include 0/7, 0/8, 0/9, and 0/10.

### 9.16.1.2 Configuring the VPC Global Settings

To configure the VPC global settings:

1. Click Switching > Virtual Port Channel > Global.

Figure 458: VPC Global Configuration

Virtual Port Channel Global Configuration	
Domain ID	1 (1 to 1)
VPC Mode	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
VPC Operational Mode	Disabled
VPC State	Disabled
Self Role	None
Peer Role	None
System MAC	00:10:18:00:00:02
Keepalive Parameters	
Keepalive Priority	10 (1 to 255)
Keepalive Timeout (Seconds)	5 (2 to 15)
Keepalive Mode	<input checked="" type="radio"/> Disable <input type="radio"/> Enable
Peer Link	
Port Channel	3/1 <input type="button" value="Edit"/>
Peer Link Status	Down
Peer Link STP Mode	Disabled
Configured VLANs	1, 10, 11, 12, 13, 14, 15, 16, 17, 100
Egress Tagging	1, 10, 11, 12, 13, 14, 15, 16, 17, 100
Peer Detection	
Peer Detected	False
Peer IP Address	192.168.0.1 (x.x.x.x)
Source IP Address	192.168.0.2 (x.x.x.x)
UDP Port	50000 (1 to 65535)
Peer Detection Mode	<input type="radio"/> Disable <input checked="" type="radio"/> Enable
<input type="button" value="Submit"/> <input type="button" value="Refresh"/> <input type="button" value="Cancel"/>	

© Broadcom Corporation 2000-2013

2. Enter 1 for the Domain ID.
3. Select Enable for the VPC Mode.
4. Enter 10 for the Keepalive Priority.
5. Click Edit in the Peer Link section to select the Peer Link port.
6. Enter 192.168.0.1 as the Peer IP Address. This configures the IP address of the peer VPC switch. This configuration is used by the dual control plane detection protocol (DCPDP) on the VPC switches. The UDP port on which the VPC switch listens to the DCPDP messages can also be configured with this command. The configurable range for the UDP port is 1 to 65535 (Default is 60000).
7. Enter 192.168.0.2 as the Source IP Address. This is the address used by DCPDP.
8. Click Submit.

### 9.16.1.3 Configuring the VPC Interface

To configure the VPC interface settings:

1. Click Switching > Virtual Port Channel > Interface Configuration.

Figure 459: VPC Interface Configuration

Display	All	rows	Showing 1 to 2 of 2 entries			Filter:
<input checked="" type="checkbox"/>	Interface	VPC Identifier	Operational Mode	VPC Interface State	Configured VLANs	
<input checked="" type="checkbox"/>	3/2	2	Disabled	Disabled	1	
<input checked="" type="checkbox"/>	3/3	1	Disabled	Disabled	1	

First Previous 1 Next Last

Refresh Add Details Remove

© Broadcom Corporation 2000-2013

- Click Add to add both Ch2 (3/2) and Ch3 (3/3).

## 9.16.2 Configuring VPC Using the CLI

To configure VPC using the CLI interface:

- Enter VLAN data base mode and create the VPC VLANs.

```
(Broadcom FASTPATH routing) #vlan database
(Broadcom FASTPATH routing) (Vlan) #vlan 10-17
```

- Create the VLAN routing interface that will be used for the Dual Control Plane detection Protocol.

```
(Broadcom FASTPATH routing) (Vlan) #vlan 100
(Broadcom FASTPATH routing) (Vlan) #vlan routing 100
(Broadcom FASTPATH routing) (Vlan) #exit
```

- Enable the VPC feature.

```
(Broadcom FASTPATH routing) #config
(Broadcom FASTPATH routing) (Config) #feature vpc
```

- Enable the keepalive protocol.

```
(Broadcom FASTPATH routing) #config
(Broadcom FASTPATH routing) (Config) #vpc domain 1
(Broadcom FASTPATH routing) (Config-VPC 1) #peer-keepalive enable
(Broadcom FASTPATH routing) (Config-VPC 1)#
```

- Configure the VPC role priority.

```
(Broadcom FASTPATH routing) (Config) #vpc domain 1
(Broadcom FASTPATH routing) (config-VPC 1) #role priority 10
```

- Create LAG1.

```
(Broadcom FASTPATH routing) (Config) #interface lag 1
(Broadcom FASTPATH routing) (Interface lag 1) #description "VPC-Peer-Link"
```

- Disable spanning tree on the LAG.

```
(Broadcom FASTPATH Routing) (Interface lag 1) #no spanning-tree port mode
```

8. Allow the LAG to participate in all VLANs and accept and send tagged frames only. This is similar to configuring a port in trunk mode.

```
(Broadcom FASTPATH Routing) (Interface lag 1) #vlan participation include 1-99
(Broadcom FASTPATH Routing) (Interface lag 1) #vlan tagging 1-99
(Broadcom FASTPATH Routing) (Interface lag 1) #vlan acceptframe vlanonly
(Broadcom FASTPATH Routing) (Interface lag 1) #vpc peer-link
(Broadcom FASTPATH Routing) (Interface lag 1) #exit
```

9. Create the peer link.

```
(Broadcom FASTPATH routing) (Config) #interface 0/1-0/2
(Broadcom FASTPATH routing) (Interface 0/1-0/2) #addport lag 1
(Broadcom FASTPATH routing) (Interface 0/1-0/2) #description "VPC-Peer-Link"
```

10. Enable UDLD (if required).

```
(Broadcom FASTPATH routing) (Interface 0/1-0/2) #udld enable
(Broadcom FASTPATH routing) (Interface 0/1-0/2) #udld port aggressive
(Broadcom FASTPATH routing) (Interface 0/1-0/2) #exit
```

11. Configure Dual Control Plane detection Protocol Configuration (if required):

- a. Configure a VLAN routing interface and assign a local IP address (independent from the peer address).

```
(Broadcom FASTPATH Routing) (Config) #interface vlan 100
```

- b. Configure the peer-switch IP address (the destination IP address). This command configures the IP address of the peer VPC switch. This configuration is used by the dual control plane detection protocol (DCPDP) on the VPC switches. The UDP port on which the VPC switch listens to the DCPDP messages can also be configured with this command. The configurable range for the UDP port 1 to 65535 (Default is 60000).

```
(Broadcom FASTPATH Routing) (Interface vlan 100) #ip address 192.168.0.2 255.255.255.0
```

Example: 192.168.0.2 255.255.255.0 for the IP address and subnet mask.

```
(Broadcom FASTPATH Routing) (Interface vlan 100) #exit
```

- c. Configure the keepalive source and destination IP address.

```
(Broadcom FASTPATH Routing) #config
(Broadcom FASTPATH Routing) #vpc domain 1
(Broadcom FASTPATH Routing) (Config-VPC 1) #peer-keepalive destination 192.168.0.1 source
192.168.0.2
```

- d. Enable Peer Detection mode. The mode starts running if VPC is globally enabled.

```
(Broadcom FASTPATH Routing) (Config-VPC 1) #peer detection enable
```

12. Configure a port-channel as VPC interface. The configurable range for the VPC ID is 1 to L7\_MAX\_NUM\_VPC.

```
(Broadcom FASTPATH Routing) (Config) #interface 0/3-0/6
(Broadcom FASTPATH Routing) (Interface 0/3-0/6) #addport lag 2
(Broadcom FASTPATH Routing) (Interface 0/3-0/6) #exit
```

```
(Broadcom FASTPATH Routing) (Config) #interface 0/7-0/10
(Broadcom FASTPATH Routing) (Interface 0/7-0/10) #addport lag 3
(Broadcom FASTPATH Routing) (Interface 0/7-0/10) #exit
```

```
(Broadcom FASTPATH Routing) (Config) #interface lag 2
(Broadcom FASTPATH Routing) (Interface lag 2) #vlan participation include 1-100
(Broadcom FASTPATH Routing) (Interface lag 2) #vlan tagging 1-100
(Broadcom FASTPATH Routing) (Interface lag 2) #vlan acceptframe vlanonly
(Broadcom FASTPATH Routing) (Interface lag 2) #vpc 1
(Broadcom FASTPATH Routing) (Interface lag 2) #exit
```



```
(Broadcom FASTPATH Routing) (Config)#interface lag 3
(Broadcom FASTPATH Routing) (Interface lag 3) #vlan participation include 1-100
(Broadcom FASTPATH Routing) (Interface lag 3) #vlan tagging 1-100
(Broadcom FASTPATH Routing) (Interface lag 3) #vlan acceptframe vlanonly
(Broadcom FASTPATH Routing) (Interface lag 3) #vpc 2
(Broadcom FASTPATH Routing) (Interface lag 3) #exit
```

The administrator must ensure that the port channel configurations on both devices are in sync before enabling VPC. After the VPC interfaces are enabled, the VPC interfaces are operationally shut down. The VPC component exchanges information regarding the port members that constitute the port-channel on each device. Once this information is populated on both devices, the VPC interfaces are operationally up and traffic forwarding on VPC interfaces is allowed. Port-channels must be configured on both devices as VPC interfaces for the VPC interface to be enabled. Also, the port-channel-number:VPC-Id pair must be the same on both the primary and secondary devices.

Member ports can be added or removed from the VPC interface. If a port is added as a port member to a VPC interface, the Primary allows the port member if the maximum criteria is satisfied. When a port member is removed from the VPC interface, the Primary decides if the minimum criteria is satisfied. If it is not, it will shut down the VPC interface on both the devices. Shutting down the VPC interface on the Secondary is not allowed. The VPC interface can only be shut down on the Primary.

The secondary switch forwards all BPDUs/LACPDUs received on the port members of the VPC interface to the primary over the Peer-Link. Events related to VPC interface and their port members are forwarded to the primary switch for handling. FDB entries learned on VPC interfaces are synced between the two devices. In the case where all VPC member ports are UP, data traffic does not traverse the peer link.

## 9.17 Virtual Routing and Forwarding Lite Operation and Configuration

### NOTICE

Virtual Routing and Forwarding (VRF) configuration can be performed only by using the CLI. Web UI and SNMP configuration options are not supported for the VRF feature.

### 9.17.1 Overview

The VRF feature enables a router to function as multiple routers. Each virtual router manages its own routing domain, with its own IP routes, routing interfaces, and host entries. Each virtual router makes its own routing decisions, independent of other virtual routers. More than one virtual routing table may contain a route to a given destination. The network administrator can configure a subset of the router's interfaces to be associated with each virtual router. The router routes packets according to the virtual routing table associated with the packet's ingress interface. Each interface can be associated with at most one virtual router.

### 9.17.2 VRF Functionality

Each virtual router behaves like an independent router. Virtual routers can be created and destroyed dynamically. The fault domains of virtual routers are isolated. Bringing down a virtual router does not impact another virtual router. Each virtual router has its own instances of routing protocols and routing applications. FASTPATH supports a maximum of 64 virtual routers. The total number of routes or host entries is still limited by the hardware capacities on the physical router, but the routes and host entries are distributed across the virtual routing domains based on the user configuration. The maximum number routes in a particular virtual router can be optionally reserved.

IP prefixes can overlap between two VR instances. The same IP address can be configured on two interfaces that are a part of different VR instances. A packet is routed based on the route table look up result in the corresponding VR instance. The VR instance is derived based on the ingress interface. There are situations, however, that require support for inter-VR routing, such as providing access to shared services syslog server, DHCP server, the Internet, etc. These cases are handled through route leaking.

In the standard FASTPATH Routing build, the VRF component must be selected to support VRF. By default, all of the standard routing software and functions are in the default router (VRID 0), which is created on startup and cannot be deleted by the user. The non-VRF routing user does not experience any disruption in using the CLI commands or in router functionality as a result of VRF configuration. Configuration migration for a system running an earlier build is supported.

The FASTPATH VRF feature depends on the Network Name Space feature in Linux. FASTPATH supports this feature in the 3.x and later Linux kernels. There is no impact on the routing feature for FASTPATH running pre-3.x kernels except that the VRF feature is not supported on them. The CLI commands for VRF are disabled in the FASTPATH builds running pre-3.x kernels.

The user manages the VRF functionality through CLI commands. There is no separate user interface for every VR instance. The user manages all the VR instances from a single CLI. The in-band management is supported through one of the interfaces on the default VR only. The FASTPATH CLI does not currently support managing VRF instances, although the CLI commands work in the default VR instance. Syslog is enhanced to support logging from different Linux processes. VRF supports logging for all the events that are already supported.

### 9.17.2.1 Route Leaking

Route leaking is the ability to install a route in one VRF that allows traffic to flow to another VRF. Although this mechanism breaks the isolation between VRFs, it is sometimes used to provide access to common services for devices inside the different VRFs. FASTPATH supports route leaking between the global default routing table and a VR, but not across VRs. FASTPATH supports route leaking only through static routes. FASTPATH does not support inter-VRF packet forwarding by connecting a wire between ports belonging to different VR instances.

#### 9.17.2.1.1 Adding Leaked Routes

Connected routes in one router that are leaked into another VR are referred to as leaked host routes. To add leaked host routes, specify the next-hop interface but not the next-hop address. For leaked routes that are not directly connected (static or dynamic routes), the next-hop address must be specified in addition to the next-hop interface. The next-hop interface is specified to identify the outgoing VR interface. If the next-hop interface is unspecified, the route is treated as an internal route to the VR.

Internal routes within a router that are added with only a next-hop interface value (and no next-hop address value) are supported only over unnumbered interfaces.

#### 9.17.2.1.2 Using Leaked Routes

The line rate forwarding continues to work the same for leaked route destinations in a router as for the internal routes in the router. For bidirectional traffic to work between VRs using leaked routes, the corresponding routes should be leaked between the VRs.

#### 9.17.2.1.3 CPU-Originated Traffic

For CPU-originated traffic from different applications (ping, traceroute, syslog, IP helper) that may use the leaked routes to access the destination or shared service, the following conditions are required to ensure proper operation:

1. The source IP address in the originated packets must be mentioned with the source IP option (e.g., ping with source option).
2. In the router where the CPU traffic originates, the route for the source option matching network must be leaked into the virtual router where the next-hop belongs so that the return traffic is directed to the traffic-originating router.

### 9.17.3 VRF and FASTPATH Feature Support

The following table lists FASTPATH features and details how they are supported by VRF Lite.

**Table 408: VRF and FASTPATH Feature Support**

Feature	VRF Support
Network Management	<p>Network management includes the ability to manage the switch using the CLI and SNMP. FASTPATH Network management is supported only via the default router. Administrators cannot log into the switch and manage the switch via one of the IP addresses on the non-default VR.</p> <p>The Service Port and the Network Port are always associated with the default router, so the customers are able to manage the switch via these interfaces.</p>
SNMP Management	Only the default router can be managed via SNMP.
AAA	The Authentication, Authorization, and Accounting protocols include services, such as the RADIUS client and the TACACS+ client. FASTPATH supports these services only on the default router.
Network Services	The Ping and the Trace Route clients are supported in the virtual router context. Other protocols are supported only in the default router. These include the SNTP client, DNS client, sFlow, RPCAP, and Auto Install.
Loopback and Tunnel Interfaces	<p>Loopback interfaces with IPv4 prefixes are supported in the virtual router. Loopback interfaces with IPv6 addresses can be configured only in the default router.</p> <p>The number of Loopback interfaces in builds containing the VRF package is increased to 64. The loopback interfaces are shared across VR instances in the system and there is no restriction on the maximum supported per VR.</p> <p>Tunnel interfaces are not supported in the virtual router.</p>
IP unnumbered interfaces	IP unnumbered interface cannot be part of non-default VRF instance. This feature is supported only in the default router.
OSPFv2	The OSPFv2 protocol is supported in the virtual router. As of the current release, a crash in the OSPFv2 protocol does not cause the switch to reboot. All OSPF features including graceful restart and NSF are supported for OSPFv2 in each VR instance.
OSPFV3	The OSPFv3 protocol is supported only in the default router.
RIP	RIP is not currently supported in the virtual router.
VRRP	<p>The Virtual Routing Redundancy Protocol is a fault-tolerance feature that enables two or more routers to appear as one router to the IP clients. If one of the VRRP routers fails, another router can take over the data forwarding with minimum interruption to client traffic.</p> <p>The VRRP protocol is supported in the virtual router context. The VRRP protocol enables two or more virtual routers running on different physical switches to form a VRRP group. The virtual routers running on the same physical switch cannot form a VRRP group with each other.</p>
BGP	<p>The Border Gateway Protocol is intended to be used by the Customer Edge (CE) switch to communicate with other CE switches and PE switches across the Provider Network. This typical VRF-Lite deployment is described in <a href="#">Section 9.17.4: "VRF Lite Deployment Scenarios"</a>. The BGP protocol runs in the Default Router context and is aware of the virtual routers. BGP is used to:</p> <ul style="list-style-type: none"> <li>• Redistribute VPN routes from virtual routers on the CE switch to the attached PE in the Provider Network.</li> <li>• Leak routes dynamically between different virtual routers on the same physical switch. This requires support for BGP extended communities and route targets.</li> </ul> <p>In the current FASTPATH implementation, BGP does not support either of the above mentioned functionalities.</p>
IPv6	The current FASTPATH release supports VRF-Lite only for IPv4. IPv6 data forwarding and protocols are not currently supported.
IP Multicast	The current FASTPATH Virtual Routing release supports only IPv4 unicast routing.

Table 408: VRF and FASTPATH Feature Support (Continued)

Feature	VRF Support
Policy Based Routing	PBR is a routing policy feature useful in overriding routing decisions with programmable rules. PBR is supported only in the default router in the current release.
DHCP Server	DHCP Server is not VR-aware in the current release.
DHCP Snooping	The IP Source Guard (IPSG) feature uses DHCP snooping to allow only packets from known sources. IPSG uses DHCP Snooping to snoop the DHCP addresses allocated to connected hosts. The tuple (IP, MAC, VLAN, Interface) uniquely identifies a host. DHCP Snooping is a layer-2 feature and is VRF-agnostic. It works in layer-2 of any VLAN irrespective of whether it belongs to a default router or any virtual router. It applies to all protocols working at L2.
IP Helper	IP Helper relays the broadcast packets received on a Routing interface in the VRF context to the configured server address. The server is looked up in the RTO specific to that VR only. Relay across VRs is not supported.
OpEN API	The applications using existing OpEN APIs are not affected by the VRF feature.
Layer-2 Features	The VRF feature does not affect the switch layer-2 features such as virtual port channels (VPC). However, if VPC is planned to be used on VRF-enabled switches, the VPC ports need to be configured to be in the same routing domain.

### 9.17.4 VRF Lite Deployment Scenarios

The following are two likely deployment scenarios for the VRF-Lite solution:

1. In the Customer edge (CE) devices that interface with the PE (Provider edge) device in the service provider backbone network to provide VPN connectivity for the Enterprise network sites spread across different geographical locations across the internet backbone.

In this scenario, the BGP protocol must be running on the device to support feature extensions required to support:

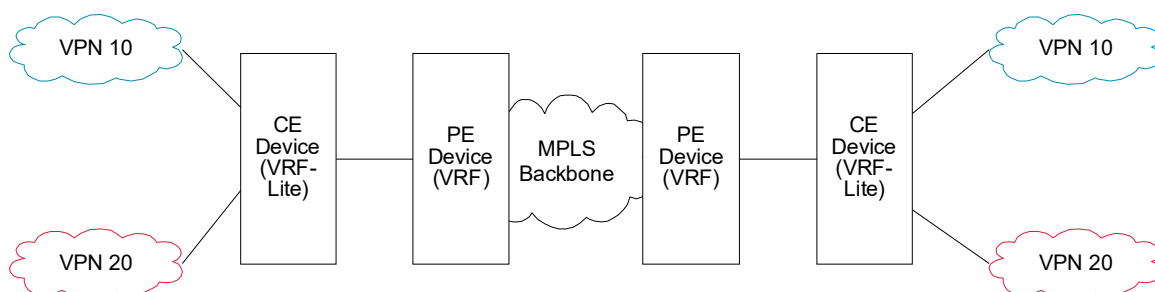
- a. Dynamic route leaking locally between the VRFs to leak the routes to shared services using Route Targets.
- b. Exchange the VPN related route information per VR with PE device using Extended communities.

2. The internal Routers in the Enterprise networks to provide isolation of different departments/offices at layer-3 or routing domain.

This scenario does not mandate that the BGP protocol be running on the device. It can still be run in this scenario to achieve dynamic route leaking only. The IGP protocol (OSPF or RIP) running in the VR instance communicates route information with corresponding peers in the same VR on other CE devices or internal Routers.

These scenarios are shown in [Figure 460: "VRF Scenarios," on page 550](#).

Figure 460: VRF Scenarios

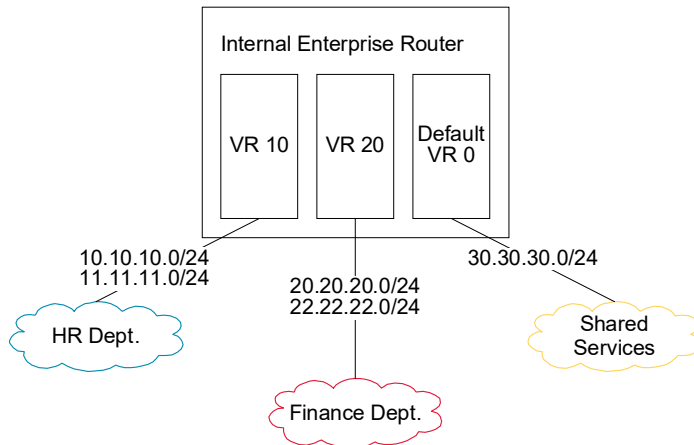


The default global routing table is also referred to as VR 0.

In the following example, subnetworks 10.10.10.0/24 and 11.11.11.0/24 belong to the virtual routing domain HR Dept and subnetworks 20.20.20.0/24 and 22.22.22.0/24 belong to virtual routing domain Finance Dept. Hence, the hosts in networks 10.10.10.0/24 can communicate only with other network 11.11.11.0/24 via the router and the hosts in networks 20.20.20.0/24 can communicate only with other network 22.22.22.0/24 via the router.

If there is a shared service printer @30.30.30.30 in the default global routing domain Shared Services, we would want the HR and Finance domains to have access to it. Therefore, we statically leak a 30.30.30.0/24 route from global routing table to VR 10 and VR 20. At the same time, we statically leak the routes 10.10.10.0/24 and 11.11.11.0/24 from VR 10 to global table (the same applies to VR 20).

Figure 461: VRF Routing With Shared Services



The route tables in both the VRs and the global domain look like the following:

```
(Router) #show ip route vrf HR
Route Codes: R - RIP Derived, O - OSPF Derived, C - Connected, S - Static
             B - BGP Derived, IA - OSPF Inter Area
             E1 - OSPF External Type 1, E2 - OSPF External Type 2
             N1 - OSPF NSSA External Type 1, N2 - OSPF NSSA External Type 2
             L - Leaked Route
C           10.10.10.0/24 [0/1] directly connected,   vlan 10
C           11.11.11.0/24 [0/1] directly connected,   vlan 11
S L        30.30.30.0/24 [1/1] directly connected,   vlan 30
S L        50.50.50.0/24 [1/1] via 30.30.30.2,     02d:22h:15m,   vlan 30
```

```
(Router) #show ip route vrf Finance
Route Codes: R - RIP Derived, O - OSPF Derived, C - Connected, S - Static
             B - BGP Derived, IA - OSPF Inter Area
             E1 - OSPF External Type 1, E2 - OSPF External Type 2
             N1 - OSPF NSSA External Type 1, N2 - OSPF NSSA External Type 2
             L - Leaked Route
C           20.20.20.0/24 [0/1] directly connected,   vlan 20
C           22.22.22.0/24 [0/1] directly connected,   vlan 22
S L        30.30.30.0/24 [1/1] directly connected,   vlan 30
S L        50.50.50.0/24 [1/1] via 30.30.30.2,     02d:22h:15m,   vlan 30
```

```
(Router) #show ip route
Route Codes: R - RIP Derived, O - OSPF Derived, C - Connected, S - Static
             B - BGP Derived, IA - OSPF Inter Area
             E1 - OSPF External Type 1, E2 - OSPF External Type 2
             N1 - OSPF NSSA External Type 1, N2 - OSPF NSSA External Type 2
             L - Leaked Route
C           30.30.30.0/24 [0/1] directly connected,   vlan 30
```

```

S L    10.10.10.0/24 [1/1] directly connected,  vlan 10
S L    11.11.11.0/24 [1/1] directly connected,  vlan 11
S L    20.20.20.0/24 [1/1] directly connected,  vlan 20
S L    22.22.22.0/24 [1/1] directly connected,  vlan 22

```

### 9.17.5 VRF Configuration Example

1. Create virtual router instances. The following commands create and name two instances and enter VRF Configuration mode for each.

In VRF Configuration mode for each VR, a description is added and the maximum number of routes allowed in each virtual instance is configured. On the Red instance, the number of routes above which a warning message is issued is also configured.

The ip routing command enables routing in each VR instance:

```

(Router) #configure
(Router) (Config)#ip vrf Red
(Router) (Config-vrf-Red)#description "finance department"
(Router) (Config-vrf-Red)#maximum routes 2048
(Router) (Config-vrf-Red)#maximum routes warn 80
(Router) (Config-vrf-Red)#ip routing
(Router) (Config-vrf-Red)#exit

(Router) (Config)#ip vrf Blue
(Router) (Config-vrf-Blue)#description "human resources department"
(Router) (Config-vrf-Blue)#maximum routes 4096
(Router) (Config-vrf-Blue)#ip routing
(Router) (Config-vrf-Blue)#exit

```

2. In Interface Config mode, assign interfaces to each virtual router:

```

(Router) (Config)#interface 0/1
(Router) (Interface 1/0/1)#ip vrf forwarding Red
Warning: routing interface moved from Default router instance to "Red" router instance.
(Router) (Interface 1/0/1)#exit

(Router) (Config)#interface 0/2
(Router) (Interface 1/0/2)#ip vrf forwarding Blue
Warning: routing interface moved from Default router instance to "Blue" router instance.
(Router) (Interface 1/0/2)#exit

```

3. Create static leaked routes as needed in the VR instances.

In the following example, subnetwork 9.0.0.0/24 is a connected subnetwork in the global route table and subnet 56.6.6.0/24 is reachable via a gateway 9.0.0.2 in the global route table. Subnet 8.0.0.0/24 is a connected subnetwork in virtual router Red.

The two routes are leaked from the global route table into the Red VR and the connected subnet 8.0.0.0/24 is leaked from the Red VR to the global route table.

The following commands also add a non-leaked static route for the 56.6.6.0/24 subnetwork scoped to the domain of Red VR.

```

(Router) (Config)#ip routing
(Router) (Config)#interface 0/27
(Router) (Interface 0/27)#routing
(Router) (Interface 0/27)#ip vrf forwarding Red
Warning: routing interface moved from Default router instance to "Red" router instance.
(Router) (Interface 0/27)#ip address 8.0.0.1 /24

(Router) (Interface 0/27)#interface 0/26
(Router) (Interface 0/26)#routing
(Router) (Interface 0/26)#ip address 9.0.0.1 /24
(Router) (Interface 0/26)#exit

```

```
(Router) (Config)#ip route 56.6.6.0 255.255.255.0 9.0.0.2
```

To leak routes from the global routing table to the VRF route table, use the following commands:

```
(Router) (Config)#ip route vrf Red 9.0.0.2 255.255.255.255 9.0.0.2 0/26
(Router) (Config)#ip route vrf Red 56.6.6.0 255.255.255.0 9.0.0.2 0/26
```

To leak routes from the VRF route table to the global routing table, use the following command:

```
(Router) (Config)#ip route 8.0.0.2 255.255.255.255 0/27
```

To configure the (non-leaked) internal route to the VRF route table, use the following command:

```
(Router) (Config)#ip route vrf Red 66.6.6.0 255.255.255.0 8.0.0.2
```

## 9.18 Bidirectional Forwarding Detection

---

### **NOTICE**

Bidirectional Forwarding Detection (BFD) configuration can be performed only by using the CLI. Web UI and SNMP configuration options are not supported for the BFD feature.

---

### 9.18.1 Overview

In a network device, BFD is presented as a service to its user applications, providing them with options to create and destroy a session with a peer device and reporting on the session status. On FASTPATH switches, BGP can use BFD for monitoring of their neighbors' availability in the network and for fast detection of connection faults with them.

BFD uses a simple 'hello' mechanism that is similar to the neighbor detection components of some well-known protocols. It establishes an operational session between a pair of network devices to detect a two-way communication path between them and serves information regarding it to the user applications. The pair of devices transmits BFD packets between them periodically, and if one stops receiving peer packets within detection time limit it considers the bidirectional path to have failed. It then notifies the application protocol using its services.

BFD allows each device to estimate how quickly it can send and receive BFD packets to agree with its neighbor upon how fast detection of failure could be done.

BFD can operate between two devices on top of any underlying data protocol (network layer, link layer, tunnels, and so on) as payload of any encapsulating protocol appropriate for the transmission medium. The FASTPATH implementation works with IPv4 and IPv6 networks and supports IPv4 and IPv6 address-based encapsulations.

## 9.18.2 Configuring BFD

The following command sequence enables BFD and configures session parameters:

1. First, globally enable BFD:

```
(Router)#configure
(Router) (Config)# feature bfd
```

2. Configure session settings. These can be configured globally or on a per-interface basis.

```
(Router) (Config)#bfd interval 100 min_rx 200 multiplier 5
(Router) (Config)#bfd slow-timer 1000
```

- The argument `interval` refers to the desired minimum transmit interval, the minimum interval that the user wants to use while transmitting BFD control packets (in ms).
- The argument `min_rx` refers to the required minimum receive interval, the minimum interval at which the system can receive BFD control packets (in ms).
- The argument `multiplier` specifies the number of BFD control packets to be missed in a row to declare a session down.
- The `slow-timer` command sets up the BFD required echo receive interval preference value (in ms). This value determines the interval the asynchronous sessions use for BFD control packets when the echo function is enabled. The `slow-timer` value is used as the new control packet interval, while the echo packets use the configured BFD intervals.

3. Configure BGP to use BFD for fast detection of faults between neighboring devices. A neighboring device IP address

```
(Router) (Config)#router bgp
(Router) (Config-router)# neighbor 172.16.11.6 fall-over bfd
(Router) (Config-router)# exit
```