

LCOS SX 5.20

WEBconfig

02/2026



LANCOM
SYSTEMS

Contents

1 Introduction.....	13
1.1 About This Document.....	13
1.1.1 Document Conventions.....	13
1.2 About Software Modules.....	13
1.3 Configuring BGP Features.....	14
2 Getting Started.....	15
2.1 Connecting the Switch to the Network.....	15
2.2 Booting the Switch / Using the Utility Menu via terminal connection.....	17
2.2.1 Utility Menu Functions.....	17
2.2.2 Select Serial Speed.....	19
2.2.3 Erase All Configuration Files.....	22
2.3 Understanding the User Interfaces.....	22
2.3.1 Using the Web Interface.....	22
2.3.2 Using the Command Line Interface.....	26
3 Configuring and viewing System Information.....	28
3.1 Viewing the Dashboard.....	28
3.2 Viewing ARP Cache.....	29
3.3 Viewing Inventory Information.....	30
3.4 Viewing the System Firmware Status.....	31
3.4.1 Dual Image Status.....	32
3.4.2 Dual Image Configuration and Upgrade.....	32
3.5 Configuring and viewing System Resources.....	33
3.6 Selecting the SDM Template.....	35
3.7 Defining General Device Information.....	37
3.7.1 System Description.....	37
3.7.2 Switch Configuration.....	38
3.7.3 IP Address Conflict Detection.....	39
3.7.4 IPv4 Network Connectivity Configuration.....	40
3.7.5 IPv6 Network Connectivity Configuration.....	41
3.7.6 Network Port IPv6 Neighbors.....	43
3.7.7 Service Port IPv4.....	44
3.7.8 Service Port IPv6.....	45
3.7.9 Service Port IPv6 Neighbors.....	47
3.7.10 DHCP Client Options.....	48
3.7.11 System Connectivity.....	49
3.7.12 Telnet Session.....	50
3.7.13 Outbound Telnet Configuration.....	51
3.7.14 Serial Port.....	52
3.7.15 CLI Banner Configuration.....	53
3.7.16 HTTP Configuration.....	54

3.7.17 HTTPS Configuration.....	55
3.7.18 SSH Configuration.....	57
3.7.19 Management Access Control and Administration List.....	59
3.7.20 User Accounts.....	61
3.7.21 Authentication Server Users.....	64
3.7.22 Logged-in Sessions.....	65
3.7.23 User Domain Name.....	66
3.7.24 Task Group.....	67
3.7.25 User Group.....	70
3.7.26 Accounting List Configuration.....	73
3.7.27 Accounting List Selection.....	75
3.7.28 Authentication List Configuration.....	76
3.7.29 Authentication List Selection.....	79
3.7.30 Authorization List Configuration.....	80
3.7.31 Enable Password.....	83
3.7.32 Password Rules.....	84
3.7.33 Denial of Service.....	85
3.8 Configuring and Searching the Forwarding Database.....	88
3.8.1 Basic Switch Configuration.....	88
3.9 Configuring LMC Configuration.....	89
3.9.1 LMC Status.....	90
3.10 Managing Logs.....	91
3.10.1 Log Configuration.....	91
3.10.2 Buffered Log.....	93
3.10.3 Event Log.....	95
3.10.4 Hosts Log Configuration.....	96
3.10.5 Syslog Source Interface Configuration.....	97
3.10.6 Persistent Log.....	98
3.11 Configuring Email Alerts.....	99
3.11.1 Email Alert Global Configuration.....	100
3.11.2 Email Alerts Server Configuration.....	101
3.11.3 Email Alert Statistics.....	102
3.11.4 Email Alert Subject Configuration.....	103
3.11.5 Email Alerts To Address Configuration.....	104
3.12 Configuring Power over Ethernet.....	105
3.12.1 PoE Summary.....	105
3.12.2 PoE Configuration.....	107
3.12.3 PoE Auto Checking.....	109
3.12.4 PoE Schedule Profile.....	110
3.12.5 PoE Power Information.....	112
3.13 Viewing Device Port Information.....	113
3.13.1 Port Summary.....	113
3.13.2 Port Description.....	116
3.13.3 Port Cable Test.....	118

3.13.4 Mirroring.....	119
3.13.5 Mirroring Summary.....	123
3.13.6 SFP Information.....	124
3.14 Configuring sFlow.....	125
3.14.1 sFlow Agent Summary.....	125
3.14.2 sFlow Receiver Configuration.....	126
3.14.3 sFlow Poller Configuration.....	128
3.14.4 sFlow Sampler Configuration.....	129
3.14.5 sFlow Source Interface Configuration.....	131
3.15 Defining SNMP Parameters.....	132
3.15.1 SNMP v1 and v2.....	133
3.15.2 SNMP v3.....	133
3.15.3 SNMP Community Configuration.....	133
3.15.4 Trap Receiver v1/v2 Configuration.....	136
3.15.5 Trap Receiver v3 Configuration.....	138
3.15.6 SNMP Supported MIBs.....	140
3.15.7 SNMP Access Control Group.....	140
3.15.8 SNMP User Security Model.....	143
3.15.9 SNMP Source Interface Configuration.....	145
3.16 Viewing System Statistics.....	146
3.16.1 Switch Statistics.....	147
3.16.2 Port Summary Statistics.....	148
3.16.3 Port Detailed Statistics.....	149
3.16.4 Port DHCPv6 Client Statistics.....	153
3.16.5 Time Based Group Statistics.....	155
3.16.6 Time Based Flow Statistics.....	157
3.16.7 Time Based Statistics.....	160
3.17 Using System Utilities.....	161
3.17.1 System Reboot.....	161
3.17.2 Ping.....	161
3.17.3 Ping IPv6.....	163
3.17.4 Traceroute.....	164
3.17.5 IP Address Conflict Detection.....	165
3.17.6 File Transfer.....	166
3.17.7 Core Dump.....	171
3.17.8 Core Dump Test.....	173
3.18 Managing SNMP Traps.....	173
3.18.1 System Trap Log.....	173
3.18.2 System Trap Flags.....	174
3.19 Configuring Time Ranges.....	175
3.19.1 Time Range Summary.....	176
3.19.2 Time Range Entry Summary.....	177
3.20 Configuring DNS.....	180
3.20.1 Global Configuration.....	180

3.20.2 DNS IP Mapping Configuration.....	181
3.20.3 DNS Source Interface Configuration.....	182
3.21 Configuring SNTP Settings.....	183
3.21.1 SNTP Global Configuration.....	184
3.21.2 SNTP Global Status.....	186
3.21.3 SNTP Server Configuration.....	187
3.21.4 SNTP Server Status.....	188
3.21.5 SNTP Source Interface Configuration.....	189
3.22 Configuring the Time Zone.....	191
3.22.1 Time Zone Configuration.....	192
3.22.2 Summer Time Configuration.....	193
3.23 Configuring and Viewing ISDP Information.....	195
3.23.1 ISDP Global Configuration.....	195
3.23.2 ISDP Cache Table.....	196
3.23.3 ISDP Interface Configuration.....	197
3.23.4 Statistics.....	198
3.24 Link Dependency.....	199
3.24.1 Link Dependency Group Status.....	199
4 Configuring Switching Information.....	203
4.1 IP Device Tracking.....	203
4.1.1 Device Tracking Global Configuration.....	203
4.1.2 Device Tracking Summary.....	204
4.1.3 Device Tracking Interface Configuration.....	205
4.1.4 Device Tracking Statistics.....	206
4.2 Loop Protection.....	207
4.2.1 Loop Protection Configuration.....	207
4.3 Managing VLANs.....	209
4.3.1 VLAN Overview.....	210
4.3.2 VLAN Port Configuration.....	212
4.3.3 VLAN Port Summary.....	214
4.3.4 VLAN Switchport Summary.....	216
4.3.5 VLAN Internal Usage.....	219
4.3.6 Configure VLAN Statistics.....	219
4.3.7 Reset VLAN Configuration.....	221
4.3.8 RSPAN Configuration.....	221
4.4 Configuring UDLD.....	222
4.4.1 UDLD Interface Configuration.....	222
4.5 MAC-Based VLAN Status.....	224
4.6 DVLAN Tunneling.....	226
4.6.1 DVLAN Configuration.....	226
4.6.2 DVLAN Summary.....	227
4.6.3 DVLAN Interface Summary.....	228
4.7 IP Subnet Based VLAN Status.....	229
4.8 Protocol-Based VLAN Configuration.....	231

4.8.1 Protocol-Based VLAN Overview.....	231
4.8.2 Protocol Based VLAN Group Configuration.....	234
4.9 Private VLAN.....	235
4.9.1 Private VLAN Configuration.....	235
4.9.2 Private VLAN Association.....	237
4.9.3 Private VLAN Interface Association.....	238
4.10 Voice VLAN Configuration.....	241
4.10.1 Voice VLAN Interface Summary.....	242
4.11 Virtual Port Channel Global Configuration.....	244
4.11.1 Virtual Port Channel Interface Configuration.....	248
4.11.2 Virtual Port Channel Statistics.....	251
4.12 Port Auto Recovery.....	253
4.12.1 Port Auto Recovery Configuration.....	253
4.13 Creating MAC Filters.....	256
4.13.1 Static MAC Filter Configuration Summary.....	256
4.14 Configuring Dynamic ARP Inspection.....	258
4.14.1 Global Configuration.....	258
4.14.2 Dynamic ARP Inspection VLAN Configuration.....	259
4.14.3 Interface Configuration.....	261
4.14.4 Dynamic ARP Inspection ACL Summary.....	262
4.14.5 Dynamic ARP Inspection ACL Configuration.....	263
4.14.6 Dynamic ARP Inspection Statistics.....	265
4.15 GARP Configuration.....	266
4.15.1 GARP Switch Configuration.....	267
4.15.2 GARP Port Configuration.....	267
4.16 Configuring DHCP Snooping.....	269
4.16.1 DHCP Snooping Configuration.....	270
4.16.2 DHCP Snooping VLAN Configuration.....	270
4.16.3 DHCP Snooping Interface Configuration.....	271
4.16.4 DHCP Snooping Static Bindings.....	274
4.16.5 DHCP Snooping Dynamic Bindings.....	275
4.16.6 DHCP Snooping Persistent Configuration.....	276
4.16.7 DHCP Snooping Statistics.....	276
4.16.8 DHCP L2 Relay Global Configuration.....	277
4.16.9 DHCP L2 Relay Interface Configuration.....	278
4.16.10 DHCP L2 Relay VLAN Configuration.....	279
4.16.11 DHCP L2 Relay Interface Statistics.....	281
4.16.12 DHCP Snooping IP Source Guard Interface Configuration.....	282
4.16.13 DHCP Snooping IP Source Guard Bindings.....	284
4.17 Configuring IGMP Snooping.....	285
4.17.1 IGMP Snooping Global Configuration and Status.....	286
4.17.2 IGMP Snooping Interface Configuration.....	286
4.17.3 IGMP Snooping Source Specific Multicast.....	288
4.17.4 IGMP Snooping VLAN Status.....	289

4.17.5 IGMP Snooping Multicast Router Configuration.....	291
4.17.6 IGMP Snooping Multicast Router VLAN Status.....	293
4.17.7 IGMP Snooping Multicast Router VLAN Configuration.....	293
4.18 Configuring IGMP Snooping Querier.....	294
4.18.1 IGMP Snooping Querier Configuration.....	295
4.18.2 IGMP Snooping Querier VLAN Configuration.....	295
4.18.3 IGMP Snooping Querier VLAN Overview.....	297
4.19 Configuring MLD Snooping.....	298
4.19.1 MLD Snooping Configuration and Status.....	298
4.19.2 MLD Snooping Interface Configuration.....	299
4.19.3 MLD Snooping Source Specific Multicast.....	301
4.19.4 MLD Snooping VLAN Status.....	301
4.19.5 Multicast Router Configuration.....	303
4.19.6 MLD Snooping Multicast Router VLAN Overview.....	304
4.20 Configuring MLD Snooping Querier.....	305
4.20.1 MLD Snooping Querier Configuration.....	306
4.20.2 MLD Snooping Querier VLAN Configuration.....	306
4.20.3 MLD Snooping Querier VLAN Overview.....	308
4.21 Creating Port Channels.....	309
4.21.1 Port Channel Summary.....	309
4.21.2 Port Channel Statistics.....	313
4.22 Viewing Multicast Forwarding Database Information.....	314
4.22.1 Multicast Forwarding Database Summary.....	314
4.22.2 Multicast Forwarding Database GMRP Table.....	316
4.22.3 Multicast Forwarding Database IGMP Snooping Table.....	316
4.22.4 MFDB Source Specific Multicast.....	317
4.22.5 Multicast Forwarding Database Statistics Source Specific Multicast Status.....	318
4.22.6 Multicast Forwarding Database Statistics.....	319
4.23 Multicast VLAN Registration.....	320
4.23.1 MVR Global Configuration.....	320
4.23.2 MVR Group Status.....	321
4.23.3 MVR Interface Status.....	322
4.23.4 MVR Statistics.....	325
4.24 Configuring Protected Ports.....	325
4.25 Priority Flow Control.....	327
4.25.1 Priority Flow Control Configuration.....	327
4.25.2 Priority Flow Control Statistics.....	329
4.26 Configuring Spanning Tree Protocol.....	330
4.26.1 Spanning Tree Switch Configuration.....	330
4.26.2 Spanning Tree CST Configuration.....	331
4.26.3 Spanning Tree CST Port Summary.....	333
4.26.4 Spanning Tree MST Summary.....	337
4.26.5 Spanning Tree MST Port Summary.....	340
4.26.6 Spanning Tree Statistics.....	343

4.26.7 PVST/RPVST Global Configuration.....	344
4.26.8 PVST/RPVST VLAN Configuration.....	345
4.26.9 PVST/RPVST Interface Configuration.....	348
4.26.10 PVST/RPVST Statistics.....	349
4.27 802.1p Priority Mapping.....	350
4.28 Configuring Port Security.....	352
4.28.1 Port Security Global Administration.....	353
4.28.2 Port Security Interface Status.....	353
4.28.3 VLAN MAC Locking Status.....	355
4.28.4 Port Security Static MAC Addresses.....	357
4.28.5 Port Security Dynamic MAC Addresses.....	359
4.29 Managing LLDP.....	359
4.29.1 LLDP Global Configuration.....	360
4.29.2 LLDP Interface Summary.....	361
4.29.3 LLDP Local Device Summary.....	363
4.29.4 LLDP Remote Device Summary.....	364
4.29.5 LLDP Statistics.....	365
4.29.6 LLDP-MED.....	367
5 Configuring Routing.....	374
5.1 Configuring ARP.....	374
5.1.1 ARP Create.....	375
5.1.2 ARP Table Configuration.....	376
5.2 Configuring Global IP Settings.....	377
5.2.1 Routing IP Configuration.....	377
5.2.2 Routing IP Interface Summary.....	378
5.2.3 Routing IP Interface Configuration.....	382
5.2.4 Routing IP Loopback Configuration.....	384
5.2.5 Routing IP Statistics.....	385
5.3 Router.....	387
5.3.1 Route Table Summary.....	388
5.3.2 Configured Route Summary.....	389
5.3.3 IP Route Summary.....	391
5.3.4 ECMP Groups Summary.....	393
5.4 Configuring IPv6 Settings.....	393
5.4.1 IPv6 Global Configuration.....	393
5.4.2 IPv6 Interface Summary.....	395
5.4.3 IPv6 Interface Configuration.....	398
5.4.4 IPv6 Loopback Configuration.....	401
5.4.5 IPv6 Global Address Table.....	402
5.4.6 IPv6 Global Address Configuration.....	403
5.4.7 IPv6 Statistics.....	404
5.4.8 IPv6 Detailed Statistics.....	405
5.4.9 IPv6 Neighbor Table.....	409
5.5 Configuring IPv6 Routes.....	411

5.5.1 IPv6 Route Table.....	411
5.5.2 IPv6 Configured Routes.....	412
5.5.3 IPv6 ECMP Groups Summary.....	414
5.5.4 IPv6 Route Summary.....	415
5.6 Configuring DHCPv4.....	417
5.6.1 DHCP Server Global Configuration.....	417
5.6.2 DHCP Server Excluded Addresses.....	418
5.6.3 DHCP Server Pool Summary.....	419
5.6.4 DHCP Server Pool Configuration.....	421
5.6.5 DHCP Server Pool Options.....	424
5.6.6 DHCP Server Bindings.....	426
5.6.7 DHCP Server Statistics.....	427
5.6.8 DHCP Server Conflicts Information.....	429
5.7 Configuring DHCPv6.....	429
5.7.1 DHCPv6 Global Configuration.....	430
5.7.2 DHCPv6 Pool Summary.....	430
5.7.3 DHCPv6 Pool Configuration.....	431
5.7.4 DHCPv6 Interface Summary.....	433
5.7.5 DHCPv6 Interface Configuration.....	434
5.7.6 DHCPv6 Binding Summary.....	435
5.7.7 DHCPv6 Statistics.....	436
5.7.8 DHCPv6 Server Conflicts Information.....	439
6 Managing Device Security.....	440
6.1 Port Access Control.....	440
6.1.1 Port Access Control Configuration.....	440
6.1.2 Port Access Control Port Summary.....	441
6.1.3 Port Access Control Port Configuration.....	442
6.1.4 Port Access Control Port Details.....	444
6.1.5 Port Access Control Statistics.....	446
6.1.6 Port Access Control Privileges Summary.....	448
6.2 RADIUS Settings.....	450
6.2.1 RADIUS Configuration.....	450
6.2.2 RADIUS Named Server Status.....	453
6.2.3 RADIUS Server Statistics.....	455
6.2.4 RADIUS Accounting Server Status.....	457
6.2.5 RADIUS Accounting Server Statistics.....	458
6.2.6 RADIUS Clear Statistics.....	460
6.2.7 RADIUS Source Interface Configuration.....	461
6.3 TACACS+ Configuration.....	462
6.3.1 TACACS+ Server Summary.....	462
6.3.2 TACACS+ Server Configuration.....	463
6.3.3 TACACS+ Source Interface Configuration.....	464
6.4 Authentication Manager.....	465
6.4.1 Authentication Manager Configuration.....	466

6.4.2 Authentication Manager Interface Configuration.....	467
6.4.3 Authentication Tiering.....	471
6.4.4 Authenticated Clients.....	472
6.4.5 Authentication Statistics.....	475
6.4.6 Authentication History.....	475
7 Configuring Quality of Service.....	478
7.1 Configuring Access Control Lists.....	478
7.1.1 IP Access Control Lists.....	478
7.1.2 Configuring Auto VoIP.....	494
7.1.3 Configuring Class of Service.....	499
7.1.4 Configuring DiffServ.....	506
8 Getting Started with Stacking.....	519
8.1 Understanding Switch Stacks.....	519
8.1.1 Switch Stack Membership.....	520
8.1.2 Stack Manager Election and Re-Election.....	520
8.1.3 Stack Member Numbers.....	520
8.1.4 Stack Member Priority Values.....	520
8.2 Switch Stack Software Compatibility Recommendations.....	521
8.3 Incompatible Software and Stack Member Image Upgrades.....	521
8.4 Switch Stack Configuration Files.....	521
8.5 Switch Stack Management Connectivity.....	522
8.5.1 Connectivity to the Switch Stack through Console Ports.....	522
8.5.2 Connectivity to the Switch Stack through Telnet.....	522
8.6 General Practices.....	522
8.7 Initial Installation and Power-up of a Stack.....	522
8.8 Removing a Unit from the Stack.....	523
8.9 Adding a Unit to an Operating Stack.....	523
8.10 Replacing the Stack Member with a New Unit.....	524
8.11 Renumbering Stack Members.....	524
8.12 Moving a Manager to a Different Unit in the Stack.....	525
8.13 Removing a Manager Unit from an Operating Stack.....	525
8.14 Initiating a Warm Failover of the Manager Unit.....	525
8.15 Merging Two Operational Stacks.....	526
8.16 Preconfiguration.....	526
8.17 Stack Links.....	527
8.18 Dynamic Load Balancing.....	527
9 Configuration Examples.....	529
9.1 VLAN Configuration Examples.....	529
9.1.1 Using the Web Interface to Configure VLANs.....	529
9.1.2 Using the CLI to Configure VLANs.....	537
9.1.3 Using SNMP to Configure VLANs.....	538
9.2 Configuring Multiple Spanning Tree Protocol.....	539
9.2.1 Using the Web UI to Configure MSTP.....	540
9.2.2 Using the CLI to Configure MSTP.....	541

9.2.3 Using SNMP to Configure MSTP.....	542
9.3 Configuring VLAN Routing.....	543
9.3.1 Using the CLI to Configure VLAN Routing.....	543
9.3.2 Using SNMP to Configure VLAN Routing.....	544
9.4 Configuring Policy-Based Routing.....	545
9.4.1 Configuring Policy-Based Routing Using the CLI.....	546
9.5 Configuring OSPF.....	548
9.5.1 Using the CLI to Configure OSPF.....	548
9.5.2 Configuring Stub and NSSA Areas.....	549
9.5.3 Using the CLI to Configure OSPF Areas.....	550
9.5.4 Using the CLI to Configure OSPFv3 Enhancements.....	551
9.6 Configuring 802.1X Network Access Control.....	552
9.6.1 Using the CLI to Configure 802.1X Port-Based Access Control.....	553
9.6.2 Using SNMP to Configure 802.1X Port-Based Access Control.....	553
9.7 Configuring Authentication Tiering.....	554
9.7.1 Configuring Authentication Tiering Using the Web Interface.....	554
9.7.2 Configuring Authentication Tiering Using the CLI.....	555
9.8 Configuring Differentiated Services for VoIP.....	556
9.8.1 Using the CLI to Configure DiffServ VoIP Support.....	556
9.8.2 Using SNMP to Configure DiffServ VoIP Support.....	557
9.9 Configuring PIM.....	558
9.9.1 Using the CLI to Configure PIM-SMv4.....	558
9.9.2 Using SNMP to Configure PIM-SMv4.....	559
9.9.3 Configuring IP Multicast Routing with PIM Sparse Mode.....	559
9.10 IGMP and MLD Snooping Switches.....	563
9.10.1 Snooping Functionality on a LCOS SX Switch.....	563
9.10.2 Snooping Switch Restrictions.....	565
9.10.3 Configuring IGMP and MLD Snooping.....	566
9.11 Multicast Snooping Example (with IP Multicast Routing).....	568
9.11.1 Snooping within a Subnet.....	568
9.11.2 Snooping on a Multicast Router.....	569
9.12 Configuring Port Mirroring.....	571
9.13 Configuring RSPAN.....	572
9.13.1 Configuring RSPAN Using the Web Interface.....	572
9.13.2 Configuring RSPAN Using the CLI.....	577
9.14 Configuring Virtual Port Channel.....	579
9.14.1 Configuring Virtual Port Channel Using the Web Interface.....	579
9.14.2 Configuring Virtual Port Channel Using the CLI.....	582
9.15 Bidirectional Forwarding Detection.....	583
9.15.1 Overview.....	584
9.15.2 Configuring BFD.....	584
9.16 Interactive SSH.....	584

Copyright

© 2026 LANCOM Systems GmbH, Würselen (Germany). All rights reserved.

While the information in this manual has been compiled with great care, it may not be deemed an assurance of product characteristics. LANCOM Systems shall be liable only to the degree specified in the terms of sale and delivery.

The reproduction and distribution of the documentation and software supplied with this product and the use of its contents is subject to written authorization from LANCOM Systems. We reserve the right to make any alterations that arise as the result of technical development.

Windows® and Microsoft® are registered trademarks of Microsoft, Corp.

LANCOM, LANCOM Systems, LCOS, LANcommunity, LANCOM Service LANcare, LANCOM Active Radio Control, and AirLancer are registered trademarks. All other names or descriptions used may be trademarks or registered trademarks of their owners. This document contains statements relating to future products and their attributes. LANCOM Systems reserves the right to change these without notice. No liability for technical errors and/or omissions.

This product contains separate open-source software components which are subject to their own licenses, in particular the General Public License (GPL). If the respective license demands, the source files for the corresponding software components will be made available on a download server upon request.

LANCOM Systems GmbH

A Rohde & Schwarz Company

Adenauerstr. 20/B2

52146 Wüerselen

Germany

www.lancom-systems.com

1 Introduction

1.1 About This Document

This guide describes how to configure the LCOS SX software features by using the Web-based GUI. The LCOS SX architecture accommodates a variety of software modules so that a platform running LCOS SX 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:

- a System administrators who are responsible for configuring and operating a network using LCOS SX software
- a Software engineers who are integrating LCOS SX software into a router or switch product
- a 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
Bold	User input and actions: for example, type exit , click OK , press Alt+C
<i>Blue Text</i>	Hyperlinked text.
Monospace	Code: <code>#include <iostream></code> Command line commands and parameters: <code>show network</code>
<i>Monospace italic</i>	Placeholders for <i>required</i> elements: <code>username name</code>
[]	Indicates <i>optional</i> command-line parameters: <code>show port [all]</code>
{ }	Indicates a choice in command-line parameters: <code>network protocol {dhcp bootp none}</code>

1.2 About Software Modules

The LCOS SX software suite includes the following modules:

- a Switching (Layer 2)
- a IPv4-IPv6 routing (Layer 3 routing)
- a Multicast
- a Quality of Service

- a Management (CLI, Web UI, SNMP, Open Ethernet Networking (OpEN) API, RESTful API, RESTCONF, NETCONF)
- a Metro
- a Stacking
- a BGP
- a IPv6 Management
- a QoS
- a Secure Management
- a Service Provider

4 Not all modules are available for all platforms or software releases. The LCOS SX 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 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. Additionally, the Border Gateway Control (BGP) features can be configured only by using the CLI. There are no web pages within LCOS SX for configuring BGP.

For information about the CLI commands you use to configure BGP, refer to the LCOS SX 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 via the CLI (SSH or Telnet), a Web browser, or SNMP, you must connect the switch to the network. We recommend connecting to the switch via LAN.

Connection via Local Area Network:

If a switch, which is in the factory state, is operated in a network without a DHCP server, the device has the IP address 172.23.56.250.

To be able to access the device via TCP/IP using a web browser, the configuration PC must have a fixed IP address from the address range 172.23.56.0/24 in its network settings (e.g. 172.23.56.249).

If a switch, which is in factory state, is operated in a network with DHCP server, the device automatically receives a local IP address from the DHCP server.

4 If you want to use static IP parameters on the switch, follow [these instructions from the LANCOM Support Knowledge Base](#).

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.

Connect via Terminal interface (EIA-232 port) to set up IP parameters:

4 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.

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:

- a Baud rate: 115200 bps
- a Data bits: 8
- a Parity: none
- a Stop bit: 1
- a Flow control: none

3. Power on the switch.

For information about the boot process, including how to access the utility menu, see [Booting the Switch / Using the Utility Menu via terminal connection](#) on page 17.

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:

- a 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.

- a To use BootP, change the protocol by entering:

```
serviceport protocol bootp
```

- a 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
```

- a 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>
```

- a To view the assigned or configured network address, enter:

```
show serviceport
```

If the unit does not have a service port:

- a To use a DHCP server to obtain the IP address, subnet mask, and default gateway information, enter:

```
network protocol dhcp.
```

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

```
network protocol bootp.
```

- a 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
```

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

```
network ipv6 address <address>/<prefix-length> [eui64]
network ipv6 gateway <gateway>
```

- a To view the network information, enter

```
show network.
```

- a 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 Booting the Switch / Using the Utility Menu via terminal connection

When the power is turned on with the local terminal already connected, the switch goes through Power-On Self-Test (POST). POST runs every time the switch is initialized and checks hardware components to determine if the switch is fully operational before completely booting.

If a critical problem is detected, the program flow stops. If POST passes successfully, a valid executable image is loaded into RAM.

POST messages are displayed on the terminal and indicate test success or failure. To boot the switch, perform the following steps:

1. Make sure that the serial cable is connected to the terminal.
2. Connect the power supply to the switch.
3. Power on the switch.

As the switch boots, the bootup test first counts the switch memory availability and then continues to boot.

4. During boot, you can use the Utility menu, if necessary, to run special procedures. To enter the Utility menu, press **2** within the first five seconds after the following message appears.

```
Select startup mode. If no selection is made within 5 seconds,
the LCOS SX Application will start automatically...

LCOS SX Startup -- Main Menu

1 - Start LCOS SX Application
2 - Display Utility Menu
Select (1, 2): 2
```

For information about the Utility menu, see [Utility Menu Functions](#) on page 17.

5. If you do not start the Utility menu, the operational code continues to load.

After the switch boots successfully, the User login prompt appears and you can use the local terminal to begin configuring the switch. However, before configuring the switch, make sure that the software version installed on the switch is the latest version. If it is not the latest version, download and install the latest version.

2.2.1 Utility Menu Functions

- 4** Utility menu functions vary on different LCOS SX products, operating systems, and switch platforms. The following example might not represent the options available on your platform.

You can perform many configuration tasks through the Utility menu, which can be invoked after the first part of the POST is completed.

Use the following procedures to display the Utility menu:

1. During the boot process, press **2** within three seconds after the following message is displayed:

```
LCOS SX Startup
Select startup mode. If no selection is made within 3seconds,
the LCOS SX Application will start automatically...

LCOS SX Startup -- Main Menu
1 - Start LCOS SX Application
2 - Display Utility Menu Select (1, 2): 2

LCOS SX Startup Options available
1 - Start LCOS SX Application
2 - Load Code Update Package
3 - Load Configuration
4 - Select Serial Speed
5 - Retrieve Error Log
6 - Erase Current Configuration
7 - Erase Permanent Storage
8 - Select Boot Method
9 - Activate Backup Image
10 - Start Diagnostic Application
11- Reboot
12- Erase All Configuration Files
Q - Quit from LCOS SX Startup
```

The following sections describe the Utility menu options. If no selection is made within 3 seconds (default), the operational code starts.

2.2.1.1 Start LCOS SX Application

Use option 1 to resume loading the LCOS SX Application code. To relaunch the boot process from the Utility menu:

1. On the **Utility menu**, select 1 and press **Enter**.

The following prompt is displayed:

```
Extracting LCOS SX from image2.....done
Loading LCOS SX ...../mnt/application
done
Linking liblua.so to /lib/liblua.so
Linking libluaconn.so to /lib/libluaconn.so
Linking libproc_libs.so to /lib/libproc_libs.so
Linking librpcclt.so to /lib/librpcclt.so
Changing lighttpd file ownership to lighttpd:lighttpd
Expanding websrc.tar.gz into /mnt/www...done
PCI unit 0: Dev 0xb624, Rev 0x11, Chip BCM56624_B0, Driver BCM56624_B0
SOC unit 0 attached to PCI device BCM56624_B0
Adding BCM transport pointers
Configuring CPUTRANS TX
Configuring CPUTRANS RX
```

2.2.1.2 Load Code Update Package

Use option 2 when a new software version must be downloaded to replace corrupted files, update, or upgrade the system software.

To download software from the Utility menu:

1. On the **Utility menu**, select 2 and press **Enter**.

The following prompt is displayed:

```
Select Mode of Transfer (Press T/X/Y/Z for TFTP/XMODEM/YMODEM/ZMODEM)
```

2. Select the transfer mode (Press T/X/Y/Z for TFTP/XMODEM/YMODEM/ZMODEM).
3. When using HyperTerminal, click **Transfer** on the **HyperTerminal** menu bar.
4. From the **Transfer** menu, click **Send File**. The **Send File** window is displayed.
5. Enter the file path for the file to be downloaded.

6. Make sure the protocol is defined per the transfer option selected in Step 2 (XMODEM/YMODEM/ZMODEM).
7. Click **Send**.

The software is downloaded. Software downloading takes several minutes. The terminal emulation application, such as HyperTerminal, may display the loading process progress.

After software downloads, the switch reboots automatically.

2.2.1.3 Load Configuration

Use option 3 when a new configuration file must be downloaded to replace the saved system configuration file. To download software from the Utility menu:

1. On the **Utility menu**, select 3 and press **Enter**.

The following prompt is displayed:

```
[Utility menu] 4
Ready to receive the file with XMODEM/CRC....
Ready to RECEIVE File tempcfg.bin in binary mode
Send several Control-X characters to cancel before transfer starts.
```

Select the transfer mode (Press T/X/Y/Z for TFTP/XMODEM/YMODEM/ZMODEM).

2. When using HyperTerminal, click **Transfer** on the **HyperTerminal** menu bar.
3. From the **Transfer** menu, click **Send File**. The **Send File** window is displayed.
4. Enter the file path for the file to be downloaded.
5. Make sure the protocol is defined per the transfer option selected in Step 2 (XMODEM/YMODEM/ZMODEM).
6. Click **Send**.

The configuration file is downloaded. The terminal emulation application, such as HyperTerminal, may display the loading process progress.

2.2.2 Select Serial Speed

Use option 4 to change the baud rate of the serial interface.

To change the baud rate from the Utility menu:

1. On the **Utility menu**, select 4 and press **Enter**.

The following prompt is displayed:

```
Select option 1-12 or Q):
1 - 2400
2 - 4800
3 - 9600
4 - 19200
5 - 38400
6 - 57600
7 - 115200
8 - Exit without change
```

4 The selected baud rate takes effect immediately.

2. The bootup process resumes.

2.2.2.1 Retrieve Error Log

Use option 5 to retrieve the event log and download it to your ASCII terminal.

To retrieve the event log from the Utility menu:

1. On the **Utility menu**, select 5 and press **Enter**.

The following prompt is displayed:

```
Select Mode of Transfer (Press T/X/Y/Z for TFTP/XMODEM/YMODEM/ZMODEM)
```

2. Select the transfer mode (Press T/X/Y/Z for TFTP/XMODEM/YMODEM/ZMODEM).

The following prompt is displayed:

```
File asciiilog.bin Ready to SEND in binary mode
Estimated File Size 169K, 1345 Sectors, 172032 Bytes
Estimated transmission time 3 minutes 20 seconds
Send several Control-X characters to cancel before transfer starts.
```

3. The bootup process resumes.

2.2.2.2 Erase Current Configuration

Use option 6 to load using the system default configuration and to boot without using the current startup configuration.

Selecting 6 from the Utility menu restores system defaults. Boot Sequence can then be started by selecting 1 from the Utility menu.

To download software from the Utility menu:

1. On the **Utility menu**, select 6 and press **Enter**.

The following prompt is displayed:

```
Are you SURE you want to delete the configuration? (y/n):y
```

2. The bootup process resumes.

2.2.2.3 Erase Permanent Storage

Use option 7 to delete the active image from the flash memory. User action is confirmed with a Y/N question before executing the command.

To delete the backup image from the Utility menu:

1. On the **Utility menu**, select 7 and press **Enter**.

The following prompt is displayed:

```
Are you SURE you want to delete operational code: image2? (y/n):y
Operational code deleted..
[Utility menu]
```

2. The bootup process resumes.

2.2.2.4 Select Boot Method

Use option 8 to select the method used to boot the system (FLASH, Network, or Serial boot). The default selection is FLASH. To select the boot method from the Utility menu:

1. On the **Utility menu**, select 8 and press **Enter**.

The following prompt is displayed:

```
Current boot method: FLASH
1 - Flash Boot
2 - Network Boot
3 - Serial Boot
4 - Exit without change
Select option (1-4):
```

2. The bootup process resumes.

2.2.2.5 Activate Backup Image

Use option 9 to activate the backup image. The active image becomes the backup when this option is selected.

To activate the backup image:

1. From the **Utility menu**, select 9 and press **Enter**.

The following message is displayed:

```
Backup image - image2 activated.
```

2. The bootup process resumes.

2.2.2.6 Start Diagnostic Application

Use option 10 to run flash diagnostics. User action is confirmed with a Y/N question before executing the command.

To perform a complete test of the flash memory from the Utility menu:

1. On the **Utility menu**, select 10 and press **Enter**.

The following prompt is displayed:

```
Do you wish to run flash diagnostics? (Boot code region will not be tested.) (y/n):y
Input number of diagnostic iterations -> 1
Testing 2 x 28F128J3 base: 0xfe000000
Iterations remaining = 1

Erasing sector 0
Verify sector 0 erased
Writing sector 0
Erasing sector 1
Verify sector 1 erased
Writing sector 1
Erasing sector 2
Verify sector 2 erased
Writing sector 2
Erasing sector 3
Verify sector 3 erased
Writing sector 3
Erasing sector 4
Verify sector 4 erased
Writing sector 4
Erasing sector 5
Verify sector 5 erased
Writing sector 5
Erasing sector 6
```

4

This process runs until all sectors have been erased, verified erased, and written.

```
Flash Diagnostics passed
[Utility menu]
```

2. The bootup process resumes.

2.2.2.7 Reboot

Use option 11 to reboot the system.

To reboot the system:

1. From the **Utility menu**, select 11 and press **Enter**.
2. The bootup process resumes.

2.2.3 Erase All Configuration Files

Use option 12 to load using the system default configuration and to boot without using the current startup configuration.

Selecting 12 from the Utility menu restores system defaults. Boot Sequence can then be started by selecting 1 from the Utility menu.

To download software from the Utility menu:

1. On the **Utility menu**, select 12 and press **Enter**.

The following prompt is displayed:

```
Are you SURE you want to delete the configuration? (y/n):y
```

2. The bootup process resumes.

2.3 Understanding the User Interfaces

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

- a Web user interface
- a Command line interface (CLI)
- a Simple Network Management Protocol (SNMP)

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

This guide describes how to use the Web-based interface to manage and monitor the system. However, in Chapter 9, CLI commands are also used to configure the examples.

For information about how to manage and monitor the system by using the CLI, refer to the *LCOS SX CLI Command Reference* and the *LCOS SX Configuration Guide*.

4 The Web configuration and monitoring pages and the CLI commands available for each platform depend on the LCOS SX software version.

2.3.1 Using the Web Interface

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

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

Use the following procedures to log in to the Web Interface shown in [Figure 1: Login Screen](#) on page 23:

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

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




Figure 2: Default Password Change Page

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: Login Screen](#) on page 23) 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 Dashboard with various system and device informations is displayed.

Figure 3: Web Interface Layout on page 24 shows the layout of the LCOS SX software Web interface. Each Web page contains three main areas: device view, the navigation menu, and the configuration status and options.




Figure 3: Web Interface Layout

2.3.1.1 Device View

The Device View is an interactive illustration that displays the ports on the switch. This illustration appears at the top of each page to provide an alternate way to navigate to port related configuration and monitoring options. The illustration also provides information about device ports, current configuration and status, table information, and feature components.

The port coloring indicates if a port is currently active. Green indicates that the port is enabled, red indicates that an error has occurred on the port, and blue indicates that the link is disabled.



Figure 4: Device View

Click the port you want to view or configure to see a menu that displays detailed port related statistics. Click the menu option to access the page that contains the configuration or monitoring options.

2.3.1.2 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 menu item, 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 25 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 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 is displayed in the main frame.

2.3.1.3 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 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: Common Command Buttons](#) on page 25 shows the command buttons that are used throughout the pages in the Web interface.

Table 1: Common Command Buttons

Button	Function
Submit	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.

Button	Function
	To save the configuration to non-volatile memory (as a start configuration), navigate to SystemConfiguration Storage and click Save .
Refresh	Refreshes the page with the most current information.
Remove	Removes the selected entry from the running configuration.
Clear	Removes all entries from a table or resets statistical counters to the default value.
Edit	Changes an existing entry.
Remove	Deletes the selected entries.
Clear Counter	Clear all the statistics counters, resetting all switch summary and detailed statistics to default values.
Log Out	Ends the session.

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

2.3.1.4 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.

2.3.1.5 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 is displayed if you click Help. [Figure 6: Help Icon](#) on page 26 shows the **Help** icon.



Figure 6: Help Icon

[Figure 3: Web Interface Layout](#) on page 24 shows the location of the Help link on the Web interface.

2.3.1.6 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):

```
\      <
/      >
*      |
?
```

2.3.2 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 SSH or Telnet.

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 *LCOS SX CLI Command Reference*.

The *LCOS SX 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:

- a The command keywords and the required and optional parameters.
- a The command mode you must be in to access the command.
- a 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.

2.3.2.1 Using SNMP

For LCOS SX, you can configure SNMP groups and users that can manage traps that the SNMP agent generates.

LCOS SX 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 that is displayed 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 *LCOS SX CLI Command Reference*.

To configure an SNMPv3 profile by using the Web interface, refer to the instructions in the [specific article in the LANCOM Support Knowledge Base](#).

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

3 Configuring and viewing System Information

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

3.1 Viewing the Dashboard

After a successful login, the Dashboard page is displayed. This page provides a brief overview of the system. To navigate to the Dashboard, click **System > Summary > Dashboard** in the navigation menu. As an alternative it is also possible to click on the device name in the upper left corner.




Figure 7: System Dashboard

Table 2: Dashboard Fields

Field	Description
System Information	
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.
System Up Time	The time in days, hours, minutes and seconds since the system was last reset.
Device Information	
Machine Type	The device hardware type or product family.

Field	Description
Machine Model	The model identifier, which is usually related to the Machine Type.
Serial Number	The unique device serial 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.
System Resource Usage	
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 Space Utilization	
Total Disk Space (Kbytes)	The total available disk space.
Free Disk Space (Kbytes)	The currently free disk space.
Used Disk Space (Kbytes)	The currently used disk space.
Disk Usage	The percentage of total available disk space that is currently in use.
Additional Fields	
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.

3.2 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.

3 Configuring and viewing System Information

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



Figure 8: ARP Cache

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 <i>Management</i> in this field.


Use the buttons to perform the following tasks:

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

3.3 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.



System Inventory Information	
Management Unit Number	1
System Description	LANCOM GS-4554XP System - 36x 1G + 12x 2.5G with 4x 10G SFP+ + 2x 40G QSFP+ PoE Switch, 5.20.0065, Linux 4.14.174
Machine Type	LANCOM GS-4554XP System - 36x 1G + 12x 2.5G with 4x 10G SFP+ + 2x 40G QSFP+ PoE Switch
Machine Model	GS-4554XP
Serial Number	4006169020000018
Part Number	BCM56170
Maintenance Level	A
Manufacturer	LANCOM
Burned In MAC Address	00:40:C7:1D:18:C4
Software Version	5.20.0065
Operating System	Linux 4.14.174
Network Processing Device	BCM56170_B0

Refresh

Figure 9: Inventory Information

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.
Part Number	The manufacturing part number.
Maintenance Level	The identification of the hardware change level.
Manufacturer	The name of 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.

Click **Refresh** to update the information on the screen with the most current data.

3.4 Viewing the System Firmware Status

The pages in the Firmware menu allow you to monitor and view the system firmware status.

3.4.1 Dual Image Status

The Dual Image feature allows the switch to have two LCOS SX 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.

Status

Configuration and Upgrade

Autoinstall

Dual Image Status

?

Unit	Active	Backup	Current Active	Next Active
1	5.20.0105	5.10.0033RU2	5.20.0105	5.20.0105

Image Description

Active	
Backup	

Refresh

Figure 10: Dual Image Status

Table 5: Dual Image Status Fields

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

Click **Refresh** to display the latest information from the router.

For information about how to update or change the system images, see [Dual Image Configuration and Upgrade](#).

3.4.2 Dual Image Configuration and Upgrade

Use the Dual Image Configuration and Upgrade feature to transfer a new firmware 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 11: 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 firmware image to activate, upgrade, delete, or describe.
Active	The active firmware version.
Backup	The backup firmware version. Use the icons to the right of the field to perform the following tasks: - a 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. - a To delete the backup image from permanent storage, click the – (minus) icon. You must confirm the action before the image is deleted.
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 firmware.
Backup Description	Use this field to specify a description to associate with the image that is currently the backup firmware.

Use the buttons to perform the following tasks:

- a If you make any changes to the page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.5 Configuring and viewing System Resources

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

- a Free memory
- a Allocated memory
- a CPU utilization by task

3 Configuring and viewing System Information

a Total CPU utilization at the following intervals:

- a Five seconds
- a One minute
- a Five minutes

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




Figure 12: System Resource Status

Table 7: 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 CPU Utilization for individual tasks in terms of percentage. The total CPU Utilization is represented by the task Total. The CPU Utilization is shown in the following intervals: - a 5 seconds - a 60 seconds - a 300 seconds

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

System Resource Configuration		
Rising Threshold (%)	0	(0 to 100, 0 = Default, 0 = Disable)
Rising Threshold Interval (Seconds)	0	(0 to 86400, 0 = Default, 0 = Disable) - Multiple of 5
Falling Threshold (%)	0	(0 to 100, 0 = Default, 0 = Disable)
Falling Threshold Interval (Seconds)	0	(0 to 86400, 0 = Default, 0 = Disable) - Multiple of 5
Free Memory Threshold (Kbytes)	0	(0 to 2008708, 0 = Default, 0 = Disable)

Submit Refresh Cancel

Figure 13: System Resources Configuration

Table 8: 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. When the CPU utilization is increasing, an SNMP trap is generated when it reaches or exceeds this level.
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. When the CPU utilization is decreasing, an SNMP trap is generated when it reaches or falls below this level.
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. If enabled, an SNMP trap is generated when the amount of free memory in the system falls below this value.

Use the buttons to perform the following tasks:

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

3.6 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.

3 Configuring and viewing System Information

In other words, SDM templates enable you to reallocate system resources to support a different mix of features based on your network requirements.

- 4 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.

SDM Template	ARP Entries	IPv4 Unicast Routes	IPv6 NDP Entries	IPv6 Unicast Routes	ECMP Next Hops	IPv4 Multicast Routes	IPv6 Multicast Routes
Dual IPv4 and IPv6	6144	10922	2560	4096	4	1024	512
IPv4-routing Default	6144	16382	0	0	4	2048	0
IPv4 Data Center	6144	10922	0	0	16	2048	0
Data Center Plus - IPv4	6144	16382	0	0	16	2048	0
Data Center - IPv4 and IPv6	6144	10922	2560	4096	16	1024	512

Figure 14: SDM Template Preference

Table 9: 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: - a Dual IPv4 and IPv6 — Both IPv4 and IPv6 are supported (default setting). - a IPv4-routing Default — Only IPv4 is supported. 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. - a IPv4 Data Center — Only IPv4 is supported. This template sets the IPv4 unicast Routes to those used in builds with IPv6. The IPv4 Data Center template increases the maximum number of ECMP next hops from 4 to 16. - a Data Center Plus - IPv4 — Only IPv4 is supported. This template maximizes the number of IPv4 unicast routes and increases the maximum number of ECMP next hops from 4 to 16. - a Data Center - IPv4 and IPv6 — Both IPv4 and IPv6 are supported. This template increases the maximum number of ECMP next hops from 4 to 16.
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.

Field	Description
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.

3.7 Defining General Device Information

The **Configuration** submenu in the **System** menu allows you to configure device parameters.

3.7.1 System Description

Use this page to configure and view general device information.

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

Field	Description
System Description	LANCOM GS-4554XP System - 36x 1G + 12x 2.5G with 4x 10G SFP+ + 2x 40G QSFP+ PoE Switch, 5.20.0065, Linux 4.14.174
System Name	BCM956170R (0 to 255 alphanumeric characters)
System Location	(0 to 255 alphanumeric characters)
System Contact	(0 to 255 alphanumeric characters)
IP Address	172.23.56.250
Service Port IP Address	0.0.0.0
System Up Time	5 days, 23 hours, 45 mins, 14 secs
Current SNTP Synchronized Time	Not Synchronized

Submit Refresh Cancel

Figure 15: System Description

Table 10: 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 alpha-numeric 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 IPv4 Network Connectivity Configuration on page 40.
System Up Time	Displays the number of days, hours, and minutes since the last system restart.

Field	Description
Current SNTP Synchronized Time	Displays currently synchronized SNTP time in UTC. If no SNTP server has been configured and the time is not synchronized, this field displays <code>Not Synchronized</code> . To specify an SNTP server, see Configuring SNTP Settings on page 183.

Use the buttons to perform the following tasks:

- a If you make any changes to the page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.1.1 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.

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

3.7.2 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 16: Switch 802.3x Flow Control

Table 11: 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: a Disabled – The switch does not send PAUSE frames if the port buffers become full.

Field	Description
	- a Symmetric – The switch can send as well as honor the PAUSE frames. The switch generates PAUSE frames towards the peer device in response to congestion at ingress and is also capable of throttling the transmit rate in response to PAUSE frames received from a peer device. - a Asymmetric – The switch can honor PAUSE frames it receives, but it will not generate PAUSE frames towards the peer device in response to congestion at ingress. The switch can throttle the transmit rate in response to the pause frames received from the peer.
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.

Use the buttons to perform the following tasks:

- a If you make any changes to the page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.3 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.

Status	
IP Address Conflict Currently Exists	False

History	
Last Conflicting IP Address	
Last Conflicting MAC Address	
Time Since Conflict Detected	

Refresh Run Detection Clear History

Figure 17: IP Address Conflict Detection

Table 12: 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. - a False – No conflict detected (the subsequent fields on this page display as N/A). - a True – Conflict was detected (the subsequent fields on this page show the relevant information).
Last Conflicting IP Address	The IP address of the interface that was last found to be conflicting. This field is displayed only if a conflict has been detected since the switch was last reset.

3 Configuring and viewing System Information

Field	Description
Last Conflicting MAC Address	The MAC address of the remote host associated with the IP address that was last found to be conflicting.
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).

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To run the tool and check for possible address conflicts, click **Run Detection**.
- a Click **Clear History** to reset the IP address conflict detection status information that was last seen by the device.

3.7.4 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 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 18: Network Connectivity Configuration for IPv4

Table 13: Network Connectivity Configuration for IPv4 Fields

Field	Description
Network Configuration Protocol	Specify what the switch should do following power-up. The factory default is DHCP. The options are as follows: - a None – Do not send any requests following power-up. - a Bootp – Transmit a Bootp request. - a DHCP – Transmit a DHCP request. Click the Refresh DHCP Lease button in the field Network Configuration Protocol to force the interface to release the current DHCP-assigned information and submit a request for new information.
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 when DHCP is enabled on the port on which the Client Identifier option is selected. Please note, that IP connectivity won't be available for a short time as the switch

Field	Description
	has to send a DHCP request. This web page will need to be refreshed when this change is made.
IP Address	The IP address of the network interface. The factory default value is 0.0.0.0 4 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.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.5 IPv6 Network Connectivity Configuration

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.

4 CLI commands are not available for all the IPv6 menus.

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

3 Configuring and viewing System Information

network by using a Web browser, SNMP, Telnet, or SSH, you must first configure the device with the appropriate IPv6 information.

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

Figure 19: IPv6 Network Connectivity Configuration

Table 14: 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. - a Enabled – The network interface can acquire an IPv6 address through IPv6 Neighbor Discovery Protocol (NDP) and through the use of Router Advertisement messages. - a Disabled – The network interface will not use the native IPv6 address autoconfiguration features to acquire an IPv6 address.
DHCPv6 Client DUID	The client identifier used by the DHCPv6 client (if enabled) when sending messages to the DHCPv6 server.
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: - a New IPv6 Address – Specify the IPv6 address to add to the interface. - a EUI Flag – Select this option to enable the Extended Universal Identifier (EUI) flag for IPv6 address, or clear the option to omit the flag. - 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.

Field	Description
Default IPv6 Routers	Lists the IPv6 address of each default router that has been automatically configured through IPv6 router discovery.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to update the information on the screen.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.6 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 20: Network Port IPv6 Neighbors

Table 15: 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.
Type	The type of the neighbor entry, which is one of the following: - a Static – The neighbor entry is manually configured. - a Dynamic – The neighbor entry is dynamically resolved. - a Local – The neighbor entry is a local entry. - a Other – The neighbor entry is unknown.
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: a 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.

Field	Description
	a 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. a 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. a Probe – A reachability confirmation is actively sought by resending neighbor solicitation messages every RetransTimer milliseconds until a reachability confirmation is received. a Unknown – The reachability status cannot be determined.
Last Updated	Displays the time since the address was confirmed to be reachable.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To add network port static IPv6 neighbor entry, click **Add** and configure the desired settings.

Add Network Port IPv6 Neighbor

IPv6 Address

(xxxxxxxx)

MAC Address

(xxxxxxxxxxxx)

Submit

Cancel

Figure 21: Add Network Port IPv6 Neighbor

Table 16: 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.

- a To remove network port static IPv6 neighbor entries, select each static neighbor entry to remove and click **Remove**.

3.7.7 Service Port IPv4

Some platforms have a built-in service port that can serve as a dedicated network management interface. 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. 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 22: Service Port IPv4 Configuration

Table 17: Service Port IPv4 Configuration Fields

Field	Description
Service Port Configuration Protocol	Specify what the switch should do following power-up. The factory default is DHCP . The options are as follows: - a None: Do not send any requests following power-up. - a BootP: Transmit a BootP request. - a DHCP: Transmit a DHCP request.
IP Address	The IP address of the network interface. The factory default value is 0.0.0.0 4 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.
Interface Status	Indicates, whether an Ethernet link has been established (Up) or not (Down).
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.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To renew the IPv4 address learned from a DHCP server on the service port, click **Renew DHCP Lease**.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.8 Service Port IPv6

Some platforms have a built-in service port that can serve as a dedicated network management interface. 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

3 Configuring and viewing System Information

be switched or routed to the operational network. For systems that have the service port, the Service Port IPv6 Configuration page allows you to configure network information for the switch.

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

Figure 23: Service Port IPv6 Configuration

Table 18: 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. Enable – The service port can acquire an IPv6 address through IPv6 Neighbor Discovery Protocol (NDP) and through the use of Router Advertisement messages. Disable – The service port will not use the native IPv6 address autoconfiguration features to acquire an IPv6 address.
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: - To add an entry to the list, click the + (plus) button to open the Add IPv6 Address dialog and provide the following: New IPv6 Address – Specify the IPv6 address to add to the service port interface. EUI Flag – Select this option to enable the Extended Universal Identifier (EUI) flag for IPv6 address, or clear the option to omit the flag. - 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 service port interface that have been dynamically configured through IPv6 autoconfiguration or DHCPv6.

Field	Description
Default IPv6 Routers	Lists the IPv6 address of each default router that has been automatically configured through IPv6 router discovery.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.9 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 24: Service Port IPv6 Neighbors

Table 19: 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: - a Static – The neighbor entry is manually configured. - a Dynamic – The neighbor entry is dynamically resolved. - a Local – The neighbor entry is a local entry. - a Other – The neighbor entry is unknown.
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: a 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.

Field	Description
	a 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. a 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. a Probe – A reachability confirmation is actively sought by resending neighbor solicitation messages every RetransTimer milliseconds until a reachability confirmation is received. a Unknown – The reachability status cannot be determined.
Last Updated	Displays the time since the address was confirmed to be reachable.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To add service port static IPv6 neighbor entry, click **Add** and configure the desired settings.

Add Service Port IPv6 Neighbor

IPv6 Address

(xxxxxxxx)

MAC Address

(xxxxxxxxxxxx)

Submit

Cancel

Figure 25: Add Service Port IPv6 Neighbor

Table 20: 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.

- a To remove service port static IPv6 neighbor entries, select each static neighbor entry to remove and click **Remove**.

3.7.10 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 26: DHCP Client Options

Table 21: 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.

Use the buttons to perform the following tasks:

- a If you change any of the DHCP client option 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.11 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 27: System Connectivity Configuration

Table 22: 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 HTTPS. 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.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.12 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.

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 28: Telnet Session Configuration

Table 23: 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. 4 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 terminated. You may enter any number from 1 to 160. The factory default is 10. 4 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 menu, select how many simultaneous Telnet sessions to allow. The maximum is 5, 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: - a Yes – Permits new Telnet sessions until the maximum number allowed is reached. - a No – New Telnet sessions will not be allowed, but existing sessions are not disconnected.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.13 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.

3 Configuring and viewing System Information

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

Figure 29: Outbound Telnet Configuration

Table 24: 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.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.14 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 30: Serial Port

Table 25: 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 terminates the connection. Enter a number between 0 and 160. The factory default is 10. 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. It is always 1 .
Flow Control	Whether hardware flow control is enabled or disabled. It is always disabled.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.15 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 31: CLI Banner Configuration**Table 26: CLI Banner Configuration Fields**

Field	Description
CLI Banner Message	Text area for creating, viewing, or updating the CLI banner message. 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.

Use the buttons to perform the following tasks:

- a 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.

3 Configuring and viewing System Information

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Clear** to delete the CLI banner message from the device. 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**.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.16 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.

HTTP Configuration	
HTTP Admin Mode	☐ Disable ☒ Enable
HTTP to HTTPS Redirect	☒ Disable ☐ Enable
HTTP Port	80 (1025 to 65535, 80 = Default)
HTTP Session Soft Time Out (Minutes)	5 (1 to 60)
HTTP Session Hard Time Out (Hours)	24 (1 to 168)
Maximum Number of HTTP Sessions	16 (1 to 16)

Figure 32: HTTP Configuration

Table 27: 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 HTTPS.
HTTP Port	The TCP port number on which the HTTP server listens for requests. Existing HTTP login sessions are terminated whenever this value is changed. All new HTTP sessions must use the new port number. 4 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. The default value is 5 minutes. The currently configured value is shown when the web page is displayed.
HTTP Session Hard Timeout	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.

Use the buttons to perform the following tasks:

- a If you make changes to the page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to refresh the page with the most current data from the switch.

- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.17 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 33: HTTPS Configuration

Table 28: 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. 4 Before changing this value, check your system (for example, using netstat) 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. a Present – The certificate has been generated and is present on the device

Field	Description
	a Absent – Certificate is not available on the device a Generation In Progress – An SSL certificate is currently being generated.  The Download Certificates button allows you to download an SSL certificate file from a remote system to the device. 4 To download SSL certificate files, SSL must be administratively disabled.  The Generate Certificate button allows you to generate an SSL certificate to use for secure communication between the web browser and the embedded web server on the device.  The Delete Certificates button allows you to delete the SSL certificate. This button is available only if an SSL certificate is present on the device.

Use the buttons to perform the following tasks:

- a If you make changes to the page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

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

Download Certificates

File Type

SSL Trusted Root Certificate PEM File

Certificate Index

0 (1 to 8; 0 for None)

Select File

Datei auswählen

Keine ausgewählt

Status

Begin Transfer

Close

Figure 34: Download Certificates

Table 29: Download Certificates Fields

Field	Description
File Type	Specify the type of file to transfer from the device to a remote system. The following options are available: - a SSL Trusted Root Certificate PEM File - a SSL Server Certificate PEM File - a SSL DH Weak Encryption Parameter PEM File

Field	Description
	a SSL DH Strong Encryption Parameter PEM File
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.

3.7.18 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 35: SSH Configuration

Table 30: 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. 4 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.
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. 4 This is the only supported SSH version and is enabled by default. Clearing this option is not permitted.

3 Configuring and viewing System Information

Field	Description
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. 4 LCOS SX 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. 4 LCOS SX only supports 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.

Use the buttons to perform the following tasks:

- a If you make changes to the page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.19 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 36: Management Access List Configuration

Table 31: User Accounts Fields

Field	Description
Profile Configuration	
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.
Profile Rule Configuration	
Interface	The port/interface or trunk ID.
Management Method	The types of action will be taken on access control list. a Permit – To allow conditions for the management access list. a Deny – To deny conditions for the management access list.
Source IP Address	IP Address of device to be permitted or denied management access.
Subnet Mask	Specifies the network mask of the source IP address.
Service	The type of service to permit or deny: a ANY a TELNET a HTTP a HTTPS

3 Configuring and viewing System Information

Field	Description
	- a SNMP - a SSH - a TFTP - a SNTP
Priority	Priority for the rule. Duplicates are not allowed.

Use the buttons to perform the following tasks:

4 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.

- a To add a new MACAL, enter a descriptive name for the **Access Profile** and click **Submit**.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Add** to create a new MACAL and specify the rule criteria in the available fields.

Figure 37: Add Profile Rule

Table 32: Add Profile Rule Fields

Field	Description
Action	Specify if access is to be denied (Deny) or to be allowed (Permit).
Interface	The port/interface or trunk ID. The following options are available: - a None - a Port - a VLAN - a Port Channel
Port	Select the physical interface to be used in the profile. Only available when setting the Interface to Port .
VLAN ID	Select the VLAN ID to be used in the profile. Only available when setting the Interface to VLAN .
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. Only available when setting the Interface to Port Channel .
Source IP Address /Subnet Mask	Tick the checkbox and enter an IP address and a corresponding subnet mask of a device to be permitted or denied management access.

Field	Description
Service	Tick the checkbox and select a service type. The following options are available: - a ANY - a TELNET - a HTTP - a HTTPS - a SNMP - a SSH - a TFTP - a SNTP
Priority	The priority for the profile rule. Duplicate priorities are not allowed. If the priority is not specified during rule creation, the rule is automatically assigned the first free priority before adding it to the management access list. The rules are displayed based on their priority within the management access list.

- a 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.
- a To remove a Profile Rule, select one or more table entries and click **Remove** to delete the selected entries.

3.7.20 User Accounts

By default, the switch contains one user account:

- a 'admin', with 'Read/Write' privileges

This account has a blank password by default. The name is not case sensitive.

- 4 The preconfigured user 'admin' is assigned to a pre-configured list named **default-usergroup-name** (see menu **Systems Users User Groups**), which you cannot delete. All newly created users are also assigned to the **default-usergroup-name** until you specifically assign them to a different list.

If you log on to the switch with the user account that Read/Write privileges (that is, 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 accounts. You can delete all accounts except for the 'admin' account.

- 4 Only a user with Read/Write privileges may alter data on this screen.

3 Configuring and viewing System Information

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

Accounts Auth Server Users Sessions User Domain Name Task Groups User Groups

Configuration

Display **All** rows Showing 1 to 1 of 1 entries Filter:

	User Name	Access Level	Lockout Status	Password Override	Password Expiration	Contained User Group	Operational Permission
☐	admin	Privilege-15	False	Disable		default-usergroup-name	AAA: Read, Write, Execute, Debug OSPF: Read, Write, Execute, Debug

First Previous 1 Next Last

User Groups

Table is Empty

Refresh Add Edit Remove

Figure 38: User Accounts

Table 33: User Accounts Fields

Field	Description
User Name	Enter the name you want to give to the new account. Users can have from 1 - 64 characters in length and are not case sensitive. Valid characters include all the alphanumeric characters and the dash ('-') and underscore ('_') characters. User name <i>default</i> is not valid. 4 The name of the <i>admin</i> user cannot be edited.
Access Level	Indicates the access or privilege level for this user. The options are: - a Privilege-0 – The user exists but is not permitted to log on to the device. - a Privilege-1 – The user can view the configuration but cannot modify any fields. - a Privilege-15 – The user can view and modify the configuration.
Lockout Status	Provides the current lockout status for this user. If the lockout status is True, the user cannot access the management interface even if the correct username and password are provided. The user has been locked out of the system due to a failure to supply the correct password within the configured number of login attempts. With the lockout status False the user can access the management interface.
Password Override	Identifies the password override complexity status for this user. - a Enable - The system does not check the strength of the password. - a Disable - When configuring a password, it is checked against the Strength Check rules configured for passwords on the Password Rules page.
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 Rules 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.

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

- a Click **Refresh** to update the information on the screen with the most current data.
- a To add a user, click **Add**. The Add new user dialog box opens. Specify the new account information in the available fields.

Figure 39: Add new User

Table 34: Add new user 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 can have from 1 - 64 characters in length and are not case sensitive. Valid characters include all the alphanumeric characters and the dash ('-') and underscore ('_') characters. User name <i>default</i> is not valid. 4 The name of the <i>admin</i> user cannot be edited.
Password	Enter the optional new or changed password for the account. The password must have between 8 to 64 characters. It does not display as it is typed, only asterisks (*) or dots (.) appear based on the browser used.
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: - a Privilege-0 – The user exists but is not permitted to log on to the device. - a Privilege-1 – The user can view the configuration but cannot modify any fields. - a Privilege-15 – The user can view and modify the configuration.
Password Override	Identifies the password override complexity status for this user. - a Enable - The system does not check the strength of the password. - a Disable - When configuring a password, it is checked against the Strength Check rules configured for passwords.
Password Strength	Shows the status of password strength check.

- a 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.
- a To remove a user, select one or more table entries and click **Remove** to delete the selected entries. The 'admin' user cannot be deleted. An error message is displayed if you try to delete it.

3.7.21 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 user names and passwords to provide authentication to users instead of using an external authentication server.

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-password
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.




Figure 40: Auth Server Users

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.

- a To add a user to the local authentication server database, click **Add** and complete the required information.

Figure 41: Add new user

Table 35: Add new 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	Reenter the password to confirm the entry.
Encrypted	Select this option to encrypt the password before it is stored on the device.

- a To change the password information for an existing user, select the user to update and click **Edit**.
- a To delete a user from the database, select each user to delete and click **Remove**.
- a To remove all users from the database, click **Clear All Users**.
- a Click **Submit** to apply the changes to the system. You must perform a save to make the changes persist across a reboot.

3.7.22 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 42: Logged In Sessions

Table 36: 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 one of the following options: - a Telnet - a Serial - a SSH - a HTTP - a HTTPS

Click **Refresh** to update the information on the screen.

3.7.23 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 43: User Domain Name

Table 37: 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\user name 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 login. When only the user name is provided, the device sends the user name as domain-name\user name, 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.

Use the buttons to perform the following tasks:

- a Click **Submit** to to apply the changes to the system. To preserve the changes after a system reboot, you must perform a save.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.24 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.

3 Configuring and viewing System Information

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 Groups** in the navigation menu.




Figure 44: Task Group Configuration


Table 38: Task Group Configuration Fields

Field	Description
Task Group	The task group name.
Description	The associated description for the task group name.
Parent Task Groups	The associated parent task groups for the task group name.
Configured Permission	The configured task permissions for the task group. The following options are available: - a AAA - a Read - a Write - a Execute - a Debug - a OSPF - a Read - a Write - a Execute - a Debug
Operational Permission	The operational task permissions. The following options are available: - a AAA - a Read - a Write - a Execute - a Debug - a OSPF - a Read

Field	Description
	- a Write - a Execute - a Debug
Task Group Name	Displays the selected task group.
Description	You can optionally enter a description.
Parent Task Group	The associated parent task groups for the task group name. To configure this parent task group, click the + icon (plus) in the header row. To remove the parent task group, click the - icon (minus) in the row.
Configured Tasks	The list of task names. To configure this task, click the + icon (plus) in the header row. To remove the task permissions, click the - icon (minus) in the row. The following tasks are available: - a AAA - a OSPF
Permissions	The task permissions. - a Read - a Write - a Execute - a Debug

Use the buttons to perform the following:

- a Click **Refresh** to update the information on the screen.
- a To add a task group, click **Add** and specify a name for the group.




The 'Add Task Group' dialog box features a title bar with a close button. It contains a text input field labeled 'Task Group Name' with a character count '(1 to 128 characters)' to its right. At the bottom right, there are 'Submit' and 'Cancel' buttons.

Figure 45: Add Task Group

- a To remove a task group, select the check box associate to the group to remove and click **Remove**.

Add Parent Task Group:

Click on the + icon (plus) next to **Parent Task Group** to assign a parent task group to the selected task group. To remove the parent task group, click the - icon (minus) in the row.



The 'Add Parent Task Group' dialog box has a title bar with a close button. It contains two fields: 'Task Group Name' and 'Parent Task Group'. The 'Parent Task Group' field is a dropdown menu currently showing 'default-taskgroup-name'. 'Submit' and 'Cancel' buttons are located at the bottom right.

Figure 46: Add Parent Task Group

Table 39: Add Parent Task Group Fields

Field	Description
Task Group Name	Displays the selected task group.
Parent Task Group	Select an existing task group.

Set Task Permission:

Click on the + icon (plus) next to **Configured Tasks** to assign permissions to the selected task group. To remove the task permissions, click the - icon (minus) in the row.

The screenshot shows a window titled "Set Task Permission". Inside, there's a form with two main sections. The first section has "Task Group Name" and "Task-Group" (a dropdown menu currently showing "AAA"). The second section lists four permissions: "Read", "Write", "Execute", and "Debug", each followed by an unchecked checkbox. At the bottom right of the window are "Submit" and "Cancel" buttons.

Figure 47: Set Task Permission

Table 40: Set Task Permission Fields

Field	Description
Task Group Name	Displays the selected task group.
Task Name	Select the task. The following tasks are available: a AAA a OSPF
Read	Activate the checkbox to allow read access.
Write	Activate the checkbox to allow write access.
Execute	Activate the checkbox to allow execute access.
Debug	Activate the checkbox to allow debug access.

3.7.25 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.

Accounts Auth Server Users Sessions User Domain Name Task Groups **User Groups**

User Group Configuration ?

Display **All** rows Showing 1 to 1 of 1 entries Filter:

	User Group	Description	Parent User Groups	Contained Task Group	Operational Permission
☐	default-usergroup-name			default-taskgroup-name	AAA: Read, Write, Execute, Debug OSPF: Read, Write, Execute, Debug

First Previous 1 Next Last

User Group Name

Description

Parent User Groups + - Contained Task Group + -

Table is Empty Table is Empty

Refresh Add Remove

Figure 48: User Group Configuration

Table 41: 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.
Contained Task Group	The associated task groups for user group.
Operational Permission	The operational task permissions for the user group. - a AAA - a Read - a Write - a Execute - a Debug - a OSPF - a Read - a Write - a Execute - a Debug
User Group Name	Displays the selected user group.
Description	You can optionally enter a description.
Parent User Groups	The associated parent user groups for the user group. To configure this parent user group, click the + icon (plus) in the header row. To remove the parent user group, click the - icon (minus) in the row.
Contained Task Group	The associated task groups for user group. To configure this task group, click the + icon (plus) in the header row. To remove the task group, click the - icon (minus) in the row.

Use the buttons to perform the following:

- a Click **Refresh** to update the information on the screen.
- a To add a user group, click **Add** and specify a name for the group.


A dialog box titled "Add User Group" with a close button (X) in the top right corner. It contains a text input field labeled "User Group Name" with a placeholder "(1 to 128 characters)". At the bottom right, there are two buttons: "Submit" and "Cancel".

Figure 49: Add User Group

- a To remove a user group, select the check box associate to the group to remove and click **Remove**.

Add Parent User Group:

Click on the + icon (plus) next to **Parent User Groups** to assign a parent user group to the selected user group. To remove the parent user group, click the - icon (minus) in the row.


A dialog box titled "Add Parent User Group" with a close button (X) in the top right corner. It contains two rows of fields. The first row has a label "User Group Name" and a text input field containing "User-Group". The second row has a label "Parent User Groups" and a dropdown menu showing "default-usergroup-name". At the bottom right, there are two buttons: "Submit" and "Cancel".

Figure 50: Add Parent User Group

Table 42: Add Parent User Group Fields

Field	Description
User Group Name	Displays the selected user group.
Parent User Groups	Select an existing user group.

Contained Task Group:

Click on the + icon (plus) next to **Contained Task Group** to assign a task group to the selected user group. To remove the contained task group, click the - icon (minus) in the row.


A dialog box titled "Add Task Group" with a close button (X) in the top right corner. It contains two rows of fields. The first row has a label "User Group Name" and a text input field containing "User-Group". The second row has a label "Task Group Name" and a dropdown menu showing "default-taskgroup-name". At the bottom right, there are two buttons: "Submit" and "Cancel".

Figure 51: Add Task Group

Table 43: Add Task Group Fields

Field	Description
User Group Name	Displays the selected user group.

Field	Description
Task Group Name	Select an existing task group created on the Task Group Configuration page.

3.7.26 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 52: Accounting List Configuration

Table 44: Accounting List Configuration Fields

Field	Description
List Name	The name of the accounting list. This field can be configured only when adding a new accounting list.
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. Dot1x – Provides accounting for DOT1X user commands, sent to an external RADIUS server.
Record Type	Indicates when to record and send information about the user activity: a 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. a StopOnly – Accounting notifications are sent at the end of an exec session or a user-executed command. a Undefined – The user has not yet been created and therefore no accounting information is available. a None – Accounting will not be notified.
Method Options	The methods used to record user activity. The possible methods are as follows: a TACACS – Accounting notifications are sent to the configured TACACS+ server.

3 Configuring and viewing System Information

Field	Description
	a Radius – Accounting notifications are sent to the configured RADIUS server (only available for Accounting Type options Exec and Dot1x).
List Type	The type of accounting list, which is one of the following: a 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. a Configured – The list has been added by a user.
Access Line	The access methods that use the list for accounting user activity. The settings for this field are configured on the Accounting List Selection page.

Use the buttons to perform the following tasks:

- a To configure a new accounting list, click **Add**.

Figure 53: Add New Accounting List

Table 45: Add New Accounting List Fields

Field	Description
Accounting List	
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. Dot1x – Provides accounting for DOT1X user commands, sent to an external RADIUS 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: a 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.

Field	Description
	- a StopOnly – Accounting notifications are sent at the end of an exec session or a user-executed command. - a Undefined – The user has not yet been created and therefore no accounting information is available. - a None – Accounting will not be notified.
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. The possible methods are as follows: - a TACACS – Accounting notifications are sent to the configured TACACS+ server. - a Radius – Accounting notifications are sent to the configured RADIUS server (only available for Accounting Type options Exec and Dot1x).
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.

- a To edit a list, select the entry to modify and click **Edit**. The settings that can be edited depend on the list type.
- a 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.
- a 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.
- a 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.

3.7.27 Accounting List Selection

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

- a **Exec** – The accounting list to record user login and logout times.
- a **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 Selection Configuration page, click **System > AAA > Accounting Selection** in the navigation menu.

Figure 54: Accounting List Selection

Table 46: Accounting List Selection Fields

Field	Description
Terminal	The access methods in this section are CLI-based. - a 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. - a Telnet — The Exec accounting list and the Commands accounting list to apply to users who access the CLI by using a Telnet session. - a 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.
Hypertext Transfer Protocol	The access methods in this section are through a web browser. - a HTTP — The Exec accounting list to apply to users who access the web-based management interface by using HTTP. - a HTTPS — The Exec accounting list to apply to users who access the web-based management interface by using HTTPS.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.28 Authentication List Configuration

Use the Authentication List Configuration 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 methods 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 Configuration page, click **System > AAA > Authentication List** in the navigation menu.

Authentication List Configuration

Display **All** rows Showing 1 to 9 of 9 entries Filter:

	List Name	Access Type	Method Options	List Type	Access Line	
☐	defaultList	Login	Local	Default	Console	⏻
☐	networkList	Login	Local	Default	Telnet,SSH	⏻
☐	noauthList	Login	None	Default		⏻
☐	enableList	Enable	Enable,None	Default	Console,Telnet,SSH	⏻
☐	enableNetList	Enable	Enable,Deny	Default		⏻
☐	noenableList	Enable	None	Default		⏻
☐	httpList	HTTP	Local	Default	HTTP	⏻
☐	httpsList	HTTPS	Local	Default	HTTPS	⏻
☐	dot1xList	Dot1x		Default	Dot1x	⏻

First Previous 1 Next Last

Refresh Add Edit

Figure 55: Authentication List Configuration

Table 47: Authentication List Configuration 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: - a 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. - a 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. - a HTTP – Management-level access to the web-based user interface by using HTTP. - a HTTPS – Management-level access to the web-based user interface by using HTTPS. - a Dot1x – Port-based access to the network through a switch port that is controlled by IEEE 802.1X.
Method Options	The methods used to authenticate a user who attempts to access the management interface or network. The possible methods are as follows: - a Enable – Uses the locally configured Enable password to verify the user's credentials. - a Local – Uses the ID and password in the Local User database to verify the user's credentials. - a RADIUS – Sends the user's ID and password to the configured RADIUS server to verify the user's credentials. - a TACACS – Sends the user's ID and password to the configured TACACS+ server to verify the user's credentials.

3 Configuring and viewing System Information

Field	Description
	- a None – No authentication is used. - a Deny - Denies authentication. - a IAS – Uses the local Internal Authentication Server (IAS) database for 802.1X port-based authentication (only available via the dot1xList).
List Type	The type of list, which is one of the following: - a Default – The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options are configurable. - a Configured – The list has been added by a user.
Access Line	The access methods that use the list for authentication. The settings for this field are configured on the Authentication Selection page.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a To configure a new authentication list, click **Add**.

Figure 56: Add New Authentication List

Table 48: Add New Authentication List Fields

Field	Description
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: - a 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. - a 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.
List Name	The name of the authentication list. This field can be configured only when adding a new authentication list.
Authentication Methods	

Field	Description
This area includes the Available Methods and Selected Methods fields. For lists that allow multiple authentication 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 authenticate the user. For example, if the selected methods are Enable, followed by None, a user who fails to authenticate with the enable password is granted access anyway because the final method indicates that no authentication is required.	
Available Methods	The authentication methods that can be used for the authentication list. Not all authentication methods are available for all lists. To set the authentication 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 authentication 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 authenticate a user. If the user fails to be authenticated using the first method, the device attempts to verify the user's credentials by using the next method in the list. No authentication 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.

- a To edit a list, select the entry to modify and click **Edit**. The settings that can be edited depend on the list type.
- a To remove a non-default authentication list, click the **Remove** icon associated with the entry. You must confirm the action before the entry is deleted.
- a To reset the Method Options for a default authentication list to the factory default values, click the **Reset** icon associated with the entry. You must confirm the action before the entry is reset.

Additional information regarding user authentication can be found in the following chapters:

- a To create a new authentication list, see [Authentication Server Users](#) on page 64.
- a To assign users to a specific authentication list, see [User Accounts](#) on page 61.
- a To configure the 802.1x port security users, see [RADIUS Settings](#) on page 450.

3.7.29 Authentication List Selection

Use the Authentication Selection 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:

- a **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.
- a **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.

3 Configuring and viewing System Information

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

Figure 57: Authentication List Selection

Table 49: Authentication List Selection 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.
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.

Use the buttons to perform the following tasks:

- a **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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.7.30 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 (with the exception of the entries with the **List Type** option **Configured**). The **List Name** and **Authorization Type** settings for any list (default or configured) cannot be changed.

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

	List Name	Authorization Type	Method Options	List Type	Access Line
☐	dfitCmdAuthList	Commands	None	Default	Console, Telnet, SSH
☐	noCmdAuthList	Commands	None	Configured	
☐	dfitExecAuthList	Exec	None	Default	Console, Telnet, SSH
☐	noExecAuthList	Exec	None	Configured	
☐	networkList	Network		Default	Dot1x

Figure 58: Authorization List Configuration

Table 50: 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. This field can be configured only when adding a new authorization list, and only the Commands and Exec authorization types can be selected. The authorization types are as follows: - a Commands – 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. - a Exec – Determines whether a user can bypass User EXEC mode and enter Privileged EXEC mode directly after a successful Login authentication. - a 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.
Method Options	The methods used to authorize a user's access to the device or network services. The possible methods are as follows: - a 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. - a 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. - a Local – Uses a list stored locally on the system to determine whether the user is authorized to access the given services. - a None – No authorization is used. If the method is None, the authorization type is effectively disabled.

3 Configuring and viewing System Information

Field	Description
List Type	The type of authorization list, which is one of the following: - a Default – The list is preconfigured on the system. This type of list cannot be deleted, and only the Method Options are configurable. - a Configured – The list has been added by a user.
Access Line	The access methods that use the list for authorization. The settings for this field are configured on the Authorization Selection page.

Use the buttons to perform the following tasks:

- a To configure a new authorization list, click **Add**.

Figure 59: Add New Authorization List

Table 51: Add New Authorization List Fields

Field	Description
Authorization List	
Authorization Type	The type of authorization list, which is one of the following: - a Commands – 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. - a Exec – Determines whether a user can bypass User EXEC mode and enter Privileged EXEC mode directly after a successful Login authentication.
List Name	The name of the authorization list. This field can be configured only when adding a new authorization list.
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.

Field	Description
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.

- a To edit a list, select the entry to modify and click **Edit**. The settings that can be edited depend on the list type. When editing an existing authentication list, only the Authorization Methods can be configured.
- a 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.
- a 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.

3.7.31 Enable Password

Use the Enable Password page to configure the enable password. This password is only used on the CLI to switch to privileged mode via the command **enable**.

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

Figure 60: Enable Password Configuration

Table 52: Enable Password Configuration Fields

Field	Description
Enable Password	Specify the password all users must enter after executing the enable command at the CLI prompt. 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.
Confirm Enable Password	Confirms the new enable password. The password appears in the ***** format.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.
- a To remove the enable password, click the **Remove** button. You must confirm the action before the password is removed.

3.7.32 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.

The screenshot shows the 'Password Rules' configuration page. At the top, there's a navigation bar with 'System > Passwords > Password Rules' and buttons for 'Save Configuration', 'Hide Device View', and 'Log Out'. Below this is a breadcrumb trail: 'System' > 'Switching' > 'Routing' > 'Security' > 'QoS' > 'Stacking'. The main content area has tabs for 'Line Password', 'Enable Password', 'Password Rules' (selected), 'Last Password', and 'Reset Passwords'. The 'Password Rules' tab is active, showing a table of settings:

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 the table is a section for 'Exclude Keyword Name' with a search bar and a table that is currently empty. At the bottom, there are 'Submit', 'Refresh', and 'Cancel' buttons.

Figure 61: Password Rules

Table 53: Password Rules Fields

Field	Description
Minimum Length	Passwords must have at least this many characters (0 to 64).
Aging (Days)	The number of days that a user password is valid from the time the password is set. When 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.
Minimum Number of Lowercase Letters	Specify the minimum number of lowercase letters a password must include.

Field	Description
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 <i>aaaa</i> .
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 <i>abcd</i> .
Minimum Character Classes	Specify the minimum number of character classes a password must contain. There are four character classes: a Uppercase a Lowercase a Numbers a Special Characters
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 <i>pass</i> is an excluded keyword, passwords such as <i>23passA2c</i> , <i>ssapword</i> , and <i>PASsWoRD</i> are prohibited. Use the plus and minus buttons to perform the following tasks: a 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 . a To remove a keyword from the list, click the – (minus) button associated with the keyword to remove and confirm the action. a To remove all keywords from the list, click the – (minus) button in the header row and confirm the action.

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.

Exclude Keyword Name:

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.

Figure 62: Add

3.7.33 Denial of Service

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

3 Configuring and viewing System Information

- a SIP=DIP: Source IP address = Destination IP address.
- a First Fragment: TCP Header size smaller then configured value.
- a TCP Fragment: IP Fragment Offset = 1.
- a 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.
- a L4 Port: Source TCP/UDP Port = Destination TCP/UDP Port.
- a ICMP: Limiting the size of ICMP Ping packets.

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

Denial of Service Configuration

TCP Settings

First Fragment	☐
TCP Port	☐
UDP Port	☐
SIP=DIP	☐
SMAC=DMAC	☐
TCP FIN and URG and PSH	☐
TCP Flag and Sequence	☐
TCP SYN	☐
TCP SYN and FIN	☐
TCP Fragment	☐
TCP Offset	☐
Port D-Disable	☐
Min TCP Hdr Size	20 (0 to 255)

ICMP Settings

ICMP	☐
Max ICMPv4 Size	512 (0 to 16376)
ICMPv6	☐
Max ICMPv6 Size	512 (0 to 16376)
ICMP Fragment	☐

Submit Refresh Cancel

Figure 63: Denial of Service Configuration

Table 54: 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.

3 Configuring and viewing System Information

Field	Description
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.
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.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.8 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.

3.8.1 Basic 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 64: Basic Switch Configuration

Table 55: 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. 4 IEEE 802.1D recommends a default of 300 seconds, which is the factory default.

Use the buttons to perform the following tasks:

- a Click **Submit** to apply the changes to the system. You must perform a save to make the changes persist across a reboot.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.9 Configuring LMC Configuration

Use this page to configure the LANCOM Management Cloud (LMC) settings.

To access the LMC Configuration Page, click **System > LMC > Configuration** in the navigation menu.

Field	Value	Range
LMC Domain	cloud.lancom.de	(1 to 255)
LMC Rollout Project ID		(0 to 36)
LMC Rollout Location ID		(0 to 36)
LMC Rollout Role		(0 to 36)
Configuration Via DHCP	☒	
DHCP Client Auto Renew	☒	
Operating	Try	
LMC Activation Code		(1 to 47)
LMC Certificate Status	Absent	

Figure 65: LMC Configuration

Table 56: LMC Configuration Fields

Field	Description
LMC Domain	Specifies the LMC Domain. The default setting cloud.lancom.de is used to establish a connection to the LANCOM Systems public cloud. Only change this setting if you are using a private cloud.
LMC Rollout Project ID	LMC Project ID rollout information string.
LMC Rollout Location ID	LMC Location ID rollout information string.
LMC Rollout Role	LMC Role rollout information string.
Configuration Via DHCP	Enable LMC configuration via DHCP option 43: a Enabled – Use URLs and rollout information configured by DHCP option 43, if present. a Disabled – Always use static configuration.
DHCP Client Auto Renew	Renew DHCP lease early when connection to LMC is not possible. Only has effect when DHCP client is enabled. a Enabled – Enable DHCP-Client-Auto-Renew a Disabled – Disable DHCP-Client-Auto-Renew
Operating	Controls if the switch is managed by the LMC or operated as a standalone device. a Yes – LMC is operating the switch. a No – Switch is not operated by the LMC. a Try – During “Try” mode LMC client is enabled after each boot for 24 hours. If any cloud management configuration functionality is triggered by LMC within that “Try” period the LMC feature will stay persistently enabled, even after rebooting the switch. Otherwise, the LMC client will automatically be disabled after 24 hours until next reboot.

3 Configuring and viewing System Information

Field	Description
LMC Activation Code	Insert the LMC Activation Code generated by the LMC to pair the device with the LMC when not using zero touch pairing. The device can only be paired as long as it is not already paired. To pair it again, delete the LMC Certificate below.
LMC Certificate Status	Present / Absent state of the LMC Device Certificate. The certificate is generated when successfully pairing with the LMC. To disconnect the device from the cloud or pair it again, the certificate can be deleted via the minus icon (-) next to the status field.

Use the buttons to perform the following tasks:

- a Click **Submit** to apply the changes to the system. You must perform a save to make the changes persist across a reboot.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.9.1 LMC Status

Use this page to view the LANCOM Management Cloud (LMC) status.

To access the LMC Status page, click **System > LMC > Status**.

Configuration

Status

LMC Status

?

Status

Management Status	Unpaired
Monitor Status	Disabled
Control Status	Disabled
Stacking Status	Stacking disabled
Config Modified	No
Activation Code Present	No
Zero Touch Support	No
Customer Device ID	
Round Trip Time	0 ms
Active LMC Domain	cloud.lancom.de
Active LMC Rollout Project ID	
Active LMC Rollout Location ID	
Active LMC Rollout Role	

Transport Status

Service Name	HTTP Requests	Request Errors	TX Bytes	RX Bytes
Control-Service	0	0	0	0
Monitor-Service	0	0	0	0

Refresh

Figure 66: LMC Status

Table 57: LMC Status Fields

Field	Description
Management Status	Shows the current Management Status of the device. Indicates whether the device is paired (controlled by the LMC) or unpaired.
Monitor Status	Shows the current Monitor Status of the device. Indicates whether the Monitor Status is Enabled or Disabled.

Field	Description
Control Status	Shows the current Control Status of the device. Indicates whether the Control Status is Enabled or Disabled.
Stacking Status	A preconfigured stack of switches can be paired with the LMC when using LCOS SX as of version 5.20 Rel. Stacking configuration via the LMC smartconfig will be available in a future LMC version.
Config Modified	Flag, Config was modified by LMC: a The LMC uses this flag to decide whether a config reset has been performed and if the configuration needs to be rolled out again.
Activation Code Present	Indicates whether an activation code is given.
Zero Touch Support	Indicates whether the zero-touch support (the possibility to assign a configuration automatically) is given.
Customer Device ID	LMC-assigned device ID.
Round Trip Time	Round trip time of last request to LMC.
Active LMC Domain	Currently used LMC Domain, for example cloud.lancom.de.
Active LMC Rollout Project ID	LMC Project ID rollout information string.
Active LMC Rollout Location ID	LMC Location ID rollout information string.
Active LMC Rollout Role	LMC Role rollout information string.

Click **Refresh** to update the screen with the most current information.

3.10 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.

3.10.1 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.

3 Configuring and viewing System Information

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

The screenshot shows the 'Log Configuration' page. It features a top navigation bar with tabs: 'Buffered Log', 'Event Log', 'Persistent Log', 'Hosts', 'Configuration' (active), 'Source Interface Configuration', and 'Statistics'. Below the tabs is a blue header 'Log Configuration' with a help icon. The main content area is divided into five sections: 'Buffered Log Configuration', 'Command Logger Configuration', 'Console Log Configuration', 'Persistent Log Configuration', and 'Syslog Configuration'. Each section contains configuration fields for 'Admin Mode' (radio buttons for 'Enable' and 'Disable') and 'Severity Filter' (dropdown menu). The 'Syslog Configuration' section also includes 'Protocol Version' (radio buttons for 'RFC 3164' and 'RFC 5424') and 'Local UDP Port' (text input with a range of 1 to 65535). At the bottom are 'Submit', 'Refresh', and 'Cancel' buttons.

Figure 67: Log Configuration

Table 58: Log Configuration Fields

Field	Description
Buffered Log Configuration	
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).
Command Logger Configuration	
Admin Mode	Enable or disable logging of the command line interface (CLI) commands issued on the device.
Console Log Configuration	
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: - a Emergency (0) – The device is unusable. - a Alert (1) – Action must be taken immediately. - a Critical (2) – The device is experiencing primary system failures. - a Error (3) – The device is experiencing non-urgent failures. - a Warning (4) – The device is experiencing conditions that could lead to system errors if no action is taken. - a Notice (5) – The device is experiencing normal but significant conditions.

Field	Description
	- a Info (6) – The device is providing non-critical information. - a Debug (7) – The device is providing debug-level information.
Persistent Log Configuration	
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.
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	Select the RFC version of the syslog protocol for your syslog server. - a RFC 3164 – The old BSD syslog standard. - a RFC 5424 – The current syslog standard. It obsoletes RFC 3164.
Local UDP Port	The UDP port on the local host from which syslog messages are sent (the default port is 514).

Use the buttons to perform the following tasks:


- a 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.
- a Click **Refresh** to update the page with the most current information.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.10.2 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, either the oldest message is deleted from the RAM when a new message is added or no new messages are added anymore. If the system restarts, all messages are cleared.

3 Configuring and viewing System Information

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



Log Index	Log Time	Severity	Component	Description
1	Jan 8 00:14:10	Info	DHCP_CLI	INFO Failed to acquire an IP address on Network Port; DHCP Server did not respond.
2	Jan 8 00:09:08	Info	DHCP_CLI	INFO Failed to acquire an IP address on Network Port; DHCP Server did not respond.
3	Jan 8 00:08:59	Info	USER_MGR	INFO HTTP Session 59 started for user admin connected from 172.23.56.12
4	Jan 8 00:04:06	Info	DHCP_CLI	INFO Failed to acquire an IP address on Network Port; DHCP Server did not respond.
5	Jan 7 23:59:04	Info	DHCP_CLI	INFO Failed to acquire an IP address on Network Port; DHCP Server did not respond.
6	Jan 7 23:59:02	Info	USER_MGR	INFO HTTP Session 58 ended for user admin connected from 172.23.56.12
7	Jan 7 23:54:02	Info	DHCP_CLI	INFO Failed to acquire an IP address on Network Port; DHCP Server did not respond.
8	Jan 7 23:49:00	Info	DHCP_CLI	INFO Failed to acquire an IP address on Network Port; DHCP Server did not respond.
9	Jan 7 23:48:35	Info	USER_MGR	INFO HTTP Session 58 started for user admin connected from 172.23.56.12
10	Jan 7 23:46:42	Info	USER_MGR	INFO HTTP Session 57 ended for user admin connected from 172.23.56.12

Figure 68: Buffered Log

Table 59: 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: - a Emergency (0) – The device is unusable. - a Alert (1) – Action must be taken immediately. - a Critical (2) – The device is experiencing primary system failures. - a Error (3) – The device is experiencing non-urgent failures. - a Warning (4) – The device is experiencing conditions that could lead to system errors if no action is taken. - a Notice (5) – The device is experiencing normal but significant conditions. - a Info (6) – The device is providing non-critical information. - a Debug (7) – The device is providing debug-level information.
Component	The component that issued the log entry.
Description	The text description for the log entry.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the screen and associated messages.
- a Click **Clear log** to clear the buffered log messages and reset the counters. The buffered log will be repopulated with new entries as they occur on the system.

3.10.3 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.

Log Index	Type	Filename	Line	Task ID	Code	Event Time
1	EVENT	bootos.c	219	05FA799C	AAAAA	0d:00:00:36
2	EVENT	bootos.c	219	0766A99C	AAAAA	0d:00:00:36
3	EVENT	bootos.c	219	061D299C	AAAAA	0d:00:00:36
4	EVENT	unitmgr.c	6768	9E7FADAC	00000000	0d:00:09:04
5	EVENT	bootos.c	217	0670999C	AAAAA	0d:00:00:50
6	EVENT	bootos.c	217	0694199C	AAAAA	0d:00:00:49
7	EVENT	bootos.c	217	0631F99C	AAAAA	0d:00:00:50
8	EVENT	bootos.c	217	0724C99C	AAAAA	0d:00:00:49
9	EVENT	bootos.c	217	0738199C	AAAAA	0d:00:00:50
10	EVENT	bootos.c	217	05BBC99C	AAAAA	0d:00:00:50

Figure 69: Event Log

Table 60: 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 LCOS SX 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.

3.10.4 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 70: Logging Hosts

Table 61: Logging Hosts Fields

Field	Description
Host	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:

- a To add a logging host, click **Add** and configure the desired settings.

Figure 71: Add Host

Table 62: 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.
Port	The UDP port on the logging host to which syslog messages are sent (default port is 514).
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.

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

3.10.5 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 72: Syslog Source Interface Configuration

Table 63: Syslog Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: - a None – The primary IP address of the originating (outbound) interface is used as the source address. - a Interface – The primary IP address of a physical port is used as the source address. - a VLAN – The primary IP address of a VLAN routing interface is used as the source address. - a 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. - a Network – The network source IP is used as the source address.
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.
Loopback Interface	When the selected Type is Loopback , select the loopback interface to use as the source interface.

Use the buttons to perform the following tasks:

- a If you change any of the settings on the page, click **Submit** to apply the changes to system.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.10.6 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 73: Persistent Log

Table 64: 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: - a Emergency (0) – The device is unusable. - a Alert (1) – Action must be taken immediately. - a Critical (2) – The device is experiencing primary system failures. - a Error (3) – The device is experiencing non-urgent failures. - a Warning (4) – The device is experiencing conditions that could lead to system errors if no action is taken. - a Notice (5) – The device is experiencing normal but significant conditions. - a Info (6) – The device is providing non-critical information. - a Debug (7) – The device is providing debug-level information.
Component	The component that has issued the log entry.
Description	The text description for the log entry.

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

3.11 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 Alert menu allow you to configure information about what type of log message are sent via email and to what address(es) the messages are emailed.

3.11.1 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 74: Email Alert Global Configuration

Table 65: Email Alert Global Configuration Fields

Field	Description
Admin Mode	Sets the administrative mode of the feature. - a Disable – The device will not send email alerts. - a Enable – The device can send email alerts to the configured SMTP server.
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: - a Emergency – Indicates, that the system is unusable. It is the highest level of severity. - a Alert – Indicates, that action must be taken immediately. - a Critical – Indicates critical conditions. - a Error – Indicates error conditions. - a Warning – Indicates warning conditions. - a Notice – Indicates normal but significant conditions. - a Informational – Indicates informational messages. - a Debug – Indicates debug-level messages.
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 Messages Severity field description for information about the security levels.
Traps Severity	Configures the severity level for trap log messages. See the Urgent Messages Severity field description for information about the security levels.

Use the buttons to perform the following tasks:

- a If you make any changes to the page, click **Submit** to apply the change to the system.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.11.2 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 75: Email Alerts Server Configuration

Table 66: Email Alert Server Configuration Fields

Field	Description
Address	Shows the IPv4/IPv6 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.
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.

Use the buttons to perform the following tasks:

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

3 Configuring and viewing System Information

- a To add an SMTP server, click **Add** and configure the desired settings.

The screenshot shows a dialog box titled "Add new Email Server". It contains the following fields and descriptions:

Field	Description
Host Name or IP Address	(Max 64 characters or x.x.x.x or x.x.x.x.x.x.x.x)
Security	☒ None ☐ TLSv1
Port	(1 to 65535) - Recommend 465 for TLSv1 Mail Server Security, 25 for None
User Name	(1 to 49 alphanumeric characters)
Password	(1 to 49 characters)

Buttons: Submit, Cancel

Figure 76: New Email Server

Table 67: Add New Email Server Fields

Field	Description
Host Name or IP Address	Specify 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.
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.

- a 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.
- a To delete a configured SMTP server from the list, select the check box associated with the entry to delete and click **Remove**.
- a If you make any changes to the page, click **Submit** to apply the change to the system.

3.11.3 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.

The screenshot shows the "Email Alert Statistics" page. It has a navigation bar with tabs: Global, Test, Server, Statistics (selected), Subject, Address. Below the tabs is a table with the following data:

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

Buttons: Refresh, Clear Counters

Figure 77: Email Alert Statistics

Table 68: 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.

Use the buttons to perform the following tasks:

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

3.11.4 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 78: Email Alert Subject Configuration**Table 69: Email Alert Subject Configuration Fields**

Field	Description
Message Type	Select the appropriate option to configure the subject line of Urgent messages, Nonurgent messages or for both types.
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 .

Use the buttons to perform the following tasks:

- a If you make any changes to the page, click **Submit** to apply the change to the system.
- a To update the page with the most current information, click **Refresh**.
- a To reset a configured Email Subject to the default setting, select the ☐ Remove check box associated with the entry and click **Delete**.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.11.5 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 79: Email Alert To Address Configuration

Table 70: Email Alert To Address Configuration Fields

Field	Description
Message Type	Displays the message type (Urgent , Non Urgent or both types) to be used for sending an email to the specified address.
To Address	Displays the email address to which the selected type of messages are sent.

Use the buttons to perform the following tasks:

- a To update the page with the most current information, click **Refresh**.
- a To add an email address to the list of email alert message recipients, click **Add** and configure the desired settings.




Figure 80: Add New Email Alert to Address

Table 71: Add new Email Alert To Address Fields

Field	Description
To Address	Specify the email address to which the selected type of messages are sent.
Message Type	Specifies whether to send Urgent , Non Urgent , or both types of email alert message to the associated address.

- a To delete an entry from the list, select the check box associated with each entry to delete and click **Remove**.

3.12 Configuring Power over Ethernet

Use the following pages to configure global PoE settings. PoE allows IP telephones, wireless LAN access points, and other appliances to receive power as well as data over existing LAN cabling without modifying the existing Ethernet infrastructure. PoE is only available on switches that contain a PoE controller.

The switches support the PoE+ specification (IEEE 802.3at) for power sourcing equipment (PSE). IEEE 802.3at allows power to be supplied to Class 4 PD devices that require power greater than 15.4 Watts and up to 30.0 Watts. This allows the PoE+ enabled network switches and routers to be used for deployment with devices that require more power than the IEEE 802.3af specification allows. PoE+ IEEE 802.3at is compatible with IEEE 802.1af.

3.12.1 PoE Summary

Use the PoE Summary page to view and configure information about PoE on the device.

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

Interface	Interface Index	PoE Port Mode	PD Class	Maximum Power	Power Allocated [W]	Power Requested [W]	Power Used [W]	Current Used [mA]	Priority	Port Status	PoE Schedule	Delay Mode
1/0/1	1	Enabled	-	30	0	0	0.0	0	Low	No PD detected	disable	Disabled (0 s)
1/0/2	2	Enabled	-	30	0	0	0.0	0	Low	No PD detected	disable	Disabled (0 s)
1/0/3	3	Enabled	-	30	0	0	0.0	0	Low	No PD detected	disable	Disabled (0 s)
1/0/4	4	Enabled	-	30	0	0	0.0	0	Low	No PD detected	disable	Disabled (0 s)
1/0/5	5	Enabled	-	30	0	0	0.0	0	Low	No PD detected	disable	Disabled (0 s)
1/0/6	6	Enabled	-	30	0	0	0.0	0	Low	No PD detected	disable	Disabled (0 s)
1/0/7	7	Enabled	-	30	0	0	0.0	0	Low	No PD detected	disable	Disabled (0 s)
1/0/8	8	Enabled	-	30	0	0	0.0	0	Low	No PD detected	disable	Disabled (0 s)
1/0/9	9	Enabled	-	30	0	0	0.0	0	Low	No PD detected	disable	Disabled (0 s)
1/0/10	10	Enabled	-	30	0	0	0.0	0	Low	No PD detected	disable	Disabled (0 s)

Figure 81: PoE Summary

Table 72: PoE Summary Fields

Field	Description
Interface	Identifies the port.
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.
PoE Port Mode	The PoE Port Mode represents the PoE operating mode for the port. - a Disabled – PoE disabled for the port. - a Enabled – Enables PoE IEEE 802.3at (Class 4 PDs limited to 30W)

3 Configuring and viewing System Information

Field	Description
PD Class	Each PD is classified according to a class that defines the maximum power the PD will use. The PD Class shows the PDs class. Five Classes are defined: - a Class 0 – Max. power 15.4 W - a Class 1 – Max. power 4.0 W - a Class 2 – Max. power 7.0 W - a Class 3 – Max. power 15.4 W - a Class 4 – Max. power 30.0 W
Maximum Power	The Maximum Power value contains a numerical value that indicates the maximum power in watts that can be delivered to a remote device. The maximum allowed value is 30 W.
Power Allocated [W]	The Power Allocated shows the amount of power the switch has allocated for the PD.
Power Requested [W]	The Power Requested shows the requested amount of power the PD wants to be reserved.
Power Used [W]	The Power Used shows how much power the PD currently is using.
Current Used [mA]	The Current Used shows how much current the PD currently is using.
Priority	The Priority shows the port's priority configured by the user. The following priority levels can occur: - a Low - a High - a Critical
Port Status	The Port Status shows the port's status. The status can be one of the following values: - a PoE not available – No PoE chip found – PoE not supported for the port. - a PoE turned OFF – PoE disabled : PoE is disabled by user. - a PoE turned OFF – Power budget exceeded – The total requested or used power by the PDs exceeds the maximum power the Power Supply can deliver, and port(s) with the lowest priority is / are powered down. - a No PD detected – No PD detected for the port. - a PoE turned OFF – PD overload – The PD has requested or used more power than the port can deliver, and is powered down. - a PoE turned OFF – PD is off. - a Invalid PD – PD detected, but is not working correctly.
PoE Schedule	Scheduled by selecting PoE Scheduling Profile.
Delay Mode	Turn on / off the power delay function. - a Enabled – Enable PoE Power Delay. - a Disabled – Disable PoE Power Delay.

Use the buttons to perform the following tasks:

- a To update the page with the most current information, click **Refresh**.

- a Select your Interface and click **Edit** to modify the Port PoE Configuration.

The screenshot shows a web-based configuration window titled "Edit Port PoE Configuration". It contains several input fields and radio buttons for configuring PoE settings on a specific interface (1/0/1). The fields include PoE Port Mode (radio buttons for Disable and Enable), Priority (a dropdown menu set to Low), Maximum Power [W] (a text box with 30 and a range of 1 to 30), PoE Schedule (a dropdown menu set to Disable), Delay Mode (radio buttons for Disable and Enable), and Delay Time (a text box with 0 and a range of 0 to 300 seconds). At the bottom right, there are "Submit" and "Cancel" buttons.

Figure 82: Edit Port PoE Configuration

Table 73: Edit Port PoE Configuration Fields

Field	Description
Interface	Identifies the port.
PoE Port Mode	The PoE Mode represents the PoE operating mode for the port. a Disable – PoE disabled for the port. a Enable – Enables PoE IEEE 802.3at (Class 4 PDs limited to 30W)
Priority	The Priority shows the port's priority configured by the user. The following options are available: a Low a High a Critical
Maximum Power [W]	The Maximum Power value contains a numerical value that indicates the maximum power in watts that can be delivered to a remote device. The maximum allowed value is 30 W.
PoE Schedule	Scheduled by selecting PoE Scheduling Profile. The following profiles are available: a Disable – No PoE schedule is used. a Profile 1 - 16 – Select one of the 16 PoE Schedule Profiles .
Delay Mode	Turn on / off the power delay function. a Enable – Enable PoE Power Delay. a Disable – Disable PoE Power Delay.
Delay Time	When rebooting, the PoE port will start to provide power to the PD after the Delay Time expires. Default: 0, range: 0-300 sec.

3.12.2 PoE Configuration

Use this page to view and configure information about PoE on the device.

3 Configuring and viewing System Information

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

The screenshot shows the 'Configuration' tab of the PoE Ethernet Configuration page. It contains a table of settings with radio button options:

PoE Firmware Version	208-211
Primary Power Supply [W]	720
Reserved Power determined by	☐ Class ☒ Allocation ☐ LLDP-Med
Power Management Mode	☐ Reserved Power ☒ Actual Consumption
Capacitor Detection	☒ Disable ☐ Enable
PoE Auto Checking	☒ Disable ☐ Enable

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

Figure 83: PoE Configuration

Table 74: PoE Configuration Fields

Field	Description
PoE Firmware Version	The version of PoE MCU firmware.
Primary Power Supply [W]	For being able to determine the amount of power the PD may use, it must be defined what amount of power a power source can deliver. The available power depends on the specific switch model.
Reserved Power determined by	There are three modes for configuring how the ports / PDs may reserve power. - a Class – In this mode each port automatically determines how much power to reserve according to the class the connected PD belongs to, and reserves the power accordingly. Several different PoE classes exist (e.g. Class 3: 15.4 W or Class 4: 30 W). In this mode the Maximum Power fields have no effect. - a Allocation – In this mode the user allocates the amount of power that each port may reserve. The allocated / reserved power for each port / PD is specified in the Maximum Power fields. - a LLDP-Med – This mode is similar to the Class mode except that each port determines the amount of power it reserves by exchanging PoE information using the LLDP protocol and reserves power accordingly. If no LLDP information is available for a port, the port will reserve power using the class mode. In this mode the Maximum Power fields have no effect. For all modes: If a port uses more power than the reserved power for the port, the port is shut down.
Power Management Mode	There are 2 modes for configuring when to shut down the ports: - a Reserved Power – In this mode the ports are shut down when total reserved power exceeds the amount of power that the power supply can deliver. In this mode the port power is not turned on if the PD requests more power than available from the power supply. - a Actual Consumption – In this mode the ports are shut down when the actual power consumption for all ports exceeds the amount of power that the power supply can deliver or if the actual power consumption for a given port exceeds the reserved power for that port. The ports are shut down according to the ports priority. If two ports have the same priority the port with the highest port number is shut down.
Capacitor Detection	Enables or disables the Legacy PD Detection mode.
PoE Auto Checking	Enable Ping Check function can detect the connection between PoE port and power device. Disable will turn off the detection.

Use the buttons to perform the following tasks:

- a If you make any changes to the page, click **Submit** to apply the changes to the system.
- a To update the page with the most current information, click **Refresh**.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.12.3 PoE Auto Checking

Use the PoE Auto Checking Page to view and configure information about PoE on the device.

To access the page, click **System > PoE > PoE Auto Checking** in the navigation menu.

	Interface	Interface Index	Ping IP Address	Startup Time	Interval Time	Retry Count	Failure Log	Failure Action	Reboot Time(sec)	Max. Reboot Times
☐	1/0/1	1	0.0.0.0	60	30	3	error:0, total=0	Nothing	15	3
☐	1/0/2	2	0.0.0.0	60	30	3	error:0, total=0	Nothing	15	3
☐	1/0/3	3	0.0.0.0	60	30	3	error:0, total=0	Nothing	15	3
☐	1/0/4	4	0.0.0.0	60	30	3	error:0, total=0	Nothing	15	3
☐	1/0/5	5	0.0.0.0	60	30	3	error:0, total=0	Nothing	15	3
☐	1/0/6	6	0.0.0.0	60	30	3	error:0, total=0	Nothing	15	3
☐	1/0/7	7	0.0.0.0	60	30	3	error:0, total=0	Nothing	15	3
☐	1/0/8	8	0.0.0.0	60	30	3	error:0, total=0	Nothing	15	3
☐	1/0/9	9	0.0.0.0	60	30	3	error:0, total=0	Nothing	15	3
☐	1/0/10	10	0.0.0.0	60	30	3	error:0, total=0	Nothing	15	3

Figure 84: PoE Auto Checking

Table 75: PoE Auto Checking Fields

Field	Description
Interface	Identifies the port.
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.
Ping IP Address	The PD's IP Address the system should ping.
Startup Time	When PD has been started up, the switch will wait the time defined here to do PoE Auto Checking. Default: 60, range: 30-600 sec.
Interval Time	Device will send checking message to PD each interval time. Default: 30, range: 10-120 sec.
Retry Count	When PoE port can't ping the PD, it will retry to send detection again. After the third time, it will trigger failure action. Default: 3, range: 1-5.
Failure Log	Failure loggings counter.
Failure Action	The action when the third fail detection. - a Nothing – Keep pinging the remote PD but does nothing further. - a Reboot Remote PD – Cut off the power of the PoE port, thus rebooting the PD.
Reboot Time(sec)	When PD has been rebooted, the PoE port restores power after the specified time. Default: 15, range: 3-120 sec.

3 Configuring and viewing System Information

Field	Description
Max. Reboot Times	When Failure Action is Reboot Remote PD, it limits times of rebooting. Default: 3, range: 0-10. 0 means without rebooting limits.

Use the buttons to perform the following tasks:

- a To update the page with the most current information, click **Refresh**.
- a Select an interface and click **Edit** to modify PoE Auto Checking settings for a specific interface. If you select multiple interfaces, these settings apply to all these interfaces.

Field	Description
Interface	1/0/1
Ping IP Address	0.0.0.0 (x.x.x.x)
Startup Time	60 (Seconds) (30 to 600)
Interval Time	30 (Seconds) (10 to 120)
Retry Count	3 (1 to 5)
Failure Action	Nothing
Reboot Time(sec)	15 (Seconds) (3 to 120)
Max. Reboot Times	3 (0 to 10)

Figure 85: Edit PoE Auto Checking Configuration

Table 76: PoE Auto Checking Configuration Fields

Field	Description
Interface	Identifies the port.
Ping IP Address	The PD's IP Address the system should ping.
Startup Time	When PD has been started up, the switch will wait the time defined here to do PoE Auto Checking. Default: 60, range: 30-600 sec.
Interval Time	Device will send checking message to PD each interval time. Default: 30, range: 10-120 sec.
Retry Count	When PoE port can't ping the PD, it will retry to send detection again. When the third time fails, it will trigger failure action. Default: 3, range: 1-5.
Failure Action	The action when the detection fails after retry count has been reached. - a Nothing – Keep pinging the remote PD but does nothing further. - a Reboot Remote PD – Cut off the power of the PoE port, thus rebooting the PD.
Reboot Time(sec)	When PD has been rebooted, the PoE port restores power after the specified time. Default: 15, range: 3-120 sec.
Max. Reboot Times	When Failure Action is Reboot Remote PD, it limits times of rebooting. Default: 3, range: 0-10. 0 means without rebooting limits.

3.12.4 PoE Schedule Profile

Use the PoE Schedule Profile page to view and configure the PoE Schedule Profile.

To access the PoE Schedule Profile page, click **System > PoE > PoE Schedule Profile** in the navigation menu.

Profile	Name	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
1	Profile 1	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
2	Profile 2	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
3	Profile 3	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
4	Profile 4	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
5	Profile 5	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
6	Profile 6	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
7	Profile 7	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
8	Profile 8	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
9	Profile 9	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
10	Profile 10	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
11	Profile 11	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
12	Profile 12	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
13	Profile 13	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
14	Profile 14	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
15	Profile 15	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00
16	Profile 16	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00	00:00 - 00:00

Figure 86: PoE Schedule Profile

Table 77: PoE Schedule Profile Fields

Field	Description
Profile	The index of the Profile. There are 16 profiles in the configuration.
Name	The name of the Profile. The default name is "Profile #".
Week Day	The day to schedule PoE (from Sunday to Saturday). The Start and End Times are shown in the fields below.

Use the buttons to perform the following tasks:

- a To update the page with the most current information, click **Refresh**.
- a Select a profile and click **Edit** to modify a PoE profile. If you select multiple profiles, these settings apply to all these profiles.

Edit PoE Schedule Profile

Profile - profileStr:1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16
 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16

Name
 Profile 1 (0 to 32)

Week Day	Start Time		End Time	
	HH	MM	HH	MM
Sunday	0	0	0	0
Monday	0	0	0	0
Tuesday	0	0	0	0
Wednesday	0	0	0	0
Thursday	0	0	0	0
Friday	0	0	0	0
Saturday	0	0	0	0

Submit Cancel

Figure 87: Edit PoE Schedule Profile

Table 78: Edit PoE Schedule Profile Fields

Field	Description
Name	The name of the Profile. The default name is "Profile #". User can define the name for identifying the profile.
Week Day	The day to schedule PoE (from Sunday to Saturday).
Start Time	The time to start PoE. The time 00:00 means the first second of this day.
End Time	The time to stop PoE. The time 00:00 means the last second of this day.

3.12.5 PoE Power Information

Use the PoE Power Information page to view and configure information about switch power on the device.

To access the PoE Power Information Page, click **System > PoE > Power Information** in the navigation menu.

Power Information		
Unit	1 ▼	
Power	A	B
Status	Good	Not Present
Temperature	24.0 C	0.0 C
Fan Speed	11488	0
Operating Mode	Boost ▼	

Submit Refresh

Figure 88: PoE Power Information**Table 79: PoE Power Information Fields**

Field	Description
Unit	Select the unit to display power information.
Power	The respective power supplies (A and B).
Status	The status of the power supply unit (PSU).
Temperature	The temperature of the PSU.
Fan Speed	The fan speed of the PSU.
Operating Mode	The operating mode of the PSU. This setting is only used, when two PSUs are present. The following options are available: - a Boost – The power of both PSUs is available. Therefore significantly more PoE devices can be supplied with power compared to Redundant mode. 5 If one PSU fails and more power is used than one PSU can handle, the switch will reboot due to overload and activate only as many PoE ports the remaining PSU can handle. - a Redundant – The power of one PSU is available. When the active PSU fails, the other PSU takes over. Therefore significantly less PoE devices can be supplied with power compared to Boost mode.

Use the buttons to perform the following tasks:

- a If you make any changes to the page, click **Submit** to apply the changes to the system.
- a To update the page with the most current information, click **Refresh**.

3.13 Viewing Device Port Information

The pages in the Port menu allow you to view and monitor the physical port information for the ports available on the switch. The Port menu contains the following submenus.

3.13.1 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.

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	10h 10f 100h 100f 1000f	Enabled	Enabled	Link Down
1/0/2	2	Normal	Enabled	Auto	Unknown	10h 10f 100h 100f 1000f	Enabled	Enabled	Link Down
1/0/3	3	Normal	Enabled	Auto	Unknown	10h 10f 100h 100f 1000f	Enabled	Enabled	Link Down
1/0/4	4	Normal	Enabled	Auto	1000 Mbps Full Duplex	10h 10f 100h 100f 1000f	Enabled	Enabled	Link Up
1/0/5	5	Normal	Enabled	Auto	Unknown	10h 10f 100h 100f 1000f	Enabled	Enabled	Link Down
1/0/6	6	Normal	Enabled	Auto	Unknown	10h 10f 100h 100f 1000f	Enabled	Enabled	Link Down
1/0/7	7	Normal	Enabled	Auto	Unknown	10h 10f 100h 100f 1000f	Enabled	Enabled	Link Down
1/0/8	8	Normal	Enabled	Auto	Unknown	10h 10f 100h 100f 1000f	Enabled	Enabled	Link Down
1/0/9	9	Normal	Enabled	Auto	Unknown	10h 10f 100h 100f 1000f	Enabled	Enabled	Link Down
1/0/10	10	Normal	Enabled	Auto	Unknown	10h 10f 100h 100f 1000f	Enabled	Enabled	Link Down

Figure 89: Port Summary

Table 80: 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	The possible values are: - a Normal – The port is a normal port, which means it is not a LAG member or configured for port mirroring. - a Trunk Member – The port is a member of a LAG.

3 Configuring and viewing System Information

Field	Description
	- a 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 Mirroring on page 119. - a 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 Mirroring on page 119.
Admin Mode	Shows the port control administration state, which can be one of the following: - a Enabled – The port can participate in the network (default). - a Disabled – The port is administratively down and does not participate in the network.
Physical Mode	Shows the speed and duplex mode at which the port is configured: - a 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 is only available on copper ports. - a <Speed> Half Duplex – The port speeds available from the menu depend on the switch model 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. - a <Speed> Full Duplex – The port speeds available from the menu depend on the switch model and which port you select. In full-duplex mode, the transmissions are two-way. In other words, the port can send and receive traffic at the same time. The physical mode for a LAG is reported as <i>LAG</i>.
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: - a Enabled - Spanning tree is enabled for this port. - a Disabled - Spanning tree is disabled for this port.
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 (the options are not available for LAG ports): - a Enabled – Specifies that the port is allowed to participate in a port channel (LAG), which is the default mode. - a Disabled – Specifies that the port cannot participate in a port channel (LAG).
Link Status	Indicates whether the Link is up or down.

Use the buttons to perform the following tasks:

- a Click **Refresh** to redisplay the most current information from the router.

- a Select a port and click **Edit** to modify port settings for a specific port.

The screenshot shows the 'Edit Port Configuration' window. The 'Interface' field is set to '1/0/1'. Under 'Admin Mode', 'Enable' is selected. Under 'Physical Mode', 'Auto Negotiate' is checked, and the 'Speed' dropdown is set to '10 Mbps Half Duplex'. For 'STP Mode', 'LACP Mode', and 'Link Trap', 'Enable' is selected. The 'MTU' is set to '1500'. The 'Broadcast Storm Recovery Level', 'Multicast Storm Recovery Level', and 'Unicast Storm Recovery Level' are all set to 'Disable' with a value of '5' and 'None' selected for the unit/drop type. 'Submit' and 'Cancel' buttons are at the bottom right.

Table 81: Edit Port Configuration Fields

Field	Description
Interface	Identifies the port to be configured.
Admin Mode	Select the port control administration state, which can be one of the following: - a Enable - a Disable
Physical Mode	Shows the speed and duplex mode at which the port is configured: - a Auto - a <Speed> Half Duplex - a <Speed> Full Duplex The physical mode for a LAG is reported as <i>LAG</i> .
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: - a Enable - Spanning tree is enabled for this port. - a Disable - Spanning tree is disabled for this port.
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 (the options are not available for LAG ports): - a Enable - a Disable
Auto Negotiate	Select this option to enable auto negotiation on the port.

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. - a Enable – Specifies that the system sends a trap when the link status changes. - a Disable – Specifies that the system does not send a trap when the link status changes.
MTU	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 Level	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 Trap 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 Trap 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 Trap option sends trap messages at approximately every 30 seconds until unicast storm control recovers.

Figure 90: Edit Port Configuration

3.13.2 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.

Interface	Physical Address	PortList Bit Offset	Interface Index	Port Description
1/0/1	00:40:C7:1D:18:C6	1	1	
1/0/2	00:40:C7:1D:18:C6	2	2	
1/0/3	00:40:C7:1D:18:C6	3	3	
1/0/4	00:40:C7:1D:18:C6	4	4	
1/0/5	00:40:C7:1D:18:C6	5	5	
1/0/6	00:40:C7:1D:18:C6	6	6	
1/0/7	00:40:C7:1D:18:C6	7	7	
1/0/8	00:40:C7:1D:18:C6	8	8	
1/0/9	00:40:C7:1D:18:C6	9	9	
1/0/10	00:40:C7:1D:18:C6	10	10	

Figure 91: Port Description

Table 82: Port Description Fields

Field	Description
Interface	Identifies the port or LAG.
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.

- a Click **Refresh** to redisplay the most current information from the router.
- a Select a port and click **Edit** to modify the port description for a specific port.

Table 83: Edit Port Description Fields

Field	Description
Interface	Shows the physical or LAG port.
Port Description	Enter a port description. You can enter a description with up to 64 characters. The field accepts alpha-numeric and special characters (-, _, and space) and is also case sensitive.

Figure 92: Edit Port Description

3.13.3 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.

4 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**.




Figure 93: Port Cable Test

Table 84: Port Cable Test Fields

Field	Description
Interface	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. 4 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. 4 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. a Normal – The cable is working correctly. a Open – The cable is disconnected or there is a faulty connector. a Open and Short – There is an electrical short in the cable. a 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.

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.

4 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.

3.13.4 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 94: Multiple Port Mirroring

Table 85: 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: - a Remote VLAN – Traffic is mirrored to the VLAN on the system that is configured as the RSPAN VLAN (Remote Switch Port Analyzer). 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. - a 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. - a None – The destination is not configured. 4 This field also identifies the status of the Remove RSPAN Tag option, which can be configured in the Destination Configuration window. When this option is set as False, packets received at the RSPAN destination

Field	Description
	port are double tagged. When the Remove RSPAN Tag option is True, the RSPAN VLAN ID tag is removed for the mirroring session.
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 ports that is sent to the probe port. Possible values are: - a Tx and Rx – Both ingress and egress traffic. - a Rx – Ingress traffic only. - a Tx – Egress traffic only.

Use the buttons to perform the following tasks:


- a To configure the port mirroring destination, click the **Edit** icon in the **Destination** field and configure the desired settings.
- a Click **Refresh** to redisplay the most current information from the router.
- a To configure the administrative mode for a port mirroring session, click **Configure Session** and configure the desired settings.
- a 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.
- a 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**.

3.13.4.1 Configuring a Port Mirroring Session

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



The image shows a 'Session Configuration' dialog box. It has a title bar with a close button. Inside, there are two fields: 'Session ID' with a text input containing the number '1', and 'Mode' with two radio buttons, 'Disable' (which is selected) and 'Enable'. At the bottom right, there are 'Submit' and 'Cancel' buttons.

2. Configure the following fields:

Table 86: Multiple Port Mirroring—Session Configuration

Field	Description
Session ID	The port mirroring session ID. Four sessions are available for configuration.
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.

- Click **Submit** to apply the changes to the system.

3.13.4.2 Configuring a Port Mirroring Source

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

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

- Configure the following fields:

Table 87: Multiple Port Mirroring-Source Configuration

Field	Description
Session ID	The port mirroring session ID. Four sessions are available for configuration.
Type	The type of interface to use as the source: - a None – The source is not configured. - a 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. - a 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. - a Interface – Traffic is mirrored from one or more physical ports on the device.
Remote VLAN	The VLAN that is configured as the RSPAN VLAN.

Field	Description
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: - a Tx/Rx – Both ingress and egress traffic. - a Rx – Ingress traffic only. - a Tx – Egress traffic only.

- Click **Submit** to apply the changes to the system.

3.13.4.3 Configuring the Destination Port for a Port Mirroring Session

4

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

- 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**.

- Configure the following fields:

Table 88: Multiple Port Mirroring-Destination Configuration

Field	Description
Session ID	The port mirroring session ID. Four sessions are available for configuration.
Type	The type of interface to use as the destination: - a None – The destination is not configured. - a 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. - a 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.
Remote VLAN	The VLAN that is configured as the RSPAN VLAN.

Field	Description
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.

3.13.4.4 Removing or Modifying a Port Mirroring Session

1. Select one or more source ports to remove from the session.
2. Click **Remove Source**.

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

3.13.5 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.

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	Disabled										
2	Disabled										
3	Disabled										
4	Disabled										

Figure 95: Multiple Port Mirroring Summary

Table 89: Multiple Port Mirroring Summary Fields

Field	Description
Session ID	The port mirroring session ID. Four sessions are available for configuration.
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 traffic to the destination. You can configure multiple source ports per session.

3 Configuring and viewing System Information

Field	Description
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: - a Tx/Rx – Both ingress and egress traffic. - a Rx – Ingress traffic only. - a Tx – Egress traffic only.
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.

3.13.6 SFP Information

Use the SFP Information Page to view the SFP module detail information.

To access the Port SFP Information Page, click **System > Port > SFP Information** in the navigation menu.

Field	Description
Interface	Select the port to display sfp information.
Connector Type	Displays the connector type, for instance, UTP, SC, ST, LC and so on.
Fiber Type	Displays the fiber mode, for instance, Multi-Mode, Single-Mode.
Vendor OUI	Displays the Manufacturer's OUI code which is assigned by IEEE.

Figure 96: Port SFP Information

Table 90: Port SFP Information Fields

Field	Description
Interface	Select the port to display sfp information.
Connector Type	Displays the connector type, for instance, UTP, SC, ST, LC and so on.
Fiber Type	Displays the fiber mode, for instance, Multi-Mode, Single-Mode.
Vendor OUI	Displays the Manufacturer's OUI code which is assigned by IEEE.

Field	Description
Vendor Name	Displays the company name of the module manufacturer.
Vendor P/N	Displays the product name and the name of the module manufacturer.
Vendor Revision	Displays the module revision.
Vendor Serial Number	Shows the serial number assigned by the manufacturer.
Date Code	Shows the date when this SFP module was manufactured.
Temperature	Shows the current temperature of the SFP module.
Vcc	Shows the working DC voltage of the SFP module.
Mon1 (Bias)	Shows the Bias current of the SFP module.
Mon2 (TX PWR)	Shows the transmit power of the SFP module.

3.14 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.

3.14.1 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.

3 Configuring and viewing System Information

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




Figure 97: sFlow Agent Summary

Table 91: 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: a MIB Version: '1.3', the version of this MIB. a Organization: Broadcom Corp. a Software Revision: Firmware version of the switch
Agent Address	The IP address associated with this agent (the switch).

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

3.14.2 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 98: sFlow Receiver Configuration

Table 92: 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 entity 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.
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 12188.
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 buttons to perform the following tasks:

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

Field	Value	Range/Constraint
Index	1	
Owner String		(Max 127 characters)
Timeout Mode	☒	
Timeout Value (Seconds)	0	(0 to 2147483647)
Maximum Datagram Size	1400	(200 to 12188)
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	

Figure 99: Edit Receiver Configuration

- a Select an sFlow Receiver and click **Clear** to reset its settings to default setting.

3.14.3 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.

3.14.3.1 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.

Agent Receiver **Poller** Sampler Source Interface Configuration

sFlow Poller Configuration

There is no Receiver Index configured.

Display **All** rows Showing 0 to 0 of 0 entries Filter:

☐	Poller Data Source	Receiver Index	Poller Interval
Table is Empty			

First Previous Next Last

Refresh Add Edit Remove

Figure 100: sFlow Poller Configuration

Use the buttons to perform the following tasks:

- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a To add an sFlow poller instance, click **Add** and complete the required information.

Add Poller

Poller Data Source 1/0/1

Receiver Index 1

Poller Interval (Seconds) 0 (0 to 86400)

Submit Cancel

Figure 101: Add Poller

Table 93: Add Poller Fields

Field	Description
Poller DataSource	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

- a 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.
- a To delete an sFlow poller instance, select one or more table entries and click **Remove**.

3.14.4 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.

3.14.4.1 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 Flows to which it belongs.

Packet Flow Sampling is accomplished as follows:

- a When a packet arrives on an interface, the Network Device makes a filtering decision to determine whether the packet should be dropped.
- a If the packet is not filtered (dropped), a destination interface is assigned by the switching/routing function.
- a 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.

3 Configuring and viewing System Information

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

Figure 102: sFlow Sampler Configuration

Table 94: 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.

Use the buttons to perform the following tasks:

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

- a To add an sFlow sampler instance, click **Add** and complete the required information.

Figure 103: Add Sampler

Table 95: 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.
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. The allowed range is 20 to 256.
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.

- a 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.

3.14.5 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 104: sFlow Source Interface Configuration

Table 96: sFlow Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: - a None – The primary IP address of the originating (outbound) interface is used as the source address. - a Interface – The primary IP address of a physical port is used as the source address. - a VLAN – The primary IP address of a VLAN routing interface is used as the source address. - a 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. - a Network – The network source IP is used as the source address.
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.
Loopback Interface	When the selected Type is Loopback , select the loopback interface to use as the source interface.

Use the buttons to perform the following tasks:

- a 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**.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.15 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.

3.15.1 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.

5 SNMPv1 and SNMPv2 are not encrypted. LANCOM Systems therefore recommends using SNMPv3 as it features encryption capabilities.

3.15.2 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:

- a **Authentication:** Provides data integrity and data origin authentication.
- a **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.
- a **Timeliness:** Protects against message delay or message redundancy. The SNMP agent compares incoming message to the message time information.
- a **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:

- a Feature Access Control
- a Security
- a 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.

3.15.3 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.

4 In LCOS SX, no SNMP communities exist by default.

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.




Table 97: SNMP 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: - a public – This SNMP community has Read Only privileges and its status set to enable. - a private – This SNMP community has Read/Write privileges and its status set to enable.
Security Name	Identifies the Security entry that associates Communities and Groups for a specific access type.
Group Name	Identifies the Group associated with this Community entry.
IP Address	Specifies the IP address that can connect with this community. If the field is left empty, access is possible from any IP address.

Figure 105: SNMP Community Configuration

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To add a community, click **Add Community** and configure the desired settings.

Figure 106: Add New Community

Table 98: Add New Community 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: - a public – This SNMP community has Read Only privileges and its status set to enable. - a private – This SNMP community has Read/Write privileges and its status set to enable.
Community Access	Specifies the access control policy for the community. The options (DefaultRead , DefaultWrite and DefaultSuper) are specified in the menu Access Control Group .
Community View	Specifies the community view for the community. If the value is empty, then no access is granted.
IP Address	Specifies the IP address that can connect with this community.

- a To add a community group, click **Add Community Group** and configure the desired settings.

Figure 107: Add New Community Group**Table 99: Add New Community Configuration Group 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: - a public: This SNMP community has Read Only privileges and its status set to enable. - a private: This SNMP community has Read/Write privileges and its status set to enable.
Group Name	Identifies the Group associated with this Community entry.
IP Address	Specifies the IP address that can connect with this community.

- a Click **Remove** to delete the selected SNMP Community or Community group.
- a If you make any changes to the page, click **Submit** to apply the changes to the system.

3.15.4 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 108: SNMP v1/v2 Trap Receivers

Table 100: 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: - a 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. - a 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.
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.

Use the buttons to perform the following tasks:

- a To add an SNMP trap receiver and configure its settings, click **Add** and complete the required information.

The screenshot shows a configuration window titled "Add SNMP v1/v2 Host". It contains the following fields and values:

Field	Value	Range/Format
Host IP Address	[Empty]	(x.x.x.x or x.x.x.x.x.x.x.x)
Community Name	[Empty]	(1 to 20 characters)
Notify Type	☒ Trap ☐ Inform	
SNMP Version	☒ SNMP V1 ☐ SNMP V2	
Retries	3	(1 to 255)
Timeout Value (Seconds)	15	(1 to 300)
Filter	[Empty]	(0 to 30 characters)
UDP Port	162	(1 to 65535)

Buttons: Submit, Cancel

Figure 109: Add SNMP v1/v2 Host

Table 101: 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: - a Trap - a Inform
SNMP Version	The version of SNMP to use, which is either SNMPv1 or SNMPv2.
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.

- a To delete one or more SNMP trap receivers from the list, select each entry to delete and click **Remove**.
- a 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**.

3.15.5 Trap Receiver v3 Configuration

Use the Trap Receiver v3 Configuration 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 110: SNMP v3 Trap Receivers

Table 102: 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: - a 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. - a 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.
Security Level	The security level associated with the SNMP user, which is one of the following: - a No Auth No Priv – No authentication and no data encryption (no security). - a 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. - a 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.
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.

Field	Description
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.

Use the buttons to perform the following tasks:

- a To add an SNMP trap receiver and configure its settings, click **Add** and complete the required information.

Figure 111: Add SNMP v3 Host

- a **Table 103: Add SNMP v3 Host**

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: - a Trap - a Inform
Security Level	The security level associated with the SNMP user, which is one of the following: - a No Auth No Priv - a Auth No Priv - a Auth Priv
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.

3 Configuring and viewing System Information

Field	Description
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.

- a To delete one or more SNMP trap receivers from the list, select each entry to delete and click **Remove**.
- a 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**.

3.15.6 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.

Community	Trap Receiver v1/v2	Trap Receiver v3	Supported MIBs	Access Control Group	User Security Model	View Entry	Source Interface Configuration

SNMP Supported MIBs

Display rows Showing 1 to 10 of 52 entries Filter:

Name	Description
DIFFSERV-DSCP-TC	The Textual Conventions defined in this module should be used whenever a Differentiated Services Code Point is used in a MIB.
DISMAN-TRACEROUTE-MIB	The Traceroute MIB (DISMAN-TRACEROUTE-MIB) provides access to the traceroute capability at a remote host.
DNS-SERVER-MIB	The MIB module for entities implementing the server side of the Domain Name System (DNS) protocol.
FASTPATH-AUTHENTICATION-MANAGER-MIB	The Broadcom Private MIB for FASTPATH authentication manager feature.
FASTPATH-BOXSERVICES-PRIVATE-MIB	The Broadcom Private MIB for FASTPATH Box Services Feature.
FASTPATH-DHCP6SERVER-PRIVATE-MIB	The Broadcom Private MIB for FASTPATH DHCPv6 Server/Relay
FASTPATH-DHCPSEVER-PRIVATE-MIB	The Broadcom Private MIB for FASTPATH DHCP Server
FASTPATH-DOT1X-ADVANCED-FEATURES-MIB	The Broadcom Private MIB for FASTPATH Dot1x Advanced Features
FASTPATH-INTERFACE-APP-MIB	The Broadcom Private MIB for FASTPATH Interface Application
FASTPATH-IPV6-LOOPBACK-MIB	The Broadcom Private MIB for FASTPATH Loopback IPV6 address configuration.

First Previous 1 2 3 4 5 Next Last

Refresh

Figure 112: SNMP Supported MIBs

Table 104: 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.

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

3.15.7 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.

Group Name	Context Name	SNMP Version	Security Level	Read	Write	Notify
DefaultRead		SNMP V1	No Auth No Priv	Default		Default
DefaultRead		SNMP V2	No Auth No Priv	Default		Default
DefaultRead		SNMP V3	No Auth No Priv	Default		Default
DefaultRead		SNMP V3	Auth No Priv	Default		Default
DefaultRead		SNMP V3	Auth Priv	Default		Default
DefaultSuper		SNMP V1	No Auth No Priv	DefaultSuper	DefaultSuper	DefaultSuper
DefaultSuper		SNMP V2	No Auth No Priv	DefaultSuper	DefaultSuper	DefaultSuper
DefaultSuper		SNMP V3	No Auth No Priv	DefaultSuper	DefaultSuper	DefaultSuper
DefaultWrite		SNMP V1	No Auth No Priv	Default	Default	Default
DefaultWrite		SNMP V2	No Auth No Priv	Default	Default	Default

Figure 113: SNMP Access Control Group

Table 105: 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. It can be one of the following versions: - a SNMP V1 - a SNMP V2 - a SNMP V3
Security Level	The security level associated with the group, which is one of the following: - a No Auth No Priv – No authentication and no data encryption (no security). This is the only Security Level available for SNMPv1 and SNMPv2 groups. - a 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. This security level is only available for SNMPv3. - a 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. This security level is only available for SNMPv3.
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

3 Configuring and viewing System Information

Field	Description
	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.

Use the buttons to perform the following tasks:

- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a To add an SNMP group, click **Add** and specify the desired setting.

Figure 114: Add new Access Control Group

- a **Table 106: Add new Access Control Group**

Field	Description
Group Name	The name that identifies the SNMP group.
SNMP Version	The SNMP version associated with the group. The following options are available: - a SNMP V1 - a SNMP V2 - a SNMP V3
Security Level	The security level associated with the group, which is one of the following: - a No Auth No Priv - a Auth No Priv - a Auth Priv
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

Field	Description
	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.
Group Access Rights	a Read a Write a Notify

- a 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.

5

The default entries cannot be selected and deleted.

- a 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**.

3.15.8 SNMP User Security Model

The SNMP User Security Model page provides the capability to configure the SNMPv3 user accounts.

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

Figure 115: SNMP User Security Model

Table 107: Add New SNMP User Fields

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	An 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.
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 an hexadecimal string in the form 0102030405.

3 Configuring and viewing System Information

Field	Description
Authentication	Specifies the authentication protocol to be used on authenticated messages on behalf of the specified user. - a None - No authentication will be used for this user. - a MD5 - MD5 protocol will be used. - a SHA - SHA protocol will be used. - a SHA-224 - SHA-224 protocol will be used. - a SHA-256 - SHA-256 protocol will be used. - a SHA-384 - SHA-384 protocol will be used. - a SHA-512 - SHA-512 protocol will be used.
Privacy	Specifies the privacy protocol to be used on encrypted messages on behalf of the specified user. This parameter is only available if the Authentication method parameter is not None. - a None - No privacy protocol will be used. - a DES - DES protocol will be used. - a AES - AES protocol will be used. - a AES-192 - AES-192 protocol will be used. - a AES-256 - AES-256 protocol will be used.

Use the buttons to perform the following tasks:

- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a To add a user, click **Add**. The Add New SNMP User dialog box opens. Specify the new account information in the available fields.

Figure 116: Add New SNMP User

Table 108: 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 an hexadecimal string in the form 0102030405.

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	An 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. - a None - No authentication will be used for this user. - a MD5 - MD5 protocol will be used. - a SHA - SHA protocol will be used. - a SHA-224 - SHA-224 protocol will be used. - a SHA-256 - SHA-256 protocol will be used. - a SHA-384 - SHA-384 protocol will be used. - a SHA-512 - SHA-512 protocol will be used.
Password	Specifies the password used to generate the key to be used in authenticating messages on behalf of this user. The password must have between 1 - 32 characters. 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 available if the Authentication method parameter is not None . - a None - No privacy protocol will be used. - a DES - DES protocol will be used. - a AES - AES protocol will be used. - a AES-192 - AES-192 protocol will be used. - a AES-256 - AES-256 protocol will be used.
Authentication Key	Specifies the password used to generate the key to be used in encrypting messages to and from this user. The password must have between 1 - 32 characters. This parameter must be specified if the Privacy parameter is not None .

- a Select an SNMP user and click **Remove** to delete it.
- a 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**.

3.15.9 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.

3 Configuring and viewing System Information

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 117: SNMP Source Interface Configuration

Table 109: SNMP Trap Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: - a None – The primary IP address of the originating (outbound) interface is used as the source address. - a Interface – The primary IP address of a physical port is used as the source address. - a VLAN – The primary IP address of a VLAN routing interface is used as the source address. - a 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. - a Network – The network source IP is used as the source address.
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.
Loopback Interface	When the selected Type is Loopback , select the loopback interface to use as the source interface.

Use the buttons to perform the following tasks:

- a 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**
- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.16 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.

3.16.1 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.

Statistics	Transmit	Receive
Octets Without Error	388515189	20750846
Packets Without Errors	284774	110833
Packets Discarded	0	0
Unicast Packets	269153	93299
Multicast Packets	8518	12809
Broadcast Packets	7103	4725

Status	FDB Entries	VLANs
Current Usage	5	1
Peak Usage	41	1
Maximum Allowed	32768	4093
Static Entries	0	1
Dynamic Entries	5	0
Total Entries Deleted	N/A	0

System	Interface
Time Since Counters Last Cleared	465 7d:19:25:23

Refresh Clear Counters

Figure 118: Switch Statistics

Table 110: 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 being delivered 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. 4 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. 4 This number does not include multicast packets.
Status	

3 Configuring and viewing System Information

Field	Description
Current Usage	In the FDB Entries column (Forwarding Database), 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.

Use the buttons to perform the following tasks:

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

3.16.2 Port Summary Statistics

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.

Note: All entries in this table indicate packet counts.

Display 10 rows Showing 1 to 10 of 118 entries Filter:

☐	Interface	Rx Good	Rx Errors	Rx Bcast	Tx Good	Tx Errors	Tx Collisions
☐	1/0/1	2299	0	93	6837	0	0
☐	1/0/2	0	0	0	0	0	0
☐	1/0/3	0	0	0	0	0	0
☐	1/0/4	97316	0	2731	340937	0	0
☐	1/0/5	0	0	0	0	0	0
☐	1/0/6	0	0	0	0	0	0
☐	1/0/7	0	0	0	0	0	0
☐	1/0/8	0	0	0	0	0	0
☐	1/0/9	0	0	0	0	0	0
☐	1/0/10	0	0	0	0	0	0

First Previous 1 2 3 4 5 Next Last

Refresh Clear Counters Clear All Counters

Figure 119: Port Summary Statistics

Table 111: 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 delivered to a higher-layer protocol.
Rx Bcast	The total number of good packets received that were directed to the broadcast address. 4 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.
- a 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.
- a Click **Clear All Counters** to clear counters for all switches in the stack.

3.16.3 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 120: Port Detailed Statistics](#) on page 150 shows some, but not all, of the fields on the Port Detailed page.

Switch

Port Summary

Port Detailed

DHCPv6 Client

Port Detailed Statistics

?

Interface	1/0/1 ▾	
Maximum Frame Size	1518	
MTU	1500	

Packet Lengths Received and Transmitted		
64 Octets	2111	
65-127 Octets	2679	
128-255 Octets	256	
256-511 Octets	451	
512-1023 Octets	638	
1024-1518 Octets	3023	
1519-1522 Octets		
1523-2047 Octets	0	
2048-4095 Octets	0	
4096-9216 Octets	0	

Basic	Transmit	Receive
Unicast Packets	4140	1884
Multicast Packets	2477	322
Broadcast Packets	242	93
Total Packets (Octets)	5117588	523714
Packets > 1518 Octets	0	0
802.3x Pause Frames	0	0
FCS Errors	0	0

Protocol	Transmit	Receive
STP BPDUs	0	0
RSTP BPDUs	0	0
MSTP BPDUs	2208	0
SSTP BPDUs	0	0
GVRP PDUs	0	0
GMRP PDUs	0	0

Figure 120: Port Detailed Statistics

Table 112: Port Detailed Statistics Fields

Field	Description
Interface	Use the 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.
MTU	Indicates MTU (Maximum Transmit Unit) of the interface. The actual frame size is calculated by adding Ethernet header size in MTU.
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).

3 Configuring and viewing System Information

Field	Description
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).
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.

3 Configuring and viewing System Information

Field	Description
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.
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

Field	Description
	Sequence (FCS) with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).
	4 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the data on the screen and display the most current statistics.
- a Click **Clear Counters** to clear all the counters. This resets all statistics for this port to the default values.
- a 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.

3.16.4 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.

3 Configuring and viewing System Information

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




Figure 121: Port DHCPv6 Client Statistics

Table 113: 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.
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.

Field	Description
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the data on the screen and display the most current statistics.
- a Click **Clear Counters** to clear all the statistics counters, resetting all switch summary and detailed statistics to default values.

3.16.5 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 122: Time Based Group Statistics

Table 114: Time Based Group Statistics Fields

Field	Description
Group ID	The Group ID is generated automatically on the basis of the used Group. The ID can be from 1 (Received) to 7 (Congestion).
Group	The type of traffic statistics to collect for the group, which is one of the following: - a Received – The number of packets received on the interfaces within the group. - a Received Errors – The number of packets received with errors on the interfaces within the group. - a Transmitted – The number of packets transmitted by the interfaces within the group. - a Transmitted Errors – The number of packets transmitted with errors by the interfaces within the group.

3 Configuring and viewing System Information

Field	Description
	- a Received Transmitted – The number of packets received and transmitted by the interfaces within the group. - a Port Utilization – The percentage of total bandwidth used by the port within the specified time period. - a Congestion – The percentage of time within the specified time range that the ports experienced congestion.
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: - a 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. - a Console – The statistics are displayed on the console. - a 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. - a Syslog – The statistics are sent to a remote syslog server. The syslog server information is configured on the Logging Hosts page.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the data on the screen with the present state of the data in the switch.
- a To add a set of time-based traffic group statistics to collect, click **Add** and configure the desired settings.

Time Based Group Configuration

Group: Received

Time Range: (1 to 31 alphanumeric characters)

Reporting Methods: None, Console, E-Mail, Syslog

Interfaces: 1/0/1, 1/0/2, 1/0/3, 1/0/4, 1/0/5, 1/0/6, 1/0/7, 1/0/8

Submit Cancel

Figure 123: Time Based Group Configuration

a **Table 115: Time Based Group Configuration Fields**

Field	Description
Group	The type of traffic statistics to collect for the group, which is one of the following: - a Received - a Received Errors - a Transmitted - a Transmitted Errors - a Received Transmitted - a Port Utilization - a Congestion
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: - a None - a Console - a E-Mail - a Syslog
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.

- a To delete one or more time-based statistics groups, select each entry to delete and click **Remove**.

3.16.6 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 SNTP) before using the time-based statistics feature.

3 Configuring and viewing System Information

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




Figure 124: Time Based Flow Statistics

Table 116: Time Based Flow Statistics Fields

Field	Description
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.

Use the buttons to perform the following tasks:

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

- a 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. The match conditions are optional, but the rule must specify at least one match condition.

Time Based Flow Configuration

Rule Id: (1 to 16)

Time Range: (1 to 31 alphanumeric characters)

Interface: 1/0/1, 1/0/2, 1/0/3, 1/0/4, 1/0/5, 1/0/6, 1/0/7, 1/0/8

Match Criteria

Match All: ☐

Source IP: (x.x.x.x)

Destination IP: (x.x.x.x)

Source MAC: (xx:xx:xx:xx:xx:xx)

Destination MAC: (xx:xx:xx:xx:xx:xx)

Source TCP Port: (1 to 65535)

Destination TCP Port: (1 to 65535)

Source UDP Port: (1 to 65535)

Destination UDP Port: (1 to 65535)

Submit Cancel

Figure 125: Time Based Flow Configuration

Table 117: Time Based Flow Configuration Fields

Field	Description
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.
Interface	The interface or interfaces the flow-based rule is applied to. Only traffic on the specified interfaces is checked against the rule.
Match Criteria	
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.

3 Configuring and viewing System Information

Field	Description
Destination TCP Port	The TCP destination port to match in the TCP header.
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.

- a To delete one or more flow-based rules for time-based statistics, select each entry to delete and click **Remove**.

3.16.7 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 126: Time Based Statistics

Table 118: 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.

3.17 Using System Utilities

The System Utilities feature menu contains links to menus that help you manage the switch.

3.17.1 System Reboot

Use the System Reboot page to reboot the system. If the platform supports stacking, you can reboot 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 127: System Reboot

Table 119: 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.

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

Click **Reboot** to initiate the system reset.

- Any configuration changes made since the last successful save are lost whenever a switch is rebooted. 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 webinterface.

3.17.2 Ping

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

3 Configuring and viewing System Information

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

Figure 128: Ping

Table 120: Ping Fields

Field	Description
Host Name or IP Address	Enter the IP address or the host name of the station you want the switch to ping. The initial value is blank.
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. The following options are available: - a None – No Source is specified (default setting). - a IP Address – Enter a specific IP Address configured on the switch to send the ICMP requests. - a Interface – Select a specific Interface to send the ICMP requests.
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.

Use the buttons to perform the following tasks:

- a Click the **Start** button to start the ping test. The device sends the specified number of ping packets to the host.
- a Click the **Stop** button to interrupt the current ping test.

3.17.3 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 129: Ping IPv6

Table 121: 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 is displayed 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 Link Local, 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.

3 Configuring and viewing System Information

Field	Description
Interface	The interface to use when sending the Echo requests packets. This field is enabled when Interface 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.

3.17.4 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 130: TraceRoute

Table 122: 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.

Field	Description
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	The source IP address or interface to use when sending the echo request packets. The following options are available: a None – No Source is specified (default setting). a IP Address – Enter a specific IP Address configured on the switch to send the TraceRoute. a Interface – Select a specific Interface to send the TraceRoute. a Loopback – Select a specific loopback interface to send the TraceRoute.
IP Address	The source IPv4 address to use when sending the TraceRoute command. This field is enabled when IP Address is selected as the source option.
Interface	The interface to use when sending the TraceRoute command. This field is enabled when Interface is selected as the source option.
Loopback	The loopback interface to use when sending the TraceRoute command. This field is enabled when Loopback is selected as the source option.
Status	The current status of the TraceRoute, which can be: a Not Started – The TraceRoute has not been initiated since viewing the page. a In Progress – The TraceRoute has been initiated and is running. a Stopped – The TraceRoute was interrupted by clicking the Stop button. a Done – The TraceRoute has completed, and information about the TraceRoute is displayed in the Results area.
Results	The results of the TraceRoute are displayed

Use the buttons to perform the following tasks:

- a Click the **Start** button to start the TraceRoute.
- a Click the **Stop** button to interrupt the TraceRoute.

3.17.5 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.

3 Configuring and viewing System Information

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

Figure 131: IP Address Conflict Detection

Table 123: Multiple Port Mirroring-Session Configuration

Field	Description
Status	
IP Address Conflict Currently Exists	Indicates whether a conflicting IP address has been detected since this status was last reset. - a False – No conflict detected (in this case the subsequent fields on this page are blank). - a True – Conflict was detected (in this case the subsequent fields on this page show the relevant information).
History	
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.

Use the buttons to perform the following tasks:

- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a Click the **Run Detection** button to activate the IP address conflict detection operation in the system.
- a Click the **Clear History** button to reset the IP address conflict detection status information that was last seen by the device.

3.17.6 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.

5

The **File Transfer** menu is designed from the viewpoint of the switch. From the viewpoint of an external target (e.g. a computer) it's the opposite. Therefore an **Upload** lets you save a file to an external target (eg. your computer) and the **Download** lets you upload a file from an external source to the switch.

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



Figure 132: File Transfer

Table 124: File Transfer Fields

Field	Description
Transfer Protocol	The protocol to use to transfer the file. The switch can upload files to and download files from an external target using one of the following protocols: - a HTTP - a TFTP - a FTP - a SCP - a SFTP
Upload	To transfer a file from the device to a remote system, 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 Begin Transfer button to begin the transfer.
Download	To transfer a file from a remote system to the device, 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 Begin Transfer button to begin the transfer.

Upload:

3 Configuring and viewing System Information

After you click the upload icon, the **File Upload** window appears (in this case for the protocol **SCP**).

The image shows a window titled "SCP File Upload". It contains a table with the following fields:

File Type	Active Code	
Server Address		(Host Name or IP Address)
File Path		(0 to 160 characters)
File Name		(1 to 33 characters)
User Name		(1 to 32 characters)
Password		(0 to 64 characters)
Status		

Below the table is a button labeled "Begin Transfer" with a right-pointing arrow. In the bottom right corner of the window is a "Close" button.

Figure 133: SCP File Upload

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

Table 125: SCP File Upload Fields

Field	Description
File Type	Specify the type of file to transfer from the device to a remote system. - a Active Code – Select this option to transfer an active image. - a Backup Code – Select this option to transfer a backup image. - a Startup Configuration – Select this option to transfer a copy of the stored startup configuration from the device to a remote system. - a Backup Configuration – Select this option to transfer a copy of the stored backup configuration from the device to a remote system. - a Script File – Select this option to transfer a custom text configuration script from the device to a remote system. - a 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. - a Crash Log – Select this option to transfer the system crash log to a remote system. - a Operational Log – Select this option to transfer the system operational log to a remote system. - a Startup Log – Select this option to transfer the system startup log to a remote system. - a Trap Log – Select this option to transfer the system trap records to a remote system. - a Error Log – Select this option to transfer the system error (persistent) log, which is also known as the event log, to a remote system. - a Buffered Log – Select this option to transfer the system buffered (in-memory) log to a remote system. - a Technical Support Log – Select this option to transfer the technical support log to a remote system.

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.
Status	Provides information about the status of the file transfer.

Download:

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

Figure 134: SFTP File Download

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

Table 126: SFTP File Download Fields

Field	Description
File Type	Specify the type of file to transfer to the device: - a Backup Code – Select this option to transfer a new image to the device. The code file is stored as the backup image. - a Startup Configuration – Select this option to update the stored startup configuration file. If the file has errors, the update will be stopped. - a Backup Configuration – Select this option to update the stored backup configuration file. If the file has errors, the update will be stopped. - a Script File – Select this option to transfer a text-based configuration script to the device. You can validate and activate the script either via the CLI (Command Line Interface) command reload configuration or via the webinterface by clicking Reload in the menu System Configuration Storage Reload.

3 Configuring and viewing System Information

Field	Description
	- a 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. - a 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. - a 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. 4 The SSH1-RSA can be downloaded, but they cannot be used. - a 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. - a 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. - a 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. - a 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. - a 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. - a 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. - a SSL Server Certificate PEM File – Select this option to transfer an SSL Server Certificate file (PEM Encoded) to the device. - a 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. - a 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. - a Public Key Image – Select this option to transfer the public key file used for code image validation to the device. - a Public Key Config – Select this option to transfer the public key file used for configuration script validation to the device. 4 - a To download SSH key files, SSH must be administratively disabled, and there can be no active SSH sessions. - a To download SSL related files, HTTPS must be administratively disabled.
Certificate Index	Index used to name a related group of certificate (PEM) or key files.
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.

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.
Status	Provides information about the status of the file transfer.

3.17.7 Core Dump

A Core Dump can be saved when *rebooting the switch*. 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.

- 5** The protocols **TFTP** and **FTP** are only available on switches with a service port. Therefore the available menu items can vary depending on the switch model used.

Figure 135: Core Dump

Table 127: Core Dump Fields

Field	Description
Core Dump Configuration	
Protocol	The protocol used to store the core dump file. User can select: a None – Disable Core Dump. a TFTP – Configure protocol to upload Core Dump to the TFTP server. a USB – Configure protocol to upload Core Dump to the USB mount point. a Local – Configure protocol to generate Core Dump on switch local file system. a FTP – Configure protocol to upload Core Dump to the FTP server.
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	User name of remote FTP server.
FTP Password	Password of remote FTP server.
File Path	File path to dump core file to FTP/TFTP server 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.
Core Dump Stack IP Address Pool	
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

Use the buttons to perform the following tasks:

- a 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**
- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a To add a stack IP address, click **Add** and configure an IP address, netmask, and gateway address.
- a To delete a configured stack IP, select each entry to delete, click **Remove**, and confirm the action.

3.17.8 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 136: Core Dump Test

Table 128: Core Dump Test Fields

Field	Description
Status	Displays test status as <code>OK</code> if test passes and <code>Error</code> if test fails.
Result	Displays detailed error information with logs.

Click Submit to initiate the Core Dump Test.

3.18 Managing SNMP Traps

The pages in the Trap Manager menu allow you to view and configure information about SNMP traps the system generates.

3.18.1 System Trap Log

Use the System Trap Log page to view the entries in the trap log.

3 Configuring and viewing System Information

To access the System Trap Log page, click **System > Advanced Configuration > Trap Manager > Trap Log** in the navigation menu.

Log	System Up Time	Trap
0	Jan 8 18:15:53 1970	HTTP session : Login to the switch is not successful, User ID: admin Source IP: Local port: 80
1	Jan 8 18:15:50 1970	HTTP session : Login to the switch is not successful, User ID: admin Source IP: Local port: 80
2	Jan 8 18:15:03 1970	Link Up: 1/0/1
3	Jan 8 18:09:11 1970	Link Down: 1/0/11
4	Jan 8 18:06:00 1970	Link Up: 1/0/4
5	Jan 8 18:05:57 1970	Link Down: 1/0/4
6	Jan 8 18:05:39 1970	Link Up: 1/0/4
7	Jan 8 17:53:08 1970	Link Up: 1/0/18
8	Jan 8 17:50:33 1970	Link Up: 1/0/11
9	Jan 8 17:40:12 1970	Link Down: 1/0/11

Figure 137: System Trap Log

Table 129: 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.

Use the buttons to perform the following tasks:

- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a Click **Clear Log** to clear all entries in the log. Subsequent displays of the log will only show new log entries.

3.18.2 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**.

System Trap Flags	
Authentication	☒
Link Up/Down	☒
Multiple Users	☒
Spanning Tree	☒
ACL Traps	☐
Fan	☒
Power Supply Module State	☒
Temperature	☒

Submit Refresh Cancel

Figure 138: System Trap Flags Configuration

Table 130: System Trap Flags Fields

Field	Description
Authentication	Specify whether to enable SNMP notifications when events involving authentication occur, such as when a user attempts to access the device management interface and fails to provide a valid username and password. The factory default is enabled.
Link Up/Down	Specify whether to enable SNMP notifications when the administrative or operational state of a physical or logical link changes. The factory default is enabled.
Multiple Users	Specify whether to enable SNMP notifications when the same user ID is logged into the device more than once at the same time (either via telnet or the serial port). The factory default is enabled.
Spanning Tree	Specify whether to enable SNMP notifications when various spanning tree events occur. The factory default is enabled.
ACL Traps	Specify whether to enable SNMP notifications when a packet matches a configured ACL rule that includes ACL logging. The factory default is disabled.
Fan	Specify whether to enable SNMP notifications when fan events occur.
Power Supply Module State	Specify whether to enable SNMP notifications when power supply events occur.
Temperature	Specify whether to enable SNMP notifications when temperature events occur.

Use the buttons to perform the following tasks:

- a 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**
- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.19 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.

3.19.1 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 139: Time Range Summary

Table 131: Time Range Summary Fields

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:

- a If you make any changes to the page, click **Submit** to apply the changes to the system.
- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a To add a time range, click **Add** and configure a name for the time range configuration.




Figure 140: Add Time Range

- a To delete a configured time range, select each entry to delete, click **Remove**, and confirm the action.

3.19.2 Time Range Entry Summary

Use this page to configure periodic and absolute time range entries and add them to named time ranges.

- 4 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**.

Time Range Name: test

Display: All rows | Showing 0 to 0 of 0 entries | Filter:

Entry Type	Starts	Ends
Table is Empty		

First Previous Next Last

Refresh Add Absolute Add Periodic Remove

Figure 141: Time Range Entry Summary

Table 132: 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.
Entry Type	The type of time range entry, which is one of the following: - a Absolute – Occurs once or has an undefined start or end period. The duration of an Absolute entry can be hours, days, or even years. Each time entry configuration can have only one Absolute entry. - a Periodic – Recurring entry that takes place at fixed intervals. This type of entry occurs at the same time on one or more days of the week.
Starts	For an Absolute entry, indicates the time, day, month, and year that the entry begins. If this field is blank, the Absolute entry became active when it was configured. For a Periodic entry, indicates the time and day(s) of the week that the entry begins.
Ends	For an Absolute entry, indicates the time, day, month, and year that the entry ends. If this field is blank, the Absolute entry does not have a defined end. For a Periodic entry, indicates the time and day(s) of the week that the entry ends.

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:

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

3 Configuring and viewing System Information

- a 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.

Add Absolute Time Range

Time Range Name

test

Start Time

☐

Start Date

Starting Time of Day

(00:00 to 23:59)

End Time

☐

End Date

Ending Time of Day

(00:00 to 23:59)

Submit

Cancel

Figure 142: Add Absolute Time Range

Table 133: Add Absolute Time Range Fields

Field	Description
Time Range Name	Displays the selected time range.
Start Time	Select this option to configure values for the Start Date and the Starting Time of Day. If this option is not selected, the entry becomes active immediately.
Start Date	Click the calendar icon to select the day, month, and year when this entry becomes active. This field can be configured only if the Start Time option is selected.
Starting Time of Day	Specify the time of day that the entry becomes active by entering the information in the field or by using the scroll bar in the Choose Time window. Click Now to use the current time of day. Click Done to close the Choose Time window. This field can be configured only if the Start Time option is selected. 5 The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.
End Time	Select this option to configure values for the End Date and the Ending Time of Day. If this option is not selected, the entry does not have an end time; after the configured Start Time begins, the entry will remain active indefinitely.
End Date	Click the calendar icon to select the day, month, and year when this entry should no longer be active. This field can be configured only if the End Time option is selected.
Ending Time of Day	Specify the time of day that the entry becomes inactive by entering the information in the field or by using the scroll bar in the Choose Time window. Click Now to use the current time of day. Click Done to close the Choose Time window. This field can be configured only if the End Time option is selected. 5 The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.

- a To add a Periodic time range entry, click **Add Periodic** and specify the days and times that the entry is in effect.

Figure 143: Add Periodic Time Range

Table 134: Add Periodic Time Range Fields

Field	Description
Time Range Name	Displays the selected time range.
Applicable Days	Specify the days when the time entry occurs: - a Daily – Has the same start and end time every day. - a Weekdays – Has the same start and end time Monday through Friday. - a Weekend – Has the same start and end time on Saturday and Sunday. - a Days of 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.
Start Days	Indicates on which days the time entry becomes active. If the selected option in the Applicable Days field is Days of Week, select one or more days on which the entry becomes active. To select multiple days, hold the Ctrl key and select each desired start day.
Starting Time of Day	Specify the time of day that the entry becomes active by entering the information in the field or by using the scroll bar in the Choose Time window. Click Now to use the current time of day. Click Done to close the Choose Time window. 5 The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.
End Days	Indicates on which days the time entry ends. If the selected option in the Applicable Days field is Days of Week, select one or more days on which the entry ends. To select multiple days, hold the Ctrl key and select each desired end day.
Ending Time of Day	Specify the time of day that the entry becomes inactive by entering the information in the field or by using the scroll bar in the Choose Time window. Click Now to use the current time of day. Click Done to close the Choose Time window.

3 Configuring and viewing System Information

Field	Description
	5 The time is based on a 24-hour clock. For example, 6:00 PM is 18:00.

- a To delete a time range entry, select each entry to delete, click **Remove**, and confirm the action.

3.20 Configuring DNS

You can use these menus to configure information about DNS servers the network uses and how the switch/router operates as a DNS client.

3.20.1 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**.

ConfigurationIP MappingSource Interface Configuration

DNS Global Configuration

Admin Mode

☒ Enable ☐ Disable

Default Domain Name

(Max 255 characters)

Retry Number

(0 to 100)

Response Timeout (secs)

(0 to 3600)

Domain List

DNS Server

Submit

Refresh

Cancel

Figure 144: DNS Global Configuration

Table 135: DNS Global Configuration Fields

Field	Description
Admin Mode	Select Enable or Disable from the menu to set the administrative status of DNS Client. The default is Enable .
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 <i>is.com</i> and the user enters <i>hotmail</i> , then <i>hotmail</i> is changed to <i>hotmail.com</i> to resolve the name). By default, no default domain name is configured in the system.
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 255 characters. Multiple default domain names can be configured using the default domain-name list.

Field	Description
	If there is no domain list, the default domain name configured is used. To add an entry to the Domain List, click on the Plus button (+) next to Domain List. Then enter a Domain Name and click Submit. Repeat this step to add multiple domains to the domain list. To remove a domain from the Domain List click on the Minus button (-) next to the item you want to remove.
DNS Server	A unique IPv4 or IPv6 address used to identify a DNS server. The order in which you add servers determines the precedence of the server. The DNS server that you add first has the highest precedence and will be used before other DNS servers that you add. To add an entry to the Domain Server list, click on the Plus button (+) next to DNS Server. Then enter a DNS Server and click Submit. Repeat this step to add multiple domains to the domain list. To remove a domain from the DNS Server list, click on the Minus button (-) next to the item you want to remove.

Use the buttons to perform the following tasks:

- a If you change any settings, click **Submit** to send the information to the system.
- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.20.2 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 > IP Mapping Summary** in the menu.

Figure 145: DNS IP Mapping Summary

Table 136: DNS IP Mapping Summary Fields

Field	Description
Entry Type	Type of DNS entry: - a Static – An entry that has been manually configured on the device. - a Dynamic – An entry that the device has learned by using a configured DNS server to resolve a hostname.

Field	Description
Host Name	The name that identifies the system. For Static entries, specify the Host Name after you click Add . A host name can contain up to 255 characters if it contains multiple levels in the domain hierarchy, but each level (the portion preceding a period) can contain a maximum of 63 characters. If the host name you specify is a single level (does not contain any periods), the maximum number of allowed characters is 63.
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.
Total Time	The number of seconds that the entry will remain in the table (dynamic entries only).
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 entries only).
Dynamic Type	The type of address in the entry, for example IP or (less common) X.121 (dynamic entries only).

Use the buttons to perform the following tasks:

- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a Click **Add** to load the Add DNS Entry page to configure the Host Name IP Mapping entries.

Figure 146: Add DNS Entry

Table 137: 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.

- a Select the DNS IP Mapping entries and click **Remove** to delete these entries.

3.20.3 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 147: DNS Source Interface Configuration

Table 138: DNS Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: - a None – The primary IP address of the originating (outbound) interface is used as the source address. - a Interface – The primary IP address of a physical port is used as the source address. - a VLAN – The primary IP address of a VLAN routing interface is used as the source address. - a Tunnel – The primary IP address of a tunnel interface is used as the source address. - a Network – The network source IP is used as the source address. - a 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.
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.
Loopback Interface	When the selected Type is Loopback , select the loopback interface to use as the source interface.
Tunnel ID	When the selected Type is Tunnel , select the tunnel interface to use as the source interface.

Use the buttons to perform the following tasks:

- a If you change any of the settings on the page, click **Submit** to apply the changes to system.
- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.21 Configuring SNTP Settings

LCOS SX 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. LCOS SX 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:

- a Stratum 0: A real time clock is used as the time source, for example, a GPS system.
- a 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.
- a 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:

- a T1: Time at which the original request was sent by the client.
- a T2: Time at which the original request was received by the server.
- a T3: Time at which the server sent a reply.
- a 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:

- a 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.
- a If more than one Unicast device responds, synchronization information is preferred from the device with the lowest stratum.
- a 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.

3.21.1 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 148: SNTP Global Configuration

Table 139: SNTP Global Configuration Fields

Field	Description
Client Mode	Use drop-down list specify the SNTP client mode, which is one of the following modes: - a Disable – SNTP is not operational. No SNTP requests are sent from the client nor are any received SNTP messages processed. - a 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. - a 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.
Port	Specifies the local UDP port to listen for responses/broadcasts. Allowed range is (1025 to 65535). Default value is 123.
Unicast Poll Interval	Specifies the number of seconds between unicast poll requests. Allowed range is (6 to 10). Default value is 6.
Broadcast Poll Interval	Specifies the number of seconds between broadcast poll requests. 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.


Use the buttons to perform the following tasks:

- a 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**
- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.21.2 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.



Global Configuration Global Status Server Configuration Server Status Source Interface Configuration	
SNTP Global Status	
Version	4
Supported Mode	Unicast and Broadcast
Last Update Time	Jan 1 00:00:00 1970
Last Attempt Time	Jan 1 00:00:00 1970
Last Attempt Status	Other
Server IP Address	
Address Type	Unknown
Server Stratum	0
Reference Clock ID	
Server Mode	Reserved
Unicast Server Max Entries	3
Unicast Server Current Entries	0
Broadcast Count	0

Refresh

Figure 149: SNTP Global Status

Table 140: 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: - a Other – None of the following enumeration values. - a Success – The SNTP operation was successful and the system time was updated. - a Request Timed Out – A directed SNTP request timed out without receiving a response from the SNTP server. - a Bad Date Encoded – The time provided by the SNTP server is not valid. - a Version Not Supported – The SNTP version supported by the server is not compatible with the version supported by the client. - a Server Unsynchronized – The SNTP server is not synchronized with its peers. This is indicated via the 'leap indicator' field on the SNTP message. - a 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.

Field	Description
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.

3.21.3 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 150: SNTP Server Configuration

Table 141: SNTP Server Configuration Fields

Field	Description
SNTP Server	The address or host name of an SNTP server the device can use to synchronize the system time.
Type	The configured SNTP server address type, which can be ipv4 , ipv6, or DNS.
Port	The UDP port on the server to which SNTP requests are sent.
Priority	The order in which to query the servers. The SNTP client on the device continues sending SNTP requests to different servers until a successful response is received or all servers are exhausted. A server entry with a lower priority value is queried before one with a higher priority. If more than one server has the same priority, the SNTP client contacts the servers in the order that they appear in the table.

Field	Description
Version	Specifies the NTP version running on the server.

Use the buttons to perform the following tasks:

- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a To add an SNTP server, click **Add** and configure the desired settings.

Host Name or IP Address (Max 64 characters or x.x.x.x or x.x.x.x.x:x.x.x.x)

Port 123 (1 to 65535)

Priority 1 (1 to 3)

Version 4 (1 to 4)

Submit Cancel

Figure 151: Add SNTP Server

Table 142: Add SNTP Server Fields

Field	Description
Host Name or IP Address	Specify the IPv4 address, IPv6 address, or DNS-resolvable host name of the SNTP server. Unicast SNTP requests will be sent to this address. The address you enter is displayed in the SNTP Server field on the main page. The address type is automatically detected. You can define up to three SNTP servers.
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.

- a To change information for an existing SNTP server, select the entry to update and click **Edit**. You cannot edit the host name or address of a server that has been added.
- a To delete a configured SNTP server from the list, select each entry to delete and click **Remove**.

3.21.4 SNTP Server Status

The SNTP Server Status page displays status information about the SNTP servers configured on your switch.

To access the SNTP Server Status page, click **System > Advanced Configuration > SNTP > Server Status** in the navigation menu.




Figure 152: SNTP Server Status

Table 143: SNTP Server Status Fields

Field	Description
Address	Specifies all the existing Server Addresses.
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: - a Other – None of the following enumeration values. - a Success – The SNTP operation was successful and the system time was updated. - a Request Timed Out – A directed SNTP request timed out without receiving a response from the SNTP server. - a Bad Date Encoded – The time provided by the SNTP server is not valid. - a Version Not Supported – The SNTP version supported by the server is not compatible with the version supported by the client. - a Server Unsynchronized – The SNTP server is not synchronized with its peers. This is indicated via the 'leap indicator' field on the SNTP message. - a 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.
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.

3.21.5 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.

3 Configuring and viewing System Information

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 153: SNTP Source Interface Configuration

Table 144: SNTP Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: - a None – The primary IP address of the originating (outbound) interface is used as the source address. - a Interface – The primary IP address of a physical port is used as the source address. - a VLAN – The primary IP address of a VLAN routing interface is used as the source address. - a 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. - a Network – The network source IP is used as the source address.
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.
Loopback Interface	When the selected Type is Loopback , select the loopback interface to use as the source interface.
Tunnel ID	When the selected Type is Tunnel, select the tunnel interface to use as the source interface.

Use the buttons to perform the following tasks:

- a 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**
- a Use the **Refresh** button to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.22 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.

Current Time	
Time	20:04:39
Zone	(UTC+0:00)
Date	January 08, 1970
Time Source	No time source

Time Zone	
Zone	
Offset	UTC+0:00

Summer Time	
Summer Time	No Summer Time
Zone	
Offset	
Status	

Refresh

Figure 154: Time Zone Summary

Table 145: Time Zone Summary Fields

Field	Description
Current Time	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 rebooted. - a 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. - a Zone — The acronym that represents the time zone. - a Date — The current date on the system. - a Time Source — The time source from which the time update is taken: - a SNTP — The time has been acquired from an SNTP server. - a No time source — The time has either been manually configured or not configured at all.
Time Zone	This section contains information about the time zone and offset. Zone — The acronym that represents the time zone. Offset — The number of hours offset from Coordinated Universal Time (UTC), which is also known as Greenwich Mean Time (GMT).

Field	Description
Summer Time	This section contains information about the summer time. - a Summer Time – The administrative status of summer time (daylight saving time). In some regions, the time shifts by one hour in the fall and spring. - a No Summer Time – Summer time is not active and the time does not shift based on the time of year. - a Recurring every year – Summer time occurs at the same time every year. The start and end times and dates for the time shift have to be configured manually. - a Non recurring Summer Time – 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. - a Zone – The acronym that represents the time zone of the summer time. - a Offset – The number of hours offset from Coordinated Universal Time (UTC), which is also known as Greenwich Mean Time (GMT). - a Status – Indicates if summer time is currently active.

Click **Refresh** to display the latest information from the router.

3.22.1 Time Zone Configuration

Use this page to manually configure the system clock settings. The [SNTP client](#) must be disabled (**Client Mode** must be set to **Disable**) 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 155: Time Zone Configuration

Table 146: Time Zone Configuration Fields

Field	Description
Time Zone	The time zone settings include the amount of time the system clock is offset from Coordinated Universal Time (UTC) and the time zone acronym. - a Offset — The number of hours the system clock is offset from UTC, which is also known as Greenwich Mean Time (GMT). - a Zone — The acronym that represents the time zone. This field is not validated against an official list of time zone acronyms.

Field	Description
Date and Time	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. a Time — The current time in hours, minutes, and seconds on the system clock. a 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.

Use the buttons to perform the following tasks:

- a Click **Submit** to apply the settings to the running configuration and cause the change to take effect.
- a Click **Refresh** to display the latest information from the router.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.22.2 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.

Summer Time Configuration

Summer Time: Disable

Date Range

Start Date:

Starting Time of Day: (00:00 to 23:59)

End Date:

Ending Time of Day: (00:00 to 23:59)

Recurring Date

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

Ending Time of Day: (00:00 to 23:59)

Zone

Offset: (1 to 1440)

Zone: (0 to 4 characters)

Submit Refresh Cancel

Figure 156: Summer Time Configuration

Table 147: Summer Time Configuration Fields

Field	Description
Summer Time	The summer time mode on the system: - a Disable – Summer time is not active, and the time does not shift based on the time of year. - a 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. - a 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. - a 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. - a 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.
Date Range	The fields in this section are available only if the Non-Recurring mode is selected from the Summer Time menu. - a 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. - a Starting Time of Day — The time, in hours and minutes, to start summer time on the specified day. - a 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. - a Ending Time of Day — The time, in hours and minutes to end summer time on the specified day.
Recurring Date	The fields in this section are available only if the Recurring mode is selected from the Summer Time menu. - a Start Week — The week of the month within which summer time begins. - a Start Day — The day of the week on which summer time begins. - a Start Month — The month of the year within which summer time begins. - a Starting Time of Day — The time, in hours and minutes, to start summer time. - a End Week — The week of the month within which summer time ends. - a End Day — The day of the week on which summer time ends. - a End Month — The month of the year within which summer time ends. - a Ending Time of Day — The time, in hours and minutes, to end summer time.
Zone	The fields in this section are available only if the Recurring or Non-Recurring modes are selected from the Summer Time menu. - a Offset — The number of minutes to shift the summer time from the standard time. - a Zone — The acronym associated with the time zone when summer time is in effect.

Use the buttons to perform the following tasks:

- a Click **Submit** to apply the settings to the running configuration and cause the change to take effect.
- a Click **Refresh** to display the latest information from the router.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.23 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.

LCOS SX participates in the CDP protocol and is able to both discover and be discovered by other CDP supporting devices.

3.23.1 ISDP Global Configuration

To access the ISDP Global Configuration page, click **System > Advanced Configuration > ISDP > Global** in the navigation menu.

ISDP Mode	☒ Disable ☐ Enable
ISDP V2 Mode	☐ Disable ☒ Enable
Message Interval (Seconds)	30 (5 to 254)
Hold Time Interval (Seconds)	180 (10 to 255)
Device ID	XS-5110F
Device ID Format Capability	Serial Number, Host Name
Device ID Format	Host Name

Submit Refresh Cancel

Figure 157: ISDP Global Configuration

The following table describes the fields available on the ISDP Global Configuration page.

Table 148: 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.

Field	Description
Device ID Format Capability	Indicates the Device ID format capability of the device. a Serial Number — Indicates that the device uses serial number as the format for its Device ID. a Host Name — Indicates that the device uses a host name as the format for its Device ID.
Device ID Format	Indicates the Device ID format of the device. a Serial Number — Indicates that the device uses serial number as the format for its Device ID. a Host Name — Indicates that the device uses a host name as the format for its Device ID.

Use the buttons to perform the following tasks:

- a Click **Submit** to apply the settings to the running configuration and cause the change to take effect.
- a Click **Refresh** to display the latest information from the router.
- a Click **Cancel** to discard changes and revert to the last saved state.

3.23.2 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.

Global	Cache Table	Interface	Statistics
ISDP Cache Table			
Display All rows		Showing 0 to 0 of 0 entries	
		Filter:	
Device ID	Interface	IP Address	Version
Hold Time (Seconds)	Capability	Platform	Port ID
Protocol Version	Last Time Changed		
Table is Empty			
First Previous Next Last			
Refresh Clear			

Figure 158: ISDP Cache Table

The following table describes the fields available on the ISDP Cache Table page.

Table 149: 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.

Field	Description
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the router.
- a Use the **Clear** button to clear all entries from the table. The table is repopulated as ISDP messages are received from neighbors.

3.23.3 ISDP Interface Configuration

From the ISDP Interface Configuration page, you can configure the ISDP settings for each interface.

- 4** 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 159: ISDP Interface Configuration

The following table describes the fields available on the ISDP Interface Configuration page.

Table 150: ISDP Interface Configuration Fields

Field	Description
Interface	Select the interface with the ISDP mode status to configure or view.

Field	Description
ISDP Mode	Use this field to enable or disable the Industry Standard Discovery Protocol on the selected interface.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the router.
- a To change the ISDP mode for one or more interfaces, select each interface to update and click **Edit**.





Figure 160: Edit ISDP Mode

3.23.4 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.



Packets Received	0
Packets Transmitted	7676
ISDPv1 Packets Received	0
ISDPv1 Packets Transmitted	0
ISDPv2 Packets Received	0
ISDPv2 Packets Transmitted	7676
Bad Header	0
Checksum Error	0
Transmission Failure	0
Invalid Format Packets Received	0
Table Full	0
ISDP IP Address Table Full	0

Figure 161: ISDP Statistics

The following table describes the fields available on the ISDP Statistics page.

Table 151: ISDP Statistics Fields

Field	Description
Packets Received	Displays the number of all ISDP protocol data units (PDUs) received.
Packets Transmitted	Displays the number of all ISDP PDUs transmitted.

Field	Description
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.
Bad Header	Displays the number of ISDP PDUs that were received with bad headers.
Checksum Error	Displays the number of ISDP PDUs that were received with checksum errors.
Transmission Failure	Displays the number of ISDP PDUs transmission failures.
Invalid Format 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 cache table but was unsuccessful because the table was full (more than 2700 entries).
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 (more than 5400 entries).

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Use the **Clear** button reset all statistics to zero.

3.24 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.

3.24.1 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 162: Link Dependency Group Status

Table 152: 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: - a Up – Downstream interfaces are up when upstream interfaces are down. - a Down – Downstream interfaces go down when upstream interfaces are down. 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: - a 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. - a Down – Link is down when the above conditions are not true.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the router.

- a To add a link dependency group, click **Add**. Then, specify a group number, link action, and the interfaces that share a dependency.

Figure 163: Add Group

- a **Table 153: Add Group Fields**

Field	Description
Group	Specify a unique link dependency group identifier.
Link Action	The action performed on downstream interfaces when the upstream interfaces are down, which can be one of the following: - a Down – Downstream interfaces go down when upstream interfaces are down. - a Up – Downstream interfaces are up when upstream interfaces are down. 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.
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.

- a To change the settings for a group, select the check box associated with the group and click **Edit**.
- a To delete a link dependency group, select the check box associated with each entry to delete and click **Remove**.

3 Configuring and viewing System Information

- a To view additional information about a group, select the check box associated with the group and click **Details**.



Table 154: Group Entry Details Fields

Fields	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. a Link Up – This field lists the downstream interfaces that currently have their link up. a Link Down – This field lists the downstream interfaces that currently have their link down.
Upstream Interfaces	The set of interfaces that determine the link state of the downstream interfaces. a Link Up – This field lists the upstream interfaces that currently have their link up. a Link Down – This field lists the upstream interfaces that currently have their link down.

Figure 164: Group Entry Details

4 Configuring Switching Information

4.1 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 [Configuring DHCP Snooping](#) on page 269) 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.

4.1.1 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 165: Device Tracking Global Configuration

Table 155: Device Tracking Global Configuration Fields

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.

4 Configuring Switching Information

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.

Use the buttons to perform the following tasks:

- a Click **Submit** to send the updated configuration to the switch.
- a Click **Refresh** to update the page with the most current information.
- a Click **Cancel** to cancel the changes.

4.1.2 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 166: Device Tracking Summary

Table 156: Device Tracking Summary Fields

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 where the device is learned.
Interface	The interface where the device is learned.
State	The state of the learned device in the IPDT table, which can be one of the following: - a Active – The device is reachable and responding to the ARP probes sent. - a Inactive – The device is not reachable.
Time Left to Inactive (Seconds)	The number of seconds remaining before an active device in the IPDT table is marked inactive.

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: - a ARP – ARP snooping detected new devices. - a DHCP – DHCP snooping detected DHCP client devices.

The administrative mode of the IPDT feature must be enabled on the device to clear the IPDT table entries.

- a Click **Refresh** to update the page with the most current information.
- a 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.

4.1.3 Device Tracking Interface Configuration

Use the Device Tracking Interface Configuration page to configure the IPDT settings on specific interfaces.

To access the Device Tracking Interface Configuration page, click **Switching** > **Device Tracking** > **Interface**.

The screenshot shows the 'Device Tracking Interface Configuration' page. It features a table with two columns: 'Interface' and 'Maximum Learned Entries'. The table contains 10 rows, each representing an interface from 1/0/1 to 1/0/10, with the value 'No Limit' in the second column. Navigation buttons (First, Previous, 1, 2, 3, 4, 5, Next, Last) are located below the table. At the bottom of the page, there are 'Refresh' and 'Edit' buttons.

Figure 167: Device Tracking Interface Configuration

Table 157: Device Tracking Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
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 0, 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.

Field	Description
	By default, there is no limit on the number of entries that can be learned per interface. You can achieve that by leaving the field empty.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the page with the most current information.
- a 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.

Table 158: Edit Port Configuration Fields

Field	Description
Interface	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 0, 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. You can achieve that by leaving the field empty.

Figure 168: Edit Port Configuration

4.1.4 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**.

Device Tracking Statistics	
ARP Entries	0
DHCP Entries	0
Active Entries	0
Inactive Entries	0
Total Entries	0

Figure 169: Device Tracking Statistics

Table 159: Device Tracking Statistics Fields

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.

4.2 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 connections to 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.

4.2.1 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.

4 Configuring Switching Information

To display this page, click **Switching > Loop Protection > Configuration** in the navigation menu.

Configuration
Loop Protection Configuration

Loop Protection ☐ Enabled ☒ Disabled

Transmission Time (Seconds) (1 to 10)

Maximum PDU Received (1 to 10)

Display rows Showing 1 to 10 of 118 entries Filter:

☐	Interface	Loop Protection	Action	Status	Loop	Loop Count	Time of Last Loop
☐	1/0/1	Disabled	Shutdown Port	Link Up		0	
☐	1/0/2	Disabled	Shutdown Port	Link Down		0	
☐	1/0/3	Disabled	Shutdown Port	Link Down		0	
☐	1/0/4	Disabled	Shutdown Port	Link Up		0	
☐	1/0/5	Disabled	Shutdown Port	Link Down		0	
☐	1/0/6	Disabled	Shutdown Port	Link Down		0	
☐	1/0/7	Disabled	Shutdown Port	Link Down		0	
☐	1/0/8	Disabled	Shutdown Port	Link Down		0	
☐	1/0/9	Disabled	Shutdown Port	Link Down		0	
☐	1/0/10	Disabled	Shutdown Port	Link Down		0	

First Previous 1 2 3 4 5 Next Last

Submit Refresh Edit Edit All

Figure 170: Loop Protection Configuration

Table 160: Loop Protection Configuration Fields

Field	Description
Loop Protection	Enables or disables the loop protection feature globally on the switch. 4 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.
Loop Protection	Indicates if loop protection is active (Enabled) or inactive (Disabled) on a specific interface.
Action	The action to be taken when a loop is detected on the port: a Log Only – Send a message to the system log but do not shut down the port. a Shutdown Port – Shut down the port for the configured Transmission Time. a Shutdown Port and Log – Shut down the port for the configured Transmission Time and send a message to the system log.
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 or that the interface doesn't have a network connection.
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.

Field	Description
Time of Last Loop	The date and time the most recent loop was detected.

Use the buttons to perform the following tasks:

- a Click **Submit** to update the switch. The changes take effect immediately but will not be retained across a power cycle unless a save is performed.
- a Click **Refresh** to update the information on the screen with the most current data.
- a To configure the settings for one or more interfaces, select each entry to modify and click **Edit**.

Figure 171: Edit Loop Protection Port Configuration

Table 161: Edit Loop Protection Port Configuration Fields

Field	Description
Interface	The port or trunk ID.
Loop Protection	Enables or disables the loop protection feature for the selected interface.
Action	The action to be taken when a loop is detected on the port: - a Log Only - a Shutdown Port - a Shutdown Port and Log

- a To apply the same settings to all interfaces, click **Edit All**.

4.3 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 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.

4.3.1 VLAN Overview

Use the VLAN Overview 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 > Overview** in the navigation menu.




Figure 172: VLAN Overview

Table 162: VLAN Overview Fields

Field	Description
VLAN ID	The VLAN Identifier (VID) of the VLAN. The range of the VLAN ID is 1 to 4093.
Name	The name of the VLAN. VLAN ID 1 is always named default .
Type	The VLAN type, which can be one of the following: - a Default – The default VLAN. This VLAN is always present, and the VLAN ID is 1. - a Static – A user-configured VLAN. - a Dynamic – A VLAN created by GVRP registration that you have not converted to static, and that GVRP may therefore remove
RSPAN	Lists the status of RSPAN, enabled or disabled.
Statistics Mode	Enable or disable the statistics collection mode on the VLAN. Statistics Mode cannot be configured on the default and dynamic VLANs. 5 This parameter is only available on XS-6128QF switches.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the router.
- a To add a VLAN, click **Add** and specify a VLAN ID in the available field.





Figure 173: Add VLAN

- a To configure a name for a VLAN or to convert a dynamic VLAN to a static VLAN, select the entry to modify and click **Edit**. Then, configure the desired VLAN settings.

- a 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.

4.3.1.1 Add a VLAN

To add a VLAN, click the **Add** button and specify a VLAN ID in the available field.




The 'Add VLAN' dialog box contains a text input field labeled 'VLAN ID or Range'. Below the field, a note states: '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.' At the bottom right, there are 'Submit' and 'Cancel' buttons.

Figure 174: Add a VLAN

Click **Submit** to add the VLAN to the system.

4.3.1.2 Edit VLAN Configuration

To edit the VLAN configuration, select the entry to modify and click the **Edit** button.



The 'Edit VLAN Configuration' dialog box displays the following fields: 'VLAN ID' with the value '20'; 'Name' with the value 'VLAN0020' and a note '(1 to 32 alphanumeric characters)'; 'Type' set to 'Static'; 'Convert VLAN Type to Static' with an unchecked checkbox; and 'Statistics Mode' with 'Enable' selected (radio button) and 'Disable' as an option. 'Submit' and 'Cancel' buttons are at the bottom right.

Figure 175: Edit VLAN Configuration

Edit the configured VLAN settings, as follows.

Table 163: Edit VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN Identifier (VID) of the VLAN. The range of the VLAN ID is 1 to 4093.
Name	For static VLANs, you can modify the VLAN name. The name can be 1 to 32 alphanumeric characters. This field is not available for other VLAN types.
Type	The VLAN type, which can be one of the following: - a Default: The default VLAN. This VLAN is always present, and the VLAN ID is 1. - a Static: A user-configured VLAN. - a Dynamic: A VLAN created by GVRP registration that you have not converted to static, and that GVRP may therefore remove
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

Field	Description
	dynamic VLANs. If you convert a dynamic VLAN to a static VLAN, it cannot be removed by GVRP.
Statistics Mode	Enable or disable the statistics collection mode on the VLAN. Statistics Mode cannot be configured on the default and dynamic VLANs.
	5 This parameter is only available on XS-6128QF switches.

Click **Submit** to submit the VLAN configuration changes. Click **Cancel** to cancel the changes.

4.3.1.3 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.

4.3.2 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 176: VLAN Port Configuration

Table 164: 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: a Include – The port is a member of the selected VLAN.

Field	Description
	a Exclude – The port is not a member of the selected VLAN.
Participation	The participation mode of the interface in the selected VLAN, which is one of the following: a Include – The port is always a member of the selected VLAN. This mode is equivalent to registration fixed in the IEEE 802.1Q standard. a Exclude – The port is never a member of the selected VLAN. This mode is equivalent to registration forbidden in the IEEE 802.1Q standard. a 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.
Tagging	The tagging behavior for all the ports in this VLAN, which is one of the following: a Tagged – The frames transmitted in this VLAN will include a VLAN ID tag in the Ethernet header. a Untagged – The frames transmitted in this VLAN will be untagged.

Use the buttons to perform the following tasks:

- a To reload the page and view the most current information, click **Refresh**.
- a 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.

Figure 177: Edit VLAN Port Configuration

Table 165: Edit VLAN Port Configuratuion Fields

Field	Description
VLAN ID	Shows the selected VLAN ID.
Interface	Shows the selected interface(s).
Participation	The participation mode of the interface in the selected VLAN, which is one of the following: a Include a Exclude a Auto Detect
Tagging	The tagging behavior for all the ports in this VLAN, which is one of the following: a Tagged a Untagged

- a To apply the same settings to all interfaces, click **Edit All** and configure the desired settings.

4.3.3 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.

Interface	Port VLAN ID	Port VLAN ID Current	Acceptable Frame Type	Ingress Filtering	Untagged VLANs	Tagged VLANs	Forbidden VLANs	Dynamic VLANs	Priority
1/0/1	1	1	Admit All	Disabled	1				0
1/0/2	1	1	Admit All	Disabled	1				0
1/0/3	1	1	Admit All	Disabled	1				0
1/0/4	1	1	Admit All	Disabled	1				0
1/0/5	1	1	Admit All	Disabled	1				0
1/0/6	1	1	Admit All	Disabled	1				0
1/0/7	1	1	Admit All	Disabled	1				0
1/0/8	1	1	Admit All	Disabled	1				0
1/0/9	1	1	Admit All	Disabled	1				0
1/0/10	1	1	Admit All	Disabled	1				0

Figure 178: VLAN Port Summary

Table 166: 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.
Port VLAN ID Current	The current VLAN ID assigned to packets received on this port. This value may differ from the configured VLAN, for example, when access to the port is managed by a Radius server in an 802.1x configuration.
Acceptable Frame Type	Indicates how the interface handles untagged and priority tagged frames. The options include the following: - a 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. - a Only Tagged – The interface discards any untagged or priority tagged frames it receives. - a Only Untagged – The interface discards any tagged frames it receives. 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. a Enable – A tagged frame is discarded if this port is not a member of the VLAN identified by the VLAN ID in the tag.

Field	Description
	a Disable – All tagged frames are accepted, which is the factory default.
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:

- a To reload the page and view the most current information, click **Refresh**.
- a 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.

Table 167: Edit Port Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
Port VLAN ID	Enter a VLAN ID to be 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 Type	Select the behavior how the interface should handle untagged and priority tagged frames. The options include the following: - a Admit All - a Only Tagged - a Only Untagged For all options, VLAN tagged frames are forwarded in accordance with the IEEE 802.1Q VLAN standard.
Ingress Filtering	Select the behavior how the port should handle tagged frames. - a Enable - a Disable

Field	Description
Priority	Identifies the default 802.1p priority assigned to untagged packets arriving at the port.

Figure 179: Edit Port Configuration

- a To apply the same settings to all interfaces, click **Edit All** and configure the desired settings.

4.3.4 VLAN Switchport Summary

Use the VLAN 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 VLAN Switchport Summary page, click **Switching > VLAN > Switchport Summary** in the navigation menu.

Overview

Port Configuration

Port Summary

Switchport Summary

Internal Usage

Reset

RSPAN

VLAN Switchport Summary

Display10rows

Showing 1 to 10 of 118 entries

Filter:

☐	Interface	Switchport Mode	Operational Switchport Mode	Access VLAN ID	Native VLAN ID	Native VLAN Tagging	Trunk Allowed VLANs
☐	1/0/1	General	General	1	1	Disabled	1-4093
☐	1/0/2	General	General	1	1	Disabled	1-4093
☐	1/0/3	General	General	1	1	Disabled	1-4093
☐	1/0/4	General	General	1	1	Disabled	1-4093
☐	1/0/5	General	General	1	1	Disabled	1-4093
☐	1/0/6	General	General	1	1	Disabled	1-4093
☐	1/0/7	General	General	1	1	Disabled	1-4093
☐	1/0/8	General	General	1	1	Disabled	1-4093
☐	1/0/9	General	General	1	1	Disabled	1-4093
☐	1/0/10	General	General	1	1	Disabled	1-4093

First

Previous

1

2

3

4

5

Next

Last

Refresh

Edit

Edit All

Figure 180: VLAN Switchport Summary**Table 168: 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	The switchport mode of the interface, which is one of the following: - a 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. - a 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. - a 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.

Field	Description
	- a 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). - a 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. - a 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. 5 The modes Private VLAN Host, Private VLAN Promiscuous and Private VLAN Promiscuous Trunk are not available in this menu but have to be configured in the Private VLAN Interface Association.
Operational Switchport Mode	The operational switchport mode of the interface, which is one of the following: - a Access - a Trunk - a General - a Private VLAN Host - a Private VLAN Promiscuous - a Private VLAN Promiscuous Trunk
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:

- a To reload the page and view the most current information, click **Refresh**.

4 Configuring Switching Information

- a 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.

Interface: 1/0/1

Switchport Mode: ☒ General ☐ Access ☐ Trunk

Access VLAN ID: 1 (1 to 4093)

Native VLAN ID: 1 (1 to 4093)

Trunk Allowed VLANs: ☒ None ☐ Add ☐ Remove

Enter VLAN ID in the range 1 to 4093.
Use ":" to specify a range and "," to separate VLAN IDs or VLAN ranges in the list.

Submit Cancel

Figure 181: Edit Switchport Configuration

Table 169: Edit Switchport Configuration Fields

Field	Description
Interface	When editing information for one or more interfaces, this field identifies the interfaces that are being configured.
Switchport Mode	The switchport mode of the interface, which is one of the following: - a General - a Access - a Trunk
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.
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.

- a To apply the same settings to all interfaces, click **Edit All** and configure the desired settings.

4.3.5 VLAN Internal Usage

Use the VLAN Internal Usage 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 182: VLAN Internal Usage

Table 170: VLAN Internal Usage 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.

Use the buttons to perform the following tasks:

- a If you change any information on the page, click **Submit** to apply the changes to the system.
- a To reload the page and view the most current information, click **Refresh**.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.3.6 Configure VLAN Statistics

Use the VLAN Statistics page to view and clear the statistical information for VLANs on which the statistics mode is enabled.

4 This feature is only supported by the LANCOM XS-6128QF.

Use the [Edit VLAN Configuration](#) on page 211 page to Enable or Disable the statistics collection mode on VLANs.

To access the VLAN Statistics page, click **Switching > VLAN > Statistics** in the navigation menu.

	VLAN ID	Rx Bytes	Rx Frames	Rx Discard Bytes	Rx Discard Frames	Tx Bytes	Tx Frames	Tx Discard Bytes	Tx Discard Frames
☐	1	15036153	79989	14910	129	15356109	79989	0	0
☐	100	0	0	0	0	0	0	0	0

Table 171:

Field	Description
VLAN ID	The VLAN ID associated with the rest of the data in the row.
Rx Bytes	The total number of bytes received on the VLAN.
Rx Frames	The total number of frames received on the VLAN.
Rx Discard Bytes	The total number of bytes received on the VLAN that were discarded.
Rx Discard Frames	The total number of frames received on the VLAN that were discarded.
Tx Bytes	The total number of bytes transmitted on the VLAN.
Tx Frames	The total number of frames transmitted on the VLAN.
Tx Discard Bytes	The total number of bytes transmitted on the VLAN that were discarded.
Tx Discard Frames	The total number of frames transmitted on the VLAN that were discarded.

Figure 183: VLAN-Statistics

Use the buttons to perform the following tasks:

- a To reload the page and view the most current information, click **Refresh**.
- a To reset the counter values on one or more VLANs to the default values, select the VLANs from the list and click the **Clear** button. This opens a modal page as shown in [Figure 184: Clear VLAN-Statistics](#) on page 220. Confirm the action to reset the counter values for the VLANs.




Figure 184: Clear VLAN-Statistics

4.3.7 Reset VLAN Configuration

Use the Reset VLAN Configuration page to return all VLAN parameters for all interfaces to the factory default values. To access the Reset VLAN Configuration page, click **Switching > VLAN > Reset** in the navigation menu.




Figure 185: Reset VLAN Configuration

When you click **Reset** a popup window is displayed and you are asked to confirm the reset. Click **OK** to restore all default VLAN settings for the ports on the system.

4.3.8 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 Configuration page, click **Switching > VLAN > RSPAN** in the navigation menu.




Figure 186: RSPAN Configuration

Table 172: RSPAN 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.

Use the buttons to perform the following tasks:

- a If you change any information on the page, click **Submit** to apply the changes to the system.

- a Click **Refresh** to display the latest information from the router.
- a If you change any information on the page, click **Submit** to apply the changes to the system.

4.4 Configuring UDLD

The UDLD feature (Unidirectional Link Detection) 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 187: Configuring UDLD

Table 173: 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 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.

Use the buttons to perform the following tasks:

- a If you change any information on the page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to display the latest information from the router.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.4.1 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.

Configuration Interface Configuration

UDLD Interface Configuration

Display 10 rows Showing 1 to 10 of 54 entries Filter:

☐	Interface	Admin Mode	UDLD Mode	UDLD Status
☐	1/0/1	Disabled	Normal	Not Applicable
☐	1/0/2	Disabled	Normal	Not Applicable
☐	1/0/3	Disabled	Normal	Not Applicable
☐	1/0/4	Disabled	Normal	Not Applicable
☐	1/0/5	Disabled	Normal	Not Applicable
☐	1/0/6	Disabled	Normal	Not Applicable
☐	1/0/7	Disabled	Normal	Not Applicable
☐	1/0/8	Disabled	Normal	Not Applicable
☐	1/0/9	Disabled	Normal	Not Applicable
☐	1/0/10	Disabled	Normal	Not Applicable

First Previous 1 2 3 4 5 Next Last

Refresh Edit UDLD Port Reset

Figure 188: UDLD Interface Configuration

Table 174: 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.
UDLD Mode	The UDLD mode for the port, which is one of the following: - a 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: - a The UDLD PDU received from a partner does not have its own details (echo). - a When there is a loopback, and information sent out on a port is received back exactly as it was sent. - a 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.
UDLD Status	The UDLD status on the port, which is one of the following: - a Not Applicable – The administrative status of UDLD is globally disabled or disabled on the interface. - a Bidirectional – UDLD has detected a bidirectional link. - a Shutdown – UDLD has detected a unidirectional link, and the port is in a disabled state. To clear the disabled state, click UDLD Port Reset.

Field	Description
	- a Undetermined – UDLD has not collected enough information to determine the state of the port. - a Unknown – The port link has physically gone down, but it is not because it was put in a disabled state by the UDLD feature.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the switch.
- a 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.

Table 175: Edit UDLD Interface Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
Admin Mode	Enable or disable the administrative mode of UDLD on the port.
UDLD Mode	The UDLD mode for the port, which is one of the following: - a Normal - a Aggressive

Figure 189: Edit UDLD Interface Configuration

- a 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.

4.5 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 190: MAC Based VLAN Status

Table 176: 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the switch.
- a To add a MAC-based VLAN, click **Add** and specify a MAC address and a VLAN ID in the available fields.

Table 177: Add MAC Based VLAN Fields

Field	Description
MAC Address	Enter the source MAC address of a 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	Enter 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.

Figure 191: Add MAC Based VLAN

- a 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.

- a 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.

4.6 DVLAN Tunneling

Double VLAN (DVLAN) Tunneling allows the use of a second tag on network traffic. The additional tag helps differentiate between customers in the Metropolitan Area Networks (MANs) while preserving individual customer's VLAN identification when they enter their own IEEE 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.

4.6.1 DVLAN Configuration

The DVLAN Configuration 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 (QinQ) > Configuration** in the navigation menu.

Figure 192: DVLAN Configuration

Table 178: DVLAN Configuration Fields

Field	Description
Primary TPID	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: - a 0x8100 – IEEE 802.1Q customer VLAN tag type - a 0x88a8 – Virtual Metropolitan Area Network (VMAN) tag type - a Custom Tag – User-defined EtherType value To change the Primary TPID, click the Edit icon and select an option from the menu.

Field	Description
	5 The options 0x88a8 and Custom Tag are only available if they have been specified as a Secondary TPID beforehand.
Secondary TPIDs	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: - a 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. - a vMAN Tag – Virtual Metropolitan Area Network (VMAN) 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. - a Custom Tag – User-defined EtherType value. If you select this option, specify the EtherType value in the available field. 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.

Click **Refresh** to display the latest information from the switch.

4.6.2 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 (QinQ) > Summary** in the navigation menu.

The screenshot shows the 'DVLAN Summary' page. At the top, there are three tabs: 'Configuration', 'Summary' (which is active), and 'Interface Summary'. Below the tabs is a header 'DVLAN Summary' with a help icon. The main content area contains a table with two rows: 'Primary TPID' with the value '0x8100' and 'Secondary TPIDs' which is currently empty. Below this table is a 'Refresh' button.

Figure 193: DVLAN Summary Fields

Table 179: DVLAN Summary Fields

Field	Description
Primary TPID	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: - a 0x8100 – IEEE 802.1Q VLAN tag type. This value indicates that the frame includes a VLAN tag. - a 0x88a8 – Virtual Metropolitan Area Network (VMAN) tag type. This value indicates that the frame is double VLAN tagged. - a Custom Tag – Any TPID value other than 0x8100 or 0x88a8 is a user-defined EtherType value.
Secondary TPIDs	The two-byte hex EtherType values configured as secondary TPIDs.

Click **Refresh** to display the latest information from the router.

4.6.3 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 IEEE 802.1Q domain.

To access the DVLAN Interface Summary page, click **Switching > DVLAN (QinQ) > Interface Summary** in the navigation menu.




Figure 194: DVLAN Interface Summary

Table 180: 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to redisplay the most current information from the switch.

- a To configure the DVLAN settings for an interface, select the interface to configure and click **Edit**.

Table 181: Edit DVLAN Interface Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
Interface Mode	Choose the behavior for double VLAN tagging by selecting the appropriate administrative mode 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).
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 values 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 values 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.

Figure 195: Edit DVLAN Interface Configuration

4.7 IP Subnet Based VLAN Status

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 196: IP Subnet Based VLAN Status

Table 182: 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to redisplay the most current information from the switch.
- a To add an IP subnet-based VLAN, click **Add** and specify an IP address, subnet mask, and VLAN ID in the available fields.

Figure 197: Add IP Subnet Based VLAN

Table 183: Add IP Subnet Based VLAN Fields

Field	Description
IP Address	Enter 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	Enter the subnet mask that defines the network portion of the IP address.
VLAN ID	Enter 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.

- a 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.
- a 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.

4.8 Protocol-Based VLAN Configuration

Protocol-based Virtual Local Area Networks (PBVLANS) is used to bridge traffic 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.

4.8.1 Protocol-Based VLAN Overview

Use this page to add and remove Protocol-based Virtual Local Area Networks (PBVLANS).

To display the Protocol Based VLAN Status page, click **Switching > Protocol Based VLAN > Status**.

The screenshot shows the 'Protocol Based VLAN Overview' page in the Configuration tab. At the top, there are tabs for 'Status' and 'Configuration'. Below the tabs is a header 'Protocol Based VLAN Overview' with a help icon. A table displays the configuration with columns: Group Name, VLAN, Protocol, and Interface. The table shows one entry: Group Name '1', VLAN '1', Protocol '0x600', and Interface '1/0/6'. Above the table, there are controls for 'Display All rows', 'Showing 1 to 1 of 1 entries', and a 'Filter' input field. Below the table, there are pagination controls: 'First', 'Previous', '1', 'Next', 'Last'. At the bottom, there are three buttons: 'Refresh', 'Add', and 'Remove'.

Figure 198: Protocol Based VLAN Overview

Table 184: Protocol Based VLAN Overview

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: - a If the frame received over a port is tagged, normal processing takes place. - a If the frame received over a port is untagged, the frame type is matched according to the protocols assigned to the group on that port. - a If a match is found, the frame is assigned the VLAN ID specified for the group. - a If a match is not found, the frame is assigned the port VID (PVID) as its VLAN ID.
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.

4 Configuring Switching Information

Field	Description
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to redisplay the most current information from the switch.
- a To add a PBVLAN, click **Add** and specify a group name, VLAN ID, protocol, and interfaces in the available fields.

The screenshot shows a window titled "Add Protocol Based VLAN Group". It has three input fields at the top: "Group Name" with a hint "(1 to 16 alphanumeric characters)", "VLAN" with a hint "(1 to 4093)", and "Protocol" with a hint "(600 to FFFF hex, or keyword) ?". Below these are two list boxes. The "Available Interfaces" list contains "1/0/1", "1/0/2", "1/0/3", "1/0/4", "1/0/5", "1/0/6", "1/0/7", and "1/0/8". The "Group Interfaces" list is currently empty. Between the two lists are two arrow buttons: a right-pointing arrow and a left-pointing arrow. At the bottom right of the window are "Submit" and "Cancel" buttons.

Figure 199: Add Protocol Based VLAN Group

Table 185: Add Protocol Based VLAN Group Fields

Field	Description
Group Name	Enter a name that identifies the PBVLAN group.
VLAN	Enter the VLAN ID to be used for the PBVLAN.
Available Interfaces	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.
Group Interfaces	The interfaces that are members of the PBVLAN group.

- a 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.

4.8.1.1 Adding a PBVLAN

To add a PBVLAN, click **Add** and specify a group name, VLAN ID, protocol, and interfaces in the available fields.

Figure 200: Add Protocol Based VLAN Group

Table 186: 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: - a If the frame received over a port is tagged, normal processing takes place. - a If the frame received over a port is untagged, the frame type is matched according to the protocols assigned to the group on that port. - a If a match is found, the frame is assigned the VLAN ID specified for the group. - a If a match is not found, the frame is assigned the port VID (PVID) as its VLAN ID.
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 PVBLAN, you can specify the EtherType hex value or the protocol keyword (for IP, ARP, and IPX).
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.

- a Click **Refresh** to display the latest information from the router.
- a 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.

4.8.2 Protocol Based VLAN Group 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 Group Configuration page, click **Switching > Protocol Based VLAN > Configuration**.

Figure 201: Protocol Based VLAN Group Configuration

Table 187: Protocol Based VLAN Group 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	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 PBVLAN Group Interfaces. When adding a protocol, you can specify the EtherType hex value or the protocol keyword (for IP, ARP, and IPX). To configure the protocols associated with a PBVLAN group, use the buttons available in the protocol table: - a To add a protocol to the group, click the + (plus) button and enter the protocol to add. - a To delete an entry from the list, click the – (minus) button associated with the entry to remove. - a To delete all entries from the list, click the – (minus) button in the heading row.
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.

Field	Description
Group Interfaces	The interfaces that are members of the PBVLAN group.

Use the buttons to perform the following tasks:

- a If you make any changes, click **Submit** to apply the change to the system.
- a Click **Refresh** to display the most current information from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.9 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.

4.9.1 Private VLAN Configuration

To access the Private VLAN Configuration page, click **Switching > Private VLAN > Configuration** in the navigation menu.

Figure 202: Private VLAN Configuration

- 4** Default VLAN and management VLAN cannot be configured as private VLANs and hence are not displayed on this page.

Table 188: Private VLAN Configuration Fields

Field	Description
VLAN ID	Displays the VLAN ID for which Private VLAN type is being set.
Type	Displays the configured Private VLAN Type. - a Unconfigured – The VLAN is not configured as a private VLAN. - a Primary – A private VLAN that forwards the traffic from the promiscuous ports to isolated ports, community ports, and other promiscuous ports in the

Field	Description
	same private VLAN. Only one primary VLAN can be configured per private VLAN. All ports within a private VLAN share the same primary VLAN. a Isolated – A secondary VLAN that carries traffic from isolated ports to promiscuous ports. Only one isolated VLAN can be configured per private VLAN. a 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the switch.
- a To add a VLAN, click **Add VLAN** and specify the VLAN IDs in the available field.


A dialog box titled "Add VLAN" with a close button (X) in the top right corner. It contains a text input field labeled "VLAN ID or Range" with the value "2". Below the input field is a small text box with instructions: "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." At the bottom right are "Submit" and "Cancel" buttons.

Figure 203: Add VLAN

Table 189: Add VLAN

Field	Description
VLAN ID or Range	Enter a VLAN ID or VLAN ID range for which Private VLAN type is being set.

- a To configure a private VLAN, select the entry to modify and click **Edit**. Then, configure the desired private VLAN setting.


A dialog box titled "Edit Private VLAN Configuration" with a close button (X) in the top right corner. It contains two fields: "VLAN ID" with the value "2" and "Type" with a dropdown menu showing "Unconfigured". At the bottom right are "Submit" and "Cancel" buttons.

Figure 204: Edit Private VLAN Configuration

Table 190: Edit Private VLAN Configuration

Field	Description
VLAN ID	Displays the VLAN ID for which Private VLAN type is being set.
Type	In the dropdown menu select the type of private VLAN. The factory default is Unconfigured . a Unconfigured – The VLAN is not configured as a private VLAN. a Primary – A private VLAN that forwards the traffic from the promiscuous ports to isolated ports, community ports, and other promiscuous ports in the same

Field	Description
	private VLAN. Only one primary VLAN can be configured per private VLAN. All ports within a private VLAN share the same primary VLAN. a Isolated – A secondary VLAN that carries traffic from isolated ports to promiscuous ports. Only one isolated VLAN can be configured per private VLAN. a 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.

- a Click **Edit All** to select the type for all configured private VLANs.

4.9.2 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 access the Private VLAN Association page, click **Switching > Private VLAN > Association** in the navigation menu.

Figure 205: Private VLAN Association

4 Isolated VLANs and Community VLANs are collectively called Secondary VLANs.

Table 191: 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the switch.

4 Configuring Switching Information

- a To configure a primary VLAN association with a secondary VLAN, select each entry to modify and click **Edit**.



Figure 206: Edit Private VLAN Association

Table 192: Edit Private VLAN Association Fields

Field	Description
Primary VLAN	The VLAN ID of the selected primary VLAN.
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.

4.9.3 Private VLAN Interface Association

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 Association page, click **Switching > Private VLAN > Interface** in the navigation menu.




Figure 207: Private VLAN Interface Association

Table 193: Private VLAN Interface Association 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: - a General – The interface is in general mode and is not a member of a private VLAN. - a 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). - a 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. - a 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. - a 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.
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.
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.
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.
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.

4 Configuring Switching Information

Field	Description
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the switch.
- a To configure the port mode and private VLAN-to-interface associations, select the entry to modify and click **Edit**.

The screenshot shows a window titled "Edit Private VLAN Interface Association". It contains a table of configuration fields for interface 1/0/1. The fields are: Interface (1/0/1), Mode (General), Host Primary VLAN, Host Secondary VLAN, Isolated Trunk Primary VLAN, Isolated Trunk Secondary VLAN, Promiscuous Primary VLAN, Promiscuous Secondary VLAN, Promiscuous Trunk Primary VLAN, Promiscuous Trunk Secondary VLAN, Trunk Native VLAN, and Trunk Allowed VLAN (20, 30). At the bottom are "Submit" and "Cancel" buttons.

Figure 208: Edit Private VLAN Interface Association

Table 194: Edit Private VLAN Interface Association 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: - a General - a Host - a Isolated Trunk - a Promiscuous - a Promiscuous Trunk
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.

Field	Description
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.

- a To apply the same settings to all interfaces, click **Edit All**.
- a 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.
- a 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.
- a 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.
- a 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.

4.10 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. QoS

based on 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 209: Voice VLAN Configuration

Table 195: Voice VLAN Configuration

Field	Description
Voice VLAN Admin Mode	Click Enable or Disable to administratively turn the Voice VLAN feature on or off for all ports. 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.

Use the buttons to perform the following tasks:

- a If you make any changes, click **Submit** to apply the change to the system.
- a Click **Refresh** to display the latest information from the router.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.10.1 Voice VLAN Interface Summary

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.

To display the Voice VLAN Interface Summary page, click **Switching** > **Voice VLAN** > **Interface Summary**.

Figure 210: Voice VLAN Interface Summary

Table 196: Voice VLAN Interface Summary 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: a Enabled – The port ignores the 802.1p priority value in the Ethernet frames it receives from connected devices. a Disabled – The port trusts the priority value in the received frame.
Voice VLAN Interface Mode	Indicates how an IP phone connected to the port should send voice traffic: a Disable – Operationally disables the Voice VLAN feature on the interface (default setting). a VLAN ID – Forward voice traffic in the specified voice VLAN. a Dot1p – Tag voice traffic with the specified 802.1p priority value. a None – Use the settings configured on the IP phone to send untagged voice traffic. a Untagged – Send untagged voice traffic.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the switch.
- a 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.

Table 197: Add Voice VLAN Fields

Field	Description
Interface	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.
CoS Override Mode	The Class of Service override mode: a Enabled – The port ignores the 802.1p priority value in the Ethernet frames it receives from connected devices. a Disabled – The port trusts the priority value in the received frame.

Field	Description
Voice VLAN Interface Mode	Indicates how an IP phone connected to the port should send voice traffic: - a Disable – Operationally disables the Voice VLAN feature on the interface (default setting). - a VLAN ID – Forward voice traffic in the specified voice VLAN. - a Dot1p – Tag voice traffic with the specified 802.1p priority value. - a None – Use the settings configured on the IP phone to send untagged voice traffic. - a Untagged – Send untagged voice traffic.
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.

Figure 211: Add Voice VLAN

- a To change the Voice VLAN settings, select the interface to modify and click **Edit**.
- a To remove the Voice VLAN configuration from one or more ports, select each entry to delete and click **Remove**.

4.11 Virtual Port Channel Global 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 (with the exception of a stack), 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.

4 This feature is only supported by the LANCOM XS-6128QF.

To display the Virtual Port Channel Global Configuration page, click **Switching > Virtual Port Channel > Global**.

Global Interface Configuration Statistics

Virtual Port Channel Global Configuration

Domain ID (1 to 255)

VPC Mode ☒ Disable ☐ Enable

Operational VPC Mode Disabled

VPC State Disabled

VPC MAC (XXXXXXXXXXXX)

Operational VPC MAC (XXXXXXXXXXXX)

VPC System Priority 32767 (1 to 65535)

Operational VPC System Priority

Self Role None

Local System MAC 00:AD:57:60:A2:4D

Keepalive Parameters

Keepalive Priority 100 (1 to 255)

Keepalive Timeout (Seconds) 5 (2 to 15)

Keepalive Mode ☒ Disable ☐ Enable

Peer

Domain ID

VPC MAC

Operational VPC MAC

VPC System Priority

Operational VPC System Priority

Peer Role None

System MAC

Peer Link

Port Channel None

Peer Link Status Down

Peer Link STP Mode Disabled

Configured VLANs None

Egress Tagging None

Peer Detection

Peer Detected False

Peer IP Address (0.0.0.0) (X.X.X.X)

Source IP Address (0.0.0.0) (X.X.X.X)

UDP Port 50000 (1 to 65535)

Peer Detection Mode ☒ Disable ☐ Enable

Tx Interval (msecs) 1000 (200 to 4000)

Operational Tx Interval (msecs)

Rx Timeout (msecs) 3500 (700 to 14000)

Operational Rx Timeout (msecs)

Submit Refresh Remove Cancel



Figure 212: Virtual Port Channel Global Configuration

Table 198: Virtual Port Channel Global 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: - a The VPC administrative mode is globally enabled. - a Peer links are configured. - a The Keepalive mode is enabled.
VPC State	The VPC state, which is one of the following: - a Disable – The VPC mode is not operational. - a Listen – The keepalive component does not advertise any packets. It listens for advertisements from a peer. - a Ready – The keepalive component starts sending periodic keepalive messages. - a 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.

4 Configuring Switching Information

Field	Description
	a 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.
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.
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	
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 to:	

Field	Description
	- a carry the keepalive messages between the two peer devices. - a carry the BPDUs and LACPDUs between the secondary and primary VPC devices. - a carry control messages like VPC member port related events, FDB/MFDB entries, and configuration details. - a 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.
Port Channel	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.  The Edit window opens and allows you to select an available port channel from the Port Channel menu. To reset the port channel to the default value, click the Reset icon.  The port channel cannot be changed or reset when the Operational VPC Mode is Enabled.
Peer Link Status	The operational status of the peer link, which is either Up or Down.
Peer Link STP Mode	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.
Configured VLANs	The VLAN ID of each VLAN in which the port channel participates.
Egress Tagging	The VLAN ID tags included in the frames transmitted from the port channel.
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.

Field	Description
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.

Use the buttons to perform the following tasks:

- a If you make any changes, click **Submit** to apply the change to the system.
- a Click **Refresh** to display the latest information from the switch.
- a To remove the configured VPC domain, click **Remove**. You must confirm the action before the domain is deleted.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.11.1 Virtual Port Channel 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 Interface Configuration page, click **Switching > Virtual Port Channel > Interface Configuration**.

Interface	VPC Identifier	Operational Mode	VPC Interface State	Configured VLANs
0/3/1	1	Disabled	Wait	1
0/3/2	2	Disabled	Wait	1

Figure 213: Virtual Port Channel Interface Configuration

Table 199: Virtual Port Channel Interface Configuration Fields

Field	Description
Interface	The interface designation 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. The following modes can occur: - a Enabled - a Disabled
VPC Interface State	The VPC interface state, which is one of the following: - a Disabled – VPC functionality is operationally disabled on the VPC interface. - a Wait – The port channel is waiting for VPC functionality to be enabled on a port channel on the peer device.

Field	Description
	- a 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. - a 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. - a Inactive – The links connected to the VPC member ports are down, but the VPC interface on the peer remains active.
Configured VLANs	The VLAN ID of each VLAN in which the port channel participates.

Use the buttons to perform the following tasks:

- a To configure a port channel as a VPC interface, click **Add** and configure the desired settings.

The screenshot shows a dialog box titled "Add Virtual Port Channel Interface Configuration". Inside the dialog, there are two input fields. The first is labeled "Interface" and has a dropdown menu showing "0/3/3". The second is labeled "VPC Identifier" and has a dropdown menu showing "3". At the bottom right of the dialog, there are two buttons: "Submit" and "Cancel".

Figure 214: Add VPC Interface Configuration

Table 200: 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.

4 Configuring Switching Information

- a To view additional details about a VPC interface, select the interface with the information to view and click **Details**.

Virtual Port Channel Interface Details

Interface	0/3/1
VPC Identifier	1
Operational Mode	Disabled
VPC Interface State	Disabled
Configured VLANs	1

Self Member

Showing 0 to 0 of 0 entries

Self Port	Status
Table is Empty	

First Previous Next Last

Peer Member

Showing 0 to 0 of 0 entries

Peer Port	Status
Table is Empty	

First Previous Next Last

Close

Figure 215: Virtual Port Channel Interface Details

Table 201: Virtual Port Channel Interface Details Fields

Field	Description
Interface	The interface designation 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. The following modes can occur: - a Enabled - a Disabled
VPC Interface State	The VPC interface state, which is one of the following: - a Disabled - a Wait - a Error - a Active - a Inactive
Configured VLANs	The VLAN ID of each VLAN in which the port channel participates.
Self Member	
The Self Member fields provide information about the VPC member ports on the local device.	
Self Port	The interface designation of each port that is a member of the port channel configured as a VPC interface.
Status	The operational status of the port.

Field	Description
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.

- a To remove the VPC functionality from one or more port channels, select each entry to change and click **Remove**.

4.11.2 Virtual Port Channel 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**.

Global	Interface Configuration	Statistics

Virtual Port Channel Statistics	
Keepalive	
Total Transmitted	0
Tx Successful	0
Tx Errors	0
Total Received	0
Rx Successful	0
Rx Errors	0
Timeout Counter	0
Peer Link	
Control Messages Transmitted	0
Control Messages Tx Errors	0
Control Messages Tx Timeout (Seconds)	0
Control Messages ACK Transmitted	0
Control Messages ACK Tx Errors	0
Control Messages Received	0
Data Messages Transmitted	0
Data Messages Tx Errors	0
Data Messages Tx Timeout (Seconds)	0
Data Messages Received	0
BPDUs Transmitted To Peer	0
BPDUs Tx Error	0
BPDUs Received From Peer	0
BPDUs Rx Error	0
LACPDUs Transmitted To Peer	0
LACPDUs Tx Error	0
LACPDUs Received From Peer	0
LACPDUs Rx Error	0

Refresh Clear Keepalive Counters Clear Peer Link Counters

Figure 216: Virtual Port Channel Statistics

Table 202: Virtual Port Channel Statistics Fields

Field	Description
Keepalive	
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.	

4 Configuring Switching Information

Field	Description
Total Transmitted	The total number of keepalive messages the local device has sent to the peer device.
Tx Successful	The number of keepalive messages that have been successfully transmitted from the local device.
Tx Errors	The number of keepalive 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 keepalive messages the local device has received from the peer device.
Rx Successful	The number of keepalive messages the local device has successfully received from the peer device.
Rx Errors	The number of keepalive messages the local device has received from the peer device that contained errors.
Timeout Counter	The number of times the keepalive timeout timer has expired.
Peer Link	
In addition to keepalive 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.
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.
BPDUs Transmitted To Peer	The number of BPDUs successfully sent to the peer device over the peer link.
BPDUs Tx Error	The number of errors encountered when sending BPUDs to the peer device.
BPDUs Received From Peer	The number of BPDUs successfully received from the peer device over the peer link.
BPDUs Rx Error	The number of errors encountered when receiving BPUDs from the peer device.
LACPDUs Transmitted To Peer	The number of LACPDUs successfully sent to the peer device over the peer link.

Field	Description
LACPDUs Tx Error	The number of errors encountered when sending LACPDUs to the peer device.
LACPDUs Received From Peer	The number of LACPDUs successfully received from the peer device over the peer link.
LACPDUs Rx Error	The number of errors encountered when receiving LACPDUs from the peer device.

Use the buttons to perform the following tasks:

- a To reload the page and view the most current information, click **Refresh**.
- a Click the button **Clear Keepalive Counters** to reset all keepalive message counters to 0.
- a Click the button **Clear Peer Link Counters** to reset all peer link message counters to 0.

4.12 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

4.12.1 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.

4 Configuring Switching Information

To access the Port Auto Recovery Configuration page, click **Switching > Auto Recovery > Configuration** in the navigation menu.

Configuration

Port Auto Recovery Configuration

Auto Recovery Components

All Components	☐
ARP Inspection	☐ Enabled ☒ Disabled
Authentication Manager	☐ Enabled ☒ Disabled
BPDU Guard	☐ Enabled ☒ Disabled
BPDU Rate Limit	☐ Enabled ☒ Disabled
Broadcast Storm Control	☐ Enabled ☒ Disabled
CoA Disable Host Port	☐ Enabled ☒ Disabled
Denial Of Service	☐ Enabled ☒ Disabled
DHCP Rate Limit	☐ Enabled ☒ Disabled
Keepalive	☐ Enabled ☒ Disabled
Link Flap	☐ Enabled ☒ Disabled
MAC Locking	☐ Enabled ☒ Disabled
Multicast Storm Control	☐ Enabled ☒ Disabled
UDLD	☐ Enabled ☒ Disabled
Unicast Storm Control	☐ Enabled ☒ Disabled

Auto Recovery Parameters

Recovery Time (Seconds) (30 to 86400, 300 = Default)

D-Disabled Interface Status

Display

All

 rows

Showing 0 to 0 of 0 entries

Filter:

Interface	Admin Mode	Port Status	Error Disable Reason	Auto Recovery Time Left (Seconds)
Table is Empty				
FirstPreviousNextLast				

Submit

Refresh

Cancel

Figure 217: Port Auto Recovery Configuration

Table 203: Port Auto Recovery Configuration Fields

Field	Description
Auto Recovery Components	
This field lists all the components that support the Auto Recovery feature. For each component, you can enable or disable Auto Recovery.	
5	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.
All Components	Enables or disables Auto Recovery for all components .
ARP Inspection	Enables or disables Auto Recovery for ARP Inspection .
Authentication Manager	Enables or disables Auto Recovery for Authentication Manager .
BPDU Guard	Enables or disables Auto Recovery for BPDU Guard .
BPDU Rate Limit	Enables or disables Auto Recovery for BPDU Rate Limit .
Broadcast Storm Control	Enables or disables Auto Recovery for Broadcast Storm Control .
CoA Disable Host Port	Enables or disables Auto Recovery for CoA Disable Host Port .
Denial Of Service	Enables or disables Auto Recovery for Denial Of Service .

Field	Description
DHCP Rate Limit	Enables or disables Auto Recovery for DHCP Rate Limit .
Keepalive	Enables or disables Auto Recovery for Keepalive .
Link Flap	Enables or disables Auto Recovery for Link Flap .
MAC Locking	Enables or disables Auto Recovery for MAC Locking .
Multicast Storm Control	Enables or disables Auto Recovery for Multicast Storm Control .
UDLD	Enables or disables Auto Recovery for UDLD .
Unicast Storm Control	Enables or disables Auto Recovery for Unicast Storm Control .
Auto Recovery Parameters	
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.
Error Disable Reason	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: - a ARP Inspection - a Authentication Manager - a BPDU Guard - a BPDU Storm - a Broadcast Storm - a CoA Disable Host Port - a Denial Of Service - a DHCP Rate Limit - a Keepalive - a Link Flap - a MAC Locking - a Multicast Storm - a UDLD - a Unicast Storm
Auto Recovery Time Left	When Auto Recovery is enabled and the interface is placed in diagnostic disabled state, then a recovery timer (in seconds) starts for that interface. When this timer expires, the device checks if the interface is in diagnostic disabled state. If yes, then the device enables the diagnostic disabled interface.

Use the buttons to perform the following tasks:

- a 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.

- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.13 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.

4.13.1 Static MAC Filter Configuration Summary

Use the Static MAC Filter Summary page to associate a MAC address with a VLAN and one or more source and/or destination ports

To access the Static MAC Filter Summary page, click **Switching > Filters > MAC Filters** in the navigation menu.

Figure 218: Static MAC Filter Configuration

Table 204: Static MAC Filter Summary 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: - a 00:00:00:00:00:00 - a 01:80:C2:00:00:00 to 01:80:C2:00:00:0F - a 01:80:C2:00:00:20 to 01:80:C2:00:00:21 - a FF:FF:FF:FF:FF:FF
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 Members	The ports 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.

Field	Description
Destination Members	The ports 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 Destination Members field. 5 The parameter Destination Members can only be configured on XS-6128QF switches.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a To add a filter, click **Add** and configure the filter criteria.

The screenshot shows the 'Add Static MAC Filter' window. It contains input fields for 'MAC Address' and 'VLAN ID'. Below these are two sets of lists: 'Available Port List' and 'Source Members' on the left, and 'Destination Members' and another 'Available Port List' on the right. Navigation arrows are placed between the lists. At the bottom right are 'Submit' and 'Cancel' buttons.

Figure 219: Add Static MAC Filter

Table 205: Add Static MAC Filter Fields

Field	Description
MAC Address	Enter a MAC address to be filtered. 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: - a 00:00:00:00:00:00 - a 01:80:C2:00:00:00 to 01:80:C2:00:00:0F - a 01:80:C2:00:00:20 to 01:80:C2:00:00:21 - a FF:FF:FF:FF:FF:FF
VLAN ID	Select the VLAN ID to be associated with the filter. The VLAN ID is used with the MAC address to fully identify the frames to filter. The VLAN ID menu only lists VLANs currently configured on the system.
Source Members	The ports 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

Field	Description
	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 Members	The ports 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 Destination Members field. 5 The parameter Destination Members can only be configured on XS-6128QF switches. Otherwise the error message MAC Filter destination port configuration is not available on this platform will be displayed.

- a To edit a filter, select the filter to update and click **Edit**. To change the MAC address or VLAN associated with a filter, you must remove and re-create the filter.
- a To remove a filter, select each entry to delete and click **Remove**. You must confirm the action for the MAC filter to be deleted.

4.14 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.

4.14.1 Global Configuration

Use the Global Configuration page to configure global DAI settings.

To display the Global Configuration page, click **Switching** > **Dynamic ARP Inspection** > **Global** in the navigation menu.

Figure 220: Global Configuration

Table 206: 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: - a 0.0.0.0 - a 255.255.255.255 - a All IP multicast addresses - a All class E addresses (240.0.0.0/4) - a Loopback addresses (in the range 127.0.0.0/8)


Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.14.2 Dynamic ARP Inspection VLAN Configuration

Use the Dynamic ARP Inspection VLAN Configuration page to select the DAI-capable VLANs for which information is to be displayed or configured.

To display the Dynamic ARP Inspection Configuration page, click **Switching > Dynamic ARP Inspection > VLAN** in the navigation menu.

**Figure 221: Dynamic ARP Inspection VLAN Configuration****Table 207: Dynamic ARP Inspection VLAN Configuration Fields**

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.

4 Configuring Switching Information

Field	Description
Log 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 ARP access control list (ACL) that the VLAN uses as the filter for ARP packet validation. The ARP ACL must have been already configured in the menu Dynamic ARP Inspection ACL Summary 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: - a 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. - a Disable – The ARP packet needs further validation by using the entries in the DHCP Snooping database.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To enable DAI on a VLAN and to configure the optional DAI settings, click **Add**.

Table 208: Add VLAN Fields

Field	Description
VLAN ID	Select the VLAN where DAI should be enabled. A VLAN does not need to exist on the system to be enabled for DAI.
Log Invalid Packets	Choose if DAI logging should be 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	Enter the name of the ARP access control list (ACL) that the VLAN uses as the filter for ARP packet validation. The ARP ACL must have been already configured in the menu Dynamic ARP Inspection ACL Summary 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: - a Enable - a Disable

Figure 222: Add VLAN

- a To change the DAI settings on VLAN, select the VLAN with the settings to update and click **Edit**.
- a 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.

4.14.3 Interface Configuration

Use the Interface Configuration page to select the DAI Interface for which information is to be displayed or configured.

To display the Interface Configuration page, click **Switching > Dynamic ARP Inspection > Interface** in the navigation menu.

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

Figure 223: Interface Configuration

Table 209: Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Trust State	Indicates whether the DAI feature should check traffic on the interface for possible ARP packet violations. This field has one of the following values: - a Enabled – The interface is trusted. ARP packets arriving on this interface are forwarded without DAI validation. - a Disabled – The interface is not trusted. ARP packets arriving on this interface are subjected to ARP inspection.
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.
Burst Interval	The number of consecutive seconds the interface is monitored for incoming ARP packet rate limit violations.

Use the buttons to perform the following tasks:

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

- a To configure DAI settings for an interface, select the entry to update and click **Edit**.

Interface	1/0/1
Trust State	☐
Rate Limiting	☒
Rate Limit	15 (0 to 300)
Burst Interval (Seconds)	1 (1 to 15)

Submit Cancel

Table 210: Edit Interface Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
Trust State	Indicates whether the DAI feature should check traffic on the interface for possible ARP packet violations. - a Check the checkbox to enable the trust state. ARP packets arriving on this interface are forwarded without DAI validation. - a Uncheck the checkbox to disable the trust state. ARP packets arriving on this interface are subjected to ARP inspection.
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.
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.
Burst Interval	The number of consecutive seconds the interface is monitored for incoming ARP packet rate limit violations.

Figure 224: Edit Interface Configuration

4.14.4 Dynamic ARP Inspection 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 Dynamic ARP Inspection ACL Configuration page, click **Switching > Dynamic ARP Inspection > ACL Summary** in the navigation menu.




Figure 225: Dynamic ARP Inspection ACL Summary

Table 211: 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:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To add an ARP ACL, click **Add** and configure the ACL name.




Figure 226: Add Access Control List

- a 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.
- a 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.

4.14.5 Dynamic ARP Inspection ACL Configuration

Use the Dynamic ARP Inspection ACL Configuration page to add rules or remove DAI ARP ACLs.

To display the Dynamic ARP Inspection ACL Configuration page, click **Switching > Dynamic ARP Inspection > ACL Configuration** in the navigation menu.




Figure 227: Dynamic ARP Inspection ACL Configuration

Table 212: Dynamic ARP Inspection ACL Configuration Fields

Field	Description
ACL Name	The menu contains the ARP ACL names that exist on the system.
Action	Action to be performed on a received ARP packet that matches both the Sender IP Address and Sender MAC Address values. A value of Permit will allow the matching ARP packet through and a value of Deny will drop the matching ARP packet.
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.
ACL Remarks	Lists the configured remarks for an ARP ACL. All remarks present in this table are applied to the next rule created with the Add Rule button.

Use the buttons to perform the following tasks:

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




Figure 228: Add ACL Rule

Table 213: Add ACL Rule

Field	Description
Sender IP Address	The IP address of a system that is sending the 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. A value of Any (0.0.0.0) will match on any Sender IP Address in the received ARP packet.
Sender MAC Address	The MAC address of a system that is sending the 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. A value of Any (00:00:00:00:00:00) will match on any Sender MAC Address in the received ARP packet.
Action	Action to be performed on a received ARP packet that matches both the Sender IP Address and Sender MAC Address values. A value of Permit will allow the matching ARP packet through and a value of Deny will drop the matching ARP packet.

- a To remove a configured rule for an ARP ACL, select the appropriate ACL from the ACL Name menu and click **Remove Rule**. You must confirm the action before the entry is deleted.

4.14.6 Dynamic ARP Inspection Statistics

Use the Dynamic ARP Inspection Statistics page to display the statistics per VLAN.

To display the Dynamic ARP Inspection Statistics page, click **Switching > Dynamic ARP Inspection > Statistics** in the navigation menu.

Figure 229: Dynamic ARP Inspection Statistics**Table 214: Dynamic ARP Inspection Statistics Fields**

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.

Field	Description
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: - a 0.0.0.0 - a 255.255.255.255 - a All IP multicast addresses - a All class E addresses (240.0.0.0/4) - a Loopback addresses (in the range 127.0.0.0/8)
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.

4.15 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.

4.15.1 GARP Switch Configuration

To access the GARP Switch Configuration page, click **Switching** > **GARP** > **Switch** in the navigation menu.

Figure 230: GARP Switch Configuration

Table 215: 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.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.15.2 GARP 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.

4 Configuring Switching Information

To access the GARP Port Configuration page, click **Switching > GARP > Port** in the navigation menu.

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

Figure 231: GARP Port Configuration

Table 216: GARP Port Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
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.

Use the buttons to perform the following tasks:

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

- a 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 217: Edit GARP Port Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
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.
GARP Timers	
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.

Figure 232: Edit GARP Port Configuration

4.16 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.

4.16.1 DHCP Snooping Configuration

Use this page to view and configure the global settings for DHCP Snooping.

To access the DHCP Snooping Configuration page, click **Switching > DHCP Snooping > Base > Global** in the navigation menu.




Figure 233: DHCP Snooping Configuration

Table 218: 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.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.16.2 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 234: DHCP Snooping VLAN Configuration

Table 219: 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 DHCP snooping mode for the VLAN. Only VLANs that are enabled for DHCP snooping appear in the list.

Use the buttons to perform the following tasks:

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




Figure 235: Add DHCP Snooping VLAN Configuration

- a 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.

4.16.3 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.

4 Configuring Switching Information

To access the DHCP Snooping Interface Configuration page, click **Switching > DHCP Snooping > Base > Interface Configuration** in the navigation menu.

Interface	Trust State	Log Invalid Packets	Rate Limit (pps)	Burst Interval (Seconds)
1/0/1	Disabled	Disabled		
1/0/2	Disabled	Disabled		
1/0/3	Disabled	Disabled		
1/0/4	Disabled	Disabled		
1/0/5	Disabled	Disabled		
1/0/6	Disabled	Disabled		
1/0/7	Disabled	Disabled		
1/0/8	Disabled	Disabled		
1/0/9	Disabled	Disabled		
1/0/10	Disabled	Disabled		

Figure 236: DHCP Snooping Interface Configuration

Table 220: DHCP Snooping Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Trust State	The trust state configured on the interface. The trust state is one of the following: - a 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: - a DHCP packets from a DHCP server (DHCP OFFER, DHCP ACK, DHCP NAK, DHCP RELEASE QUERY) are dropped. - a 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. - a DHCP packets are dropped when the source MAC address does not match the client hardware address if MAC Address Validation is globally enabled. - a Enabled – The interface is considered to be trusted and forwards DHCP server messages without validation.
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.

Use the buttons to perform the following tasks:

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

Edit DHCP Snooping Interface Configuration	
Interface	1/0/1
Trust State	☐
Log Invalid Packets	☐
Rate Limit (pps)	☐ (0 to 300)
Burst Interval (Seconds)	☐ (1 to 15)
Submit Cancel	

Table 221: Edit DHCP Snooping Interface Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
Trust State	The trust state configured on the interface. The trust state is one of the following: - a Check the checkbox to disable the trust state. 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: - a DHCP packets from a DHCP server (DHCP OFFER, DHCP ACK, DHCP NAK, DHCP RELEASE QUERY) are dropped. - a 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. - a DHCP packets are dropped when the source MAC address does not match the client hardware address if MAC Address Validation is globally enabled. - a Check the checkbox to enable the trust state. DHCP server messages are forwarded without validation.
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.

Figure 237: Edit DHCP Snooping Interface Configuration

4.16.4 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 > DCHP Snooping > Base > Static Bindings** in the navigation menu.




Figure 238: DHCP Snooping Static Bindings

Table 222: DHCP Snooping Static Bindings Fields

Field	Description
Interface	The interface where the DHCP client is authorized.
MAC Address	The MAC address associated with the client. This is the key to the binding database.
VLAN ID	The VLAN ID the client is authorized to use.
IP Address	The IP address of the client.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To add a static entry to the DHCP snooping bindings table, click **Add** and specify the desired settings.

Add DHCP Snooping Static Binding

Interface

1/0/1

MAC Address

(xx:xx:xx:xx:xx:xx)

VLAN ID

1

IP Address

(x.x.x.x)

Submit

Cancel

Table 223: Add DHCP Snooping Static Binding Fields

Field	Description
Interface	Select the interface where the DHCP client is to be authorized.
MAC Address	Enter the MAC address of the client to be authenticated. This is the key to the binding database.
VLAN ID	Enter the VLAN ID the client is authorized to use.

Field	Description
IP Address	Enter the IP address of the client.

Figure 239: Add DHCP Snooping Static Binding

- a To remove one or more static entries, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

4.16.5 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 240: DHCP Snooping Dynamic Bindings**Table 224: DHCP Snooping Dynamic Bindings Fields**

Field	Description
Interface	The interface where 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 DHCP lease time for the client.

Use the buttons to perform the following tasks:

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

4.16.6 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 switch must be able to reach the IP address of the remote system to send bindings to a remote database.

- 5** The binding database is sent to a remote system via TFTP. There the remote system must act as as TFTP server.

To access the DHCP Snooping Persistent Configuration page, click **Switching > DCHP Snooping > Base > Persistent** in the navigation menu.

Figure 241: DHCP Snooping Persistent Configuration

Table 225: 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 where 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.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.16.7 DHCP Snooping Statistics

Use this page to view and clear per-interface statistics about the DHCP messages filtered by the DHCP snooping feature. Only interfaces that are enabled for DHCP snooping and are untrusted appear in the table.

To access the DHCP Snooping Statistics page, click **Switching > DHCP Snooping > Base > Statistics** in the navigation menu.

Interface	MAC Verify Failures	Client Ifc Mismatch	DHCP Server Msgs Received
1/0/1	0	0	0
1/0/2	0	0	0
1/0/3	0	0	0
1/0/4	0	0	0
1/0/5	0	0	0
1/0/6	0	0	0
1/0/7	0	0	0
1/0/8	0	0	0
1/0/9	0	0	0
1/0/10	0	0	0

Figure 242: DHCP Snooping Statistics

Table 226: 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 ((DHCPOFFER, DHCPACK, DHCPNAK, DHCPRELEASEQUERY) that have been dropped on an untrusted port.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To reset the statistics to zero for one or more interfaces, select each interface with the data to reset and click the **Clear Counters** button. You must confirm the action before the entry is deleted.

4.16.8 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 243: DHCP L2 Relay Global Configuration

Table 227: DHCP L2 Relay Global Configuration

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.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.16.9 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 244: DHCP L2 Relay Interface Configuration

Table 228: DHCP L2 Relay Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
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: - a Enabled – 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. - a Disabled – 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.

Use the buttons to perform the following tasks:

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

Table 229: Edit DHCP L2 Relay Interface Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
L2 Relay Mode	Choose 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	Choose the L2 relay trust mode for the interface, which is one of the following: - a Enable - a Disable

Figure 245: Edit DHCP L2 Relay Interface Configuration

4.16.10 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), 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 246: DHCP L2 Relay VLAN Configuration

Table 230: DHCP L2 Relay VLAN Configuration

Field	Description
VLAN ID	The VLAN associated with the rest of the data in the row.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To enable a VLAN for DHCP L2 relay, click **Add** and select the VLAN from the available menu.

Figure 247: Add DHCP L2 Relay VLAN Configuration

Table 231: Add DHCP L2 Relay VLAN Configuration Fields

Field	Description
VLAN ID	Select the VLAN ID to be used in the DHCP L2 Relay.
Circuit ID	Choose 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 than to all ports in the VLAN).
Remote ID	Enter 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.

- a 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.
- a To disable one or more VLANs as DHCP L2 relay agents, select the appropriate VLANs and click **Remove**. You must confirm the action.

4.16.11 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.

Global

Interface Configuration

VLAN Configuration

Statistics

DHCP L2 Relay Interface Statistics

Display10rows

Showing 1 to 10 of 118 entries

Filter:

☐	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
☐	1/0/1	0	0	0	0
☐	1/0/2	0	0	0	0
☐	1/0/3	0	0	0	0
☐	1/0/4	0	0	0	0
☐	1/0/5	0	0	0	0
☐	1/0/6	0	0	0	0
☐	1/0/7	0	0	0	0
☐	1/0/8	0	0	0	0
☐	1/0/9	0	0	0	0
☐	1/0/10	0	0	0	0

First

Previous

1

2

3

4

5

Next

Last

Refresh

Clear Counters

Figure 248: DHCP L2 Relay Interface Statistics

Table 232: 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.
Untrusted 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.

Use the buttons to perform the following tasks:

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

4.16.12 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 access the DHCP Snooping IP Source Guard Interface Configuration page, click **Switching** > **DHCP Snooping** > **IP Source Guard** > **Interface Configuration** in the navigation menu.

Interface	IP Source Guard	Port Security
1/0/1	Disabled	Disabled
1/0/2	Disabled	Disabled
1/0/3	Disabled	Disabled
1/0/4	Disabled	Disabled
1/0/5	Disabled	Disabled
1/0/6	Disabled	Disabled
1/0/7	Disabled	Disabled
1/0/8	Disabled	Disabled
1/0/9	Disabled	Disabled
1/0/10	Disabled	Disabled

Figure 249: DHCP Snooping IP Source Guard Interface Configuration

Table 233: DHCP Snooping IP Source Guard Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
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.

Use the buttons to perform the following tasks:

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

Table 234: Edit DHCP Snooping IP Source Guard Interface Configuration Fields


Field	Description
Interface	Shows the selected interface(s).
IP Source Guard	Choose 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	Choose 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.

Figure 250: Edit-DHCP Snooping IP Source Guard Interface Configuration

4.16.13 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 251: DHCP Snooping IP Source Guard Bindings****Table 235: DHCP Snooping IP Source Guard Bindings Fields**

Field	Description
Interface	The interface where 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, which is one of the following: - a IP – Only the IP address configured in the binding is used as the source ID to filter IP packets. - a IP-MAC – Both the IP address and its associated MAC address are used to verify whether the IP packets are allowed on the interface. The MAC address is used for IP source enforcement when IPSG port security is enabled on the interface.

Field	Description
Binding Type	The binding type, which is either dynamically learned or statically configured by an administrator.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To add a static entry to the bindings database, click **Add** and specify the desired settings.




Figure 252: Add DHCP Snooping IP Source Guard Bindings

Table 236: Add DHCP Snooping IP Source Guard Bindings Fields

Field	Description
Interface	Select the interface where the sender IP address is to be authorized.
VLAN ID	Enter the VLAN ID to be used for the binding rule.
MAC Address	Enter the authorized sender MAC address for the binding rule.
IP Address	Enter the authorized source IP address for the binding rule.

- a 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.

4.17 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

unrequested 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.

4.17.1 IGMP Snooping 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 Global Configuration and Status page, click **Switching > IGMP Snooping > Configuration** in the navigation menu.

Figure 253: IGMP Snooping Global Configuration and Status

Table 237: IGMP Snooping Global Configuration and Status Fields

Field	Description
Admin Mode	Select Enable or Disable in the Admin Mode field and click Submit to turn the feature on or off. 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 IGMP Snooping Interface Configuration on page 286.
Data Frames Forwarded by the CPU	Shows the number of data frames forwarded by the CPU.

Use the buttons to perform the following tasks:

- a 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.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.17.2 IGMP Snooping 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.

Interface	Admin Mode	Group Membership Interval	Max Response Time	Multicast Router Expiration Time	Fast Leave Admin Mode
1/0/1	Disable	260	10	0	Disable
1/0/2	Disable	260	10	0	Disable
1/0/3	Disable	260	10	0	Disable
1/0/4	Disable	260	10	0	Disable
1/0/5	Disable	260	10	0	Disable
1/0/6	Disable	260	10	0	Disable
1/0/7	Disable	260	10	0	Disable
1/0/8	Disable	260	10	0	Disable
1/0/9	Disable	260	10	0	Disable
1/0/10	Disable	260	10	0	Disable

Figure 254: IGMP Snooping Interface Configuration

Table 238: IGMP Snooping Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring IGMP snooping settings, this field identifies the interface(s) that are being configured.
Admin Mode	The administrative mode of IGMP snooping on the interface. IGMP snooping must be enabled globally and on an interface for the interface to be able to snoop IGMP 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 IGMP 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 must 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 IGMP leave message for a multicast group without first sending out MAC-based general queries.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the screen and associated messages.

a Select an entry and click **Edit** to configure IGMP Snooping for a specific interface.

The screenshot shows a configuration window titled "Edit IGMP Snooping Interface Configuration". It contains a table with the following fields and values:

Interface	1/0/1
Admin Mode	☒ Disable ☐ Enable
Group Membership Interval (Seconds)	260 (2 to 3600)
Max Response Time (Seconds)	10 (1 to 25) Must be less than Group Membership Interval
Multicast Router Expiration Time (Seconds)	0 (0 to 3600)
Fast Leave Admin Mode	☒ Disable ☐ Enable

At the bottom right, there are "Submit" and "Cancel" buttons.

Figure 255: Edit IGMP Snooping Interface Configuration

Table 239: Edit IGMP Snooping Interface Configuration Fields

Fields	Description
Interface	Shows the selected physical or LAG interface(s).
Admin Mode	Select the interface mode for the selected interface for IGMP Snooping for the switch from the 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 between 1 and 25. The default value 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 a particular interface from the menu. The default is Disable .

If you make any changes on the page, click **Submit** to apply the new settings to the switch.

4.17.3 IGMP Snooping Source Specific Multicast

This page displays information about multicast groups discovered by snooping IGMPv3 reports.

To access the IGMP Snooping Source Specific Multicast page, click **Switching > IGMP Snooping > Source Specific Multicast** in the navigation menu.




Figure 256: IGMP Snooping Source Specific Multicast

Table 240: IGMP Snooping Source Specific Multicast Fields

Field	Description
VLAN ID	VLAN where the IGMPv3 report is received.
Group	The IPv4 multicast group address.
Interface	The interface where the IGMPv3 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.

4.17.4 IGMP Snooping 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 IGMP Snooping VLAN Status page, click **Switching > IGMP Snooping > VLAN Overview** in the navigation menu.




Figure 257: IGMP Snooping VLAN Overview

Table 241: IGMP Snooping VLAN Status

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 a 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 must be less than the Group Membership Interval.
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: - a Enabled – Only the first IGMP report from all hosts for a group IGMP report is forwarded to the multicast routers. - a Disabled – The device forwards all IGMP reports from all hosts in a multicast group to the multicast routers.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To enable IGMP snooping on a VLAN, click **Add** and configure the settings in the available fields.

Table 242: Add IGMP Snooping VLAN Configuration Fields

Field	Description
VLAN ID	Select the VLAN ID to be used for IGMP Snooping.
Fast Leave Admin Mode	Enable or disable Fast Leave for the VLAN.
Group Membership Interval (Seconds)	Specify the amount of time you want the switch to wait for a report for a particular group on a particular VLAN 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 (Seconds)	Specify the amount of time you want the switch to wait after sending a query on a VLAN because it did not receive a report for a particular group on that VLAN. Enter a value between 1 and 25. The default value is 10 seconds. The configured value must be less than the Group Membership Interval .
Multicast Router Expiration Time (Seconds)	Specify the amount of time the VLAN should wait to receive a query before it is removed from the list of VLANs with multicast routers attached. Enter a value between 0 and 3600 seconds. The default is 0 seconds.
Report Suppression Mode	Enable or disable the IGMPv1 and IGMPv2 report suppression mode.

Figure 258: Add IGMP Snooping VLAN Configuration

- a To change the IGMP snooping settings for an IGMP-snooping enabled VLAN, select the entry with the settings to change and click **Edit**.
- a 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.

4.17.5 IGMP Snooping 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 IGMP Snooping Multicast Router Configuration page to manually configure an interface as a static multicast router interface.

4 Configuring Switching Information

To access the IGMP Snooping Multicast Router Configuration page, click **Switching > IGMP Snooping > Multicast Router Configuration** in the navigation menu.




Figure 259: IGMP Snooping Multicast Router Configuration

Table 243: IGMP Snooping Multicast Router Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Multicast Router	Indicates whether the interface is enabled or disabled as a multicast router interface. The following options are available: - a Enabled – The port is a multicast router interface. - a Disabled – The port does not have a multicast router configured.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a To change the multicast router mode for one or more interfaces, select each entry to modify and click **Edit**.



Table 244: Edit IGMP Snooping Multicast Router Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
Multicast Router	Choose whether the interface is enabled or disabled as a multicast router interface. a Check the checkbox to enable the multicast router interface on the selected ports(s).

Field	Description
	a Uncheck the checkbox to disable the multicast router on the selected port(s).

Figure 260: Edit IGMP Snooping Multicast Router Configuration

4.17.6 IGMP Snooping 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 IGMP Snooping Multicast Router VLAN Status page, click **Switching** > **IGMP Snooping** > **Multicast Router VLAN Overview** in the navigation menu.

Figure 261: IGMP Snooping Multicast Router VLAN Overview

Table 245: IGMP Snooping Multicast Router VLAN Overview 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 VLAN ID configured as enabled for multicast routing on the associated interface.

Use the buttons as follows:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To disable all VLANs as multicast router interfaces for one or more physical ports or LAGs, select each entry to modify and click **Remove**.
- a To enable or disable specific VLANs as multicast router interfaces for a physical port or LAG, use the **Add** and **Edit** buttons. In this case you are redirected to the [IGMP Snooping Multicast Router VLAN Configuration page](#). A multicast router interface faces a multicast router or IGMP querier and receives multicast traffic.

4.17.7 IGMP Snooping 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 262: IGMP Snooping Multicast Router VLAN Configuration

Table 246: IGMP Snooping Multicast Router VLAN Configuration Fields

Field	Description
Interface	Select the port or LAG where a VLAN multicast routing interface should be enabled or disabled.
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.

Use the buttons to perform the following tasks:

- a If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- a Click **Refresh** to refresh the page with the most current data from the switch.

4.18 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 queries 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.

4.18.1 IGMP Snooping Querier Configuration

To access the IGMP Snooping Querier Configuration page, click **Switching > IGMP Snooping Querier > Configuration** in the navigation menu.

Figure 263: IGMP Snooping Querier Configuration

Table 247: 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 where the query is being sent.
IGMP Version	The IGMP protocol version used in periodic IGMP queries.
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.

Use the buttons to perform the following tasks:

- a If you make any changes to this page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.18.2 IGMP Snooping Querier 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 264: IGMP Snooping Querier VLAN Configuration

Table 248: IGMP Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN where the IGMP snooping querier is enabled. Only VLANs that have been configured on the system and are not already enabled for the IGMP snooping querier appear in the menu.
Querier Election Participation	The participation mode for the IGMP snooping querier election process: - a 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. - a 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.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To enable the IGMP snooping querier feature on a VLAN, click **Add** and specify the desired settings.

Figure 265: Add IGMP Snooping Querier VLAN Configuration

Table 249: Add IGMP Snooping Querier VLAN Configuration Fields


Field	Description
VLAN ID	Select the VLAN where the IGMP snooping querier is to be enabled. Only VLANs that have been configured on the system and are not already enabled for the IGMP snooping querier appear in the menu.
Querier Election Participation	Choose the participation mode for the IGMP snooping querier election process: - a Enabled - a Disabled
Querier VLAN IP Address	Optionally enter 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.

- a To change the IGMP snooping querier settings for a VLAN, select the entry to modify and click **Edit**.
- a 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.

4.18.3 IGMP Snooping Querier VLAN Overview

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 Overview page, click **Switching > IGMP Snooping Querier > VLAN Overview** in the navigation menu.

**Figure 266: IGMP Snooping Querier VLAN Overview****Table 250: IGMP Snooping Querier VLAN Overview 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: a 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.

Field	Description
	a 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. a 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.
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.

4.19 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.

4.19.1 MLD Snooping Configuration and Status

To access the MLD Snooping Configuration and Status page, click **Switching > MLD Snooping > Configuration** in the navigation menu.

Figure 267: MLD Snooping Configuration and Status

Table 251: MLD Snooping Configuration and Status Fields

Field	Description
MLD Snooping Admin Mode	The administrative mode of MLD snooping on the device. Select Enable or Disable for the MLD Snooping Admin Mode field and click Submit to turn the feature on or off.

Field	Description
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 where 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 where MLD snooping is administratively enabled.

Use the buttons to perform the following tasks:

- a If you make any changes to this page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.19.2 MLD Snooping Interface Configuration

Use this page to configure MLD snooping settings on specific interfaces.

To access the MLD Snooping Interface Configuration page, click **Switching > MLD Snooping > Interface Configuration** in the navigation menu.

Configuration	Interface Configuration	Source Specific Multicast	VLAN Overview	Multicast Router Configuration	Multicast Router VLAN Overview

MLD Snooping Interface Configuration

Display 10 rows Showing 1 to 10 of 118 entries Filter:

	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

First Previous 1 2 3 4 5 Next Last

Refresh Edit

Figure 268: MLD Snooping Interface Configuration

Table 252: MLD Snooping Interface Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
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.

4 Configuring Switching Information

Field	Description
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 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a 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.

The screenshot shows a window titled "Edit Interface Configuration". Inside, there is a table of configuration fields:

Interface	1/0/1
Admin Mode	☒ Disable ☐ Enable
Group Membership Interval (Seconds)	260 (2 to 3600)
Max Response Time (Seconds)	10 (1 to 65) Must be less than Group Membership Interval
Multicast Router Expiration Time (Seconds)	0 (0 to 3600)
Fast Leave Admin Mode	☒ Disable ☐ Enable

At the bottom right of the dialog are "Submit" and "Cancel" buttons.

Table 253: Edit Interface Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
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 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.

Figure 269: Edit Interface Configuration

4.19.3 MLD Snooping Source Specific Multicast

This page displays Source Specific Multicast (SSM) information learned by snooping MLDv2 reports. MLDv2 includes support for SSM, where 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 MLD Snooping Source Specific Multicast page, click **Switching > MLD Snooping > Source Specific Multicast** in the navigation menu.

Figure 270: MLD Snooping Source Specific Multicast

Table 254: MLD Snooping Source Specific Multicast Fields

Field	Description
VLAN ID	The VLAN where the MLDv2 report is received.
Group	The IPv6 multicast group address of the multicast group the host belongs to.
Interface	The interface where 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: - a Include – The receiver has expressed interest in receiving multicast traffic for the multicast group from the source or sources in the Source Address List. - a 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.
Source Address List	The source IPv6 address or addresses source filtering is requested for.

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

4.19.4 MLD Snooping 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 MLD Snooping VLAN Status page, click **Switching > MLD Snooping > VLAN Overview** in the navigation menu.




Figure 271: MLD Snooping VLAN Status

Table 255: MLD Snooping VLAN Status

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.
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 MLD Snooping 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To enable MLD snooping on a VLAN, click **Add** and configure the settings in the available fields.




Figure 272: Add MLD Snooping VLAN Configuration

Table 256: Add MLD Snooping VLAN Configuration Fields

Field	Description
VLAN ID	Select the VLAN ID to be used for MLD Snooping.
Group Membership Interval	Specify the amount of time you want the switch to wait for a report for a particular group on a particular VLAN 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 a VLAN because it did not receive a report for a particular group on that VLAN. Enter a value between 1 and 25. The default value is 10 seconds. The configured value must be less than the Group Membership Interval .
Multicast Router Expiration Time	Specify the amount of time the VLAN should wait to receive a query before it is removed from the list of VLANs with multicast routers attached. Enter a value between 0 and 3600 seconds. The default is 0 seconds.
Fast Leave Admin Mode	Enable or disable Fast Leave for the VLAN.

- a To change the MLD snooping settings for an MLD-snooping enabled VLAN, select the entry with the settings to change and click **Edit**.
- a 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.

4.19.5 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.

Configuration | Interface Configuration | Source Specific Multicast | VLAN Overview | **Multicast Router Configuration** | Multicast Router VLAN Overview

MLD Snooping Multicast Router Configuration

Display **10** rows | Showing 1 to 10 of 118 entries | Filter:

☐	Interface	Multicast Router
☐	1/0/1	Disabled
☐	1/0/2	Disabled
☐	1/0/3	Disabled
☐	1/0/4	Disabled
☐	1/0/5	Disabled
☐	1/0/6	Disabled
☐	1/0/7	Disabled
☐	1/0/8	Disabled
☐	1/0/9	Disabled
☐	1/0/10	Disabled

First Previous 1 2 3 4 5 Next Last

Figure 273: Multicast Router Configuration

Table 257: MLD Snooping Multicast Router Configuration Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Multicast Router	Indicates whether the interface is enabled or disabled as a multicast router interface.

Use the button to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To change the multicast router mode for one or more interfaces, select each entry to modify and click **Edit**.

Table 258: Edit Multicast Router Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
Multicast Router	Choose whether the interface should be enabled or disabled as a multicast router interface.

Figure 274: Edit Multicast Router Configuration

4.19.6 MLD Snooping Multicast Router VLAN Overview

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 MLD Snooping Multicast Router VLAN Status page, click **Switching > MLD Snooping > Multicast Router VLAN Status** in the navigation menu.

Figure 275: MLD Snooping Multicast Router VLAN Overview

Table 259: MLD Snooping Multicast Router VLAN Overview 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 each VLAN configured as enabled as a multicast router interface on the associated interface.

Use the buttons to perform the following tasks:

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

Figure 276: Add Multicast Router VLAN Configuration**Table 260: Add Multicast Router VLAN Configuration Fields**

Field	Description
Interface	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).

- a 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**.
- a 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.

4.20 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 queries 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.

4.20.1 MLD Snooping Querier Configuration

To access the MLD Snooping Querier Configuration page, click **Switching > MLD Snooping Querier > Configuration** in the navigation menu.

Figure 277: Configuring MLD Snooping Querier Configuration

Table 261: Configuring 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. Only MLDv1 is available.
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.

Use the buttons to perform the following tasks:

- a If you make any changes to this page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to refresh the page with the most current data from the switch.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.20.2 MLD Snooping Querier 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 278: MLD Snooping Querier VLAN Configuration

Table 262: MLD Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	The VLAN where the MLD snooping querier is enabled.
Querier Election Participation	The participation mode for the MLD snooping querier election process: - a 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. - a 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.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to refresh the page with the most current data from the switch.
- a To enable the MLD snooping querier feature on a VLAN, click **Add** and specify the desired settings.

Table 263: Add MLD Snooping Querier VLAN Configuration Fields

Field	Description
VLAN ID	Select the VLAN where the MLD snooping querier is to be enabled. When enabling the MLD snooping querier on a VLAN, use this menu to select

Field	Description
	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	Choose the behavior for the participation mode for the MLD snooping querier election process: - a Tick the checkbox to enable the participation of the MLD snooping querier on this VLAN 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. - a Untick the checkbox to disable the participation of the MLD snooping querier on this VLAN. If it sees other queriers of the same version in the VLAN, the snooping querier moves to the non-querier state and stops sending periodic queries.
Querier VLAN IPv6 Address	Enter 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.

Figure 279: Add MLD Snooping Querier VLAN Configuration

- a To change the MLD snooping querier settings for a VLAN, select the entry to modify and click **Edit**.
- a 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.

4.20.3 MLD Snooping Querier VLAN Overview

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 Overview** in the navigation menu.




Figure 280: MLD Snooping Querier VLAN Overview

Table 264: MLD Snooping Querier VLAN Overview 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: - a 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. - a 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. - a 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.
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.

4.21 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. When a port is added to a LAG as a static member, it neither transmits nor receives LACPDUs (Link Aggregation Control Protocol Data Unit).

4 If you configure the maximum number of dynamic port-channels (LAGs) that your platform supports, additional port-channels that you configure are automatically static.

4.21.1 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 treats the port-channel as if it were a single link.

4 Configuring Switching Information

To access the Port Channel Summary page, click **Switching > Port Channel / LAG > Summary** in the navigation menu.

Name	Type	Admin Mode	STP Mode	Link State	Link Trap	Local Preference Mode	Members	Active Ports	Load Balance
☐ ch1	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype
☐ ch2	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype
☐ ch3	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype
☐ ch4	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype
☐ ch5	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype
☐ ch6	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype
☐ ch7	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype
☐ ch8	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype
☐ ch9	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype
☐ ch10	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype

Figure 281: Port Channel Summary

Table 265: Port Channel Summary Fields

Field	Description
Name	A unique name to identify the port channel. Depending on the type of port channel, this name is automatically assigned by the system or can be configured by a system administrator.
Type	The type of port channel: - a 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. - a 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. 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 (LACP).
Admin Mode	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 .
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.
Local Preference Mode	The local preference mode for the port channel: a Enabled – Known unicast traffic that is destined for a LAG egresses only out of members (if it has any) of the LAG interface on the local unit. This ensures that the LAG-destined known unicast traffic does not cross the external stack

Field	Description
	link when the LAG has members on the local unit. Unknown unicast, broadcast and multicast traffic behavior remains unchanged. a Disabled – Known unicast traffic that is destined for a LAG may egress out of any of the member ports depending upon the traffic pattern and the configured LAG hashing algorithm for the LAG interface. It is possible that this traffic may egress out of a member port on another unit. In this case, the traffic has to cross the external stacking link, which results in unnecessary bandwidth utilization of the external stack link.
Members	Lists the ports that are members of the Port Channel, in Unit/Slot/Port notation. There can be a maximum of 8 ports assigned to a Port Channel.
Active Ports	Lists the ports that are actively participating members of this Port Channel, in Unit/Slot/Port notation.
Load Balance	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: - a Source MAC, VLAN, EtherType, Incoming Port - a Destination MAC, VLAN, EtherType, Incoming Port - a Source/Destination MAC, VLAN, EtherType, Incoming Port - a Source IP and Source TCP/UDP Port Fields - a Destination IP and Destination TCP/UDP Port Fields - a Source/Destination IP and TCP/UDP Port Fields - a Enhanced Hashing Mode – In contrast to the other modes the Enhanced Hashing Mode uses more criteria for loadbalancing thus leading to a more efficient data transfer. 5 The option Enhanced Hashing Mode is only available on XS-6128QF switches.

Use the buttons to perform the following tasks:

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

4 Configuring Switching Information

- a To change the existing port channels select the entry with the settings to change and click **Edit**.

The screenshot shows the 'Edit Existing Port Channel' configuration window. The 'Port Channel Name' is 'ch1'. The 'Admin Mode', 'STP Mode', 'Static Mode', and 'Link Trap' are all set to 'Enable'. The 'Local Preference Mode' is set to 'Disable'. The 'Load Balance' dropdown is set to 'Source/Destination MAC, VLAN, Ethertype, Incoming Port'. Below these settings are two lists: 'Port List' and 'Members'. The 'Port List' contains ports 1/0/1 through 1/0/8. The 'Members' list is empty. At the bottom right are 'Submit' and 'Cancel' buttons.

Figure 282: Port Channel Configuration

Table 266: Port Channel Configuration Fields

Field	Description
Port Channel Name	You can either use the preconfigured name for the Port Channel or change it according to your needs. You may enter any string of up to 15 alphanumeric characters. You must specify a valid name to create the Port Channel.
Admin Mode	Select Enable or Disable from the 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	Select the Spanning Tree Protocol (STP) Administrative Mode associated with the Port Channel: - a Disable: Spanning tree is disabled for this Port Channel. - a Enable: Spanning tree is enabled for this Port Channel.
Static Mode	Select Enable or Disable from the menu. The factory default is Enable . - a Disable: The port channel is dynamically maintained (LACP). The interface transmits and processes LAGPDUs and requires a partner system. - a 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.
Link Trap	Specify whether you want to have a trap sent when link status changes. When Enable is selected, a trap will be sent after a link status change.
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

Field	Description
	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: - a Source MAC, VLAN, EtherType, and source port - a Destination MAC, VLAN, EtherType and source port - a Source/Destination MAC, VLAN, EtherType, and source port - a Source IP and Source TCP/UDP Port - a Destination IP and Destination TCP/UDP Port - a Source/Destination IP and source/destination TCP/UDP Port - a Enhanced hashing mode 5 The option Enhanced Hashing Mode is only available on XS-6128QF switches.
Port List	Contains the available physical ports to select for an LAG.
Members	The ports that are members of a port channel. Each port channel can have a maximum of 8 member ports. To add ports to the port channel, select one or more ports from the Port List field (CTRL + click to select multiple ports). Then, use the appropriate arrow icon to move the selected ports to the Members field.

4.21.2 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 / LAG > Statistics** in the navigation menu.

Summary

Statistics

Port Channel Statistics

?

Display

10

rows

Showing 1 to 10 of 64 entries

Filter:

	Interface	Channel Name	Type	Flap Count
☐	0/3/1	ch1	Port Channel	0
☐	0/3/2	ch2	Port Channel	0
☐	0/3/3	ch3	Port Channel	0
☐	0/3/4	ch4	Port Channel	0
☐	0/3/5	ch5	Port Channel	0
☐	0/3/6	ch6	Port Channel	0
☐	0/3/7	ch7	Port Channel	0
☐	0/3/8	ch8	Port Channel	0
☐	0/3/9	ch9	Port Channel	0
☐	0/3/10	ch10	Port Channel	0

First

Previous

1

2

3

4

5

Next

Last

Refresh

Clear Counters

Figure 283: Port Channel Statistics

Table 267: 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the router.
- a Click the **Clear Counters** button to reset the flap counters for all port channels and member ports to 0.

4.22 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.

4.22.1 Multicast Forwarding Database Summary

Use the Multicast Forwarding Database Summary 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 Multicast Forwarding Database Summary page, click **Switching > Multicast Forwarding Database > Summary** in the navigation menu.




Figure 284: Multicast Forwarding Database Summary Table

Table 268: Multicast Forwarding Database Summary Fields

Field	Description
VLAN ID	The VLAN ID (the first two groups of hexadecimal digits) associated with the entry in the MFDB.
MAC Address	The 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: - a 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. - a 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. - a 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. - a Static Filtering – A static MAC filter that was manually added to the address table by an administrator.
Type	The type of entry, which is one of the following: - a Static – The entry has been manually added to the MFDB by an administrator. - a Dynamic – The entry has been added to the MFDB as a result of a learning process or protocol.
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.
Forwarding Interface(s)	The list of forwarding interfaces. This list does not include any interfaces that are listed as static filtering interfaces.

- a Click **Refresh** to update the information on the screen with the most current data.

4.22.2 Multicast Forwarding Database GMRP Table

Use the Multicast Forwarding Database 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 Multicast Forwarding Database Table page, click **Switching > Multicast Forwarding Database > GMRP** in the navigation menu.




Figure 285: Multicast Forwarding Database GMRP Table

Table 269: Multicast Forwarding Database GMRP Table 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: - a Static – The entry has been manually added to the MFDB by an administrator. - a 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.
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.

4.22.3 Multicast Forwarding Database 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 Multicast Forwarding Database IGMP Snooping Table page, click **Switching > Multicast Forwarding Database > IGMP Snooping** in the navigation menu.




Figure 286: Multicast Forwarding Database IGMP Snooping Table

Table 270: Multicast Forwarding Database IGMP Snooping Table 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: - a Static – The entry has been manually added to the MFDB by an administrator. - a 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.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a To remove all IGMP snooping entries from the MFDB table, click **Clear Entries**. The table is repopulated as new addresses are discovered by the IGMP snooping feature.

4.22.4 MFDB Source Specific Multicast

This page displays the entries in the multicast forwarding database (MFDB) for source specific multicast. They were added because they were discovered by the IGMP Snooping or MLD Snooping feature.

To access the MFDB Source Specific Multicast page, click **Switching > Multicast Forwarding Database > Source Specific Multicast** in the navigation menu.




Figure 287: MFDB Source Specific Multicast

Table 271: MFDB Source Specific Multicast Fields

Field	Description
Type	Type of snooping. The values can be either IGMP Snooping or MLD Snooping.
VLAN ID	VLAN where 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 where an incoming packet is forwarded.

Click **Refresh** to update the information on the screen with the most current data.

4.22.5 Multicast Forwarding Database Statistics Source Specific Multicast Status

This page displays the entries in the multicast forwarding database (MFDB) for source specific multicast. They were added because they were discovered by the IGMP Snooping or MLD Snooping feature.

To access the MFDB Source Specific Multicast Status page, click **Switching > Multicast Forwarding Database > Source Specific Multicast Status** in the navigation menu.




Figure 288: MFDB Source Specific Multicast Status

Table 272: MFDB 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.
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.

4.22.6 Multicast Forwarding Database Statistics

Use the Multicast Forwarding Database Statistics page to view statistical information about the MFDB table.

To access the Statistics page, click **Switching > Multicast Forwarding Database > Statistics** in the navigation menu.

Summary

GMRP

IGMP Snooping

Source Specific Multicast

Source Specific Multicast Status

Statistics

Multicast Forwarding Database Statistics

?

MFDB Max Table Entries	2048
MFDB Most Entries Since Last Reset	0
MFDB Current Entries	0

Refresh

Figure 289: Multicast Forwarding Database Statistics**Table 273: 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 rebooted. 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.

4.23 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.

4.23.1 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 290: MVR Global Configuration

Table 274: 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: - a 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. - a Dynamic – MVR learns source ports membership from IGMP queries. MVR forwards the IGMP Joins and Leaves from the hosts to the router. 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.

Use the buttons to perform the following tasks:

- a If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.23.2 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 291: MVR Group Status

Table 275: MVR Group Status Fields

Field	Description
Group	The multicast group address.
Status	The status of the group, which can be one of the following: - a Active – Group has one or more MVR ports participating. - a Inactive – Group has no MVR ports participating.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a To add a group, click **Add** and specify a group address in the available field.




Figure 292: Add MVR Group

Table 276: Add MVR Group Fields

Field	Description
Group	Enter the multicast group address.

Field	Description
Contiguous Group Count	Specify the desired number of groups to be created starting with the entered group address. The default contiguous group count is 1.

- a To edit a configured group, select the entry to modify and click **Edit**. Then, configure which interfaces should be members of that group.




Figure 293: Edit MVR Group

Table 277: Edit MVR Group Fields

Field	Description
Group	The multicast group address.
Available Interfaces	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	The interfaces that are members of the MVR group.

- a 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.

4.23.3 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**.

Interface	MVR Interface Mode	Type	Status	Immediate Leave
1/0/1	Disabled	None	Active	Disabled
1/0/2	Disabled	None	Inactive	Disabled
1/0/3	Disabled	None	Inactive	Disabled
1/0/4	Disabled	None	Active	Disabled
1/0/5	Disabled	None	Inactive	Disabled
1/0/6	Disabled	None	Inactive	Disabled
1/0/7	Disabled	None	Inactive	Disabled
1/0/8	Disabled	None	Inactive	Disabled
1/0/9	Disabled	None	Inactive	Disabled
1/0/10	Disabled	None	Inactive	Disabled

Figure 294: MVR Interface Status

Table 278: MVR Interface Status Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
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: - a Source – The port where multicast traffic is flowing to. It must be a member of the multicast VLAN. - a Receiver – The port where listening host is connected to the switch. It must not be a member of the multicast VLAN. - a None – The port is not an MVR port.
Status	The active state of the interface, which can be one of the following: - a Active – The link for the port is up and it is in the forwarding state. - a Inactive – The link for the port may not be up, it may not be in the forwarding state, or both.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.

4 Configuring Switching Information

- a 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.

Edit MVR Port Configuration

Interface

1/0/1

MVR Interface Mode

☐ Enable

☒ Disable

Type

☐ Source

☐ Receiver

☒ None

Immediate Leave

☐ Enable

☐ Disable

Submit

Cancel


Table 279: Edit MVR Port Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
MVR Interface Mode	Choose 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	Choose the type of interface, which can be one of the following: - a Source – The port where multicast traffic is flowing to. It must be a member of the multicast VLAN. - a Receiver – The port where listening host is connected to the switch. It must not be a member of the multicast VLAN. - a None – The port is not an MVR port.
Immediate Leave	Choose 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.

Figure 295: Edit MVR Port Configuration

4.23.4 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**.



Statistics	Transmit	Receive
IGMP Queries	0	0
IGMPv1 Reports	0	0
IGMPv2 Reports	0	0
IGMP Leaves	0	0
Packet Failures	0	0

Refresh

Figure 296: MVR Statistics

Table 280: 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.

4.24 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 297: Protected Ports Configuration

Table 281: 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a To create a protected ports group and add ports to the group, click **Add** and configure the settings in the available fields.




Table 282: Add Group Fields

Field	Description
Group Name	Enter the name of the protected ports group.
Protected Ports	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 ports to the desired field.

Figure 298: Add Group

- a To change the name or the port members for an existing group, select the group to update and click **Edit**.
- a 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.

4.25 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.

4 This feature is only supported by the LANCOM XS-6128QF.

4.25.1 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.

Interface	Admin Mode	Status	Compatible Configuration Count	Incompatible Configuration Count	Peer Configuration Compatible	No-Drop Priorities
1/0/1	Disabled	Inactive	0	0	N/A	
1/0/2	Disabled	Inactive	0	0	N/A	
1/0/3	Disabled	Inactive	0	0	N/A	
1/0/4	Disabled	Inactive	0	0	N/A	
1/0/5	Disabled	Inactive	0	0	N/A	
1/0/6	Disabled	Inactive	0	0	N/A	
1/0/7	Disabled	Inactive	0	0	N/A	
1/0/8	Disabled	Inactive	0	0	N/A	
1/0/9	Disabled	Inactive	0	0	N/A	
1/0/10	Disabled	Inactive	0	0	N/A	

Figure 299: Priority Flow Control Configuration

Table 283: Priority Flow Control Configuration Fields

Field	Description
Interface	The physical port associated with the rest of the data in the row.
Admin Mode	The administrative mode of PFC on the interface: a 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.

4 Configuring Switching Information

Field	Description
	a 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.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a 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.

Priority Flow Control Action Mapping

You can configure any 2 priorities for No-Drop.

Interface	1/0/1
Admin Mode	☒ Disable ☐ Enable
Priority 0 Action	☒ Drop ☐ No-Drop
Priority 1 Action	☒ Drop ☐ No-Drop
Priority 2 Action	☒ Drop ☐ No-Drop
Priority 3 Action	☒ Drop ☐ No-Drop
Priority 4 Action	☒ Drop ☐ No-Drop
Priority 5 Action	☒ Drop ☐ No-Drop
Priority 6 Action	☒ Drop ☐ No-Drop
Priority 7 Action	☒ Drop ☐ No-Drop

Submit Cancel

Figure 300: Edit Priority Flow Control Action Mapping

Table 284: Edit Priority Flow Control Action Mapping Field

Field	Description
Interface	When configuring one or more ports, the Interface field in the Priority Flow Control Mapping Action window identifies the ports that are being configured.
Action Mode	Choose the administrative mode of PFC on the interface: a 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.

Field	Description
	a 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.
Priority Action (0 – 7)	Choose 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.

- a To apply the same PFC settings to all interfaces, click **Edit All** and configure the desired settings.

4.25.2 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.

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.

Figure 301: Priority Flow Control Statistics

Table 285: 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a Click the **Clear** button to reset all PFC statistics values for one or more interfaces to the default values.

4.26 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 [Spanning Tree CST Port Summary](#) on page 333.

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. An MSTP bridge can be configured to behave entirely as a RSTP bridge or an STP bridge.

- 4 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.

4.26.1 Spanning Tree Switch Configuration

The Spanning Tree Switch Configuration page contains fields for enabling STP on the switch.

To display the Spanning Tree Switch Configuration page, click **Switching** > **Spanning Tree** > **Switch** in the navigation menu.

Figure 302: Spanning Tree Switch Configuration

Table 286: 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.

Field	Description
Force Protocol Version	The STP version the device uses, which is one of the following: - a IEEE 802.1d – Classic STP provides a single path between end stations, avoiding and eliminating loops. - a 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. - a 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.
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.

Use the buttons to perform the following tasks:

- a If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.26.2 Spanning Tree CST Configuration

Use the Spanning Tree 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.

4 Configuring Switching Information

To display the Spanning Tree CST Configuration page, click **Switching > Spanning Tree > CST** in the navigation menu.

Spanning Tree CST Configuration	
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	☐
BPDUs Filter	☐
Spanning Tree Tx Hold Count	6 (1 to 10)
Bridge Identifier	80:00:00:40:C7:1D:18:C4
Time Since Topology Change	7d:21:29:15
Topology Change Count	0
Topology Change	False
Designated Root	80:00:00:40:C7:1D:18:C4
Root Path Cost	0
Root Port	00:00
Max Age	20
Forward Delay	15
Hold Time	6
CST Regional Root	80:00:00:40:C7:1D:18:C4
CST Path Cost	0

Submit Refresh Cancel

Figure 303: Spanning Tree CST Configuration

Table 287: Spanning Tree CST Configuration 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 in seconds a bridge waits before implementing a topological change.
Bridge Hello Time	The amount of time in seconds the root bridge waits between sending hello BPDUs. This is hardcoded to 2 seconds.
Bridge Forward Delay	The amount of time in seconds 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	When enabled, BPDU 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.

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. It is displayed in the format days:hours:minutes:seconds .
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 in seconds 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 in seconds 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.

Use the buttons to perform the following tasks:

- a If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- a Click **Refresh** to update the screen with most recent data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.26.3 Spanning Tree CST Port Summary

Use the Spanning Tree CST Port Summary page to view and configure the Common Spanning Tree (CST) settings for each interface on the device.

4 Configuring Switching Information

To display the Spanning Tree CST Port Summary page, click **Switching > Spanning Tree > CST Port** in the navigation menu.

Interface	Port Role	Port Forwarding State	Port Priority	Port Path Cost	Description
1/0/1	Designated	Forwarding	0x0080	20000	
1/0/2	Disabled	Disabled	0x0080	0	
1/0/3	Disabled	Disabled	0x0080	0	
1/0/4	Designated	Forwarding	0x0080	20000	
1/0/5	Disabled	Disabled	0x0080	0	
1/0/6	Disabled	Disabled	0x0080	0	
1/0/7	Disabled	Disabled	0x0080	0	
1/0/8	Disabled	Disabled	0x0080	0	
1/0/9	Disabled	Disabled	0x0080	0	
1/0/10	Disabled	Disabled	0x0080	0	

Figure 304: Spanning Tree CST Port Summary

Table 288: Spanning Tree CST Port Summary Fields

Field	Description
Interface	The port or link aggregation group (LAG) associated with the rest of the data in the row.
Port Role	The role of the port within the CST, which is one of the following: - a Root – A port on the non-root bridge that has the least-cost path to the root bridge. - a Designated – A port that has the least-cost path to the root bridge on its segment. - a Alternate – A blocked port that has an alternate path to the root bridge. - a Backup – A blocked port that has a redundant path to the same network segment as another port on the bridge. - a Master – The port on a bridge within an MST instance that links the MST instance to other STP regions. - a Disabled – The port is administratively disabled and is not part of the spanning tree.
Port Forwarding State	- a 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. - a 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. - a 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. - a Forwarding – The port sends and receives user traffic. - a Disabled – The port is administratively disabled and is not part of the spanning tree.

Field	Description
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a 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**.

Figure 305: Edit CST Port Entry

Table 289: Edit CST Port Entry Fields

Field	Description
Interface	Shows the selected interface(s).
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

4 Configuring Switching Information

Field	Description
	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.
Admin Edge Port	Select this option to 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.
Port Path Cost	The path cost from the port to the root bridge.
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 in seconds 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 Filter	When enabled, this feature filters the BPDU traffic on the edge ports. Edge ports do not need to participate in the spanning tree, so BPDU filtering allows BPDU packets received on edge ports to be dropped.
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.
Port Up Time Since Counters Last Cleared	The amount of time that the port has been up since the counters were cleared. It is displayed in the format days:hours:minutes:seconds .
Port Mode	The administrative mode of spanning tree on the port.
Port Forwarding State	- a Blocking - a Listening - a Learning - a Forwarding - a Disabled
Port Role	The role of the port within the CST, which is one of the following: - a Root - a Designated - a Alternate - a Backup

Field	Description
	a Master a Disabled
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 (Topology Change Notification) 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.

- a Click the **Details** button to view additional information for the selected interface in the CST topology.

4.26.4 Spanning Tree MST Summary

Use the Spanning Tree MST Summary 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

4 Configuring Switching Information

with a better distribution of network traffic and an increase in available bandwidth when compared to classic STP.

To display the Spanning Tree MST Summary page, click **Switching > Spanning Tree > MST** in the navigation menu.




Figure 306: Spanning Tree MST Summary


Table 290: 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. It is displayed in the format days:hours:minutes:seconds .
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the screen with most recent data.

- a To configure a new MSTI, click **Add** and specify the desired settings.



The 'Add MST Entry' dialog box contains the following fields and controls:


- MST ID:** A text input field with a value of 1 and a range indicator '(1 to 4094)'.
- Priority:** A text input field with a value of 8000 and a range indicator '(0 to F000 hex)'.
- # of Associated VLANs:** A label for the VLAN selection area.
- Available VLANs:** A list box containing VLANs 1 through 8.
- Members:** An empty list box for associated VLANs.
- Navigation:** Two arrow buttons (right and left) between the Available VLANs and Members list boxes.
- Buttons:** 'Submit' and 'Cancel' buttons at the bottom right.

Figure 307: Add MST Entry

Table 291: Add MST Entry 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.
Available VLANs	Lists the VLAN ID of each VLAN that can be associated with the MSTI so that the VLAN(s) are no longer associated with the common and internal spanning tree. To move a VLAN between the Available VLANs and Members fields, click the VLAN (or CTRL + click to select multiple VLANs), and then click the left or right arrow to move the selected VLAN(s) to the appropriate field.
Members	Lists the VLAN ID of each VLAN associated with the MSTI. The number of VLANs equals the value in # of Associated VLANs .

- a To change the Priority or the VLAN associations for an existing MSTI, select the entry to modify and click **Edit**.



The 'Edit MST Entry' dialog box contains the following fields and controls:

- MST ID:** A text input field with a value of 5.
- Priority:** A text input field with a value of 8000 and a range indicator '(0 to F000 hex)'.
- # of Associated VLANs:** A label for the VLAN selection area.
- Available VLANs:** A list box containing VLANs 2 through 9.
- Members:** A list box containing the value 1.
- Navigation:** Two arrow buttons (right and left) between the Available VLANs and Members list boxes.
- Bridge Information:** A section with the following fields:
 - Bridge Identifier:** 80:05:00:A0:57:53:6F:C5
 - Time Since Topology Change:** 0d:05:56:10
 - Topology Change Count:** 0
 - Topology Change:** False
 - Designated Root:** 80:05:00:A0:57:53:6F:C5
 - Root Path Cost:** 0
 - Root Port:** 00:00
- Buttons:** 'Submit' and 'Cancel' buttons at the bottom right.

Figure 308: Edit MST Entry

Table 292: Edit MST Entry Fields

Field	Description
MST ID	The number that identifies the MST instance. It cannot be changed when editing an existing entry.
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.
Available VLANs	Lists the VLAN ID of each VLAN that can be associated with the MSTI so that the VLAN(s) are no longer associated with the common and internal spanning tree. To move a VLAN between the Available VLANs and Members fields, click the VLAN (or CTRL + click to select multiple VLANs), and then click the left or right arrow to move the selected VLAN(s) to the appropriate field.
Members	Lists the VLAN ID of each VLAN associated with the MSTI. The number of VLANs equals the value in # of Associated VLANs .
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. It is displayed in the format days:hours:minutes:seconds .
Topology Change Count	The number of times the topology of the MSTI has changed.
Topology Change	Indicates whether a topology change is in progress on any port assigned to the MSTI. 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 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.

- a To remove one or more MSTIs, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

4.26.5 Spanning Tree MST Port Summary

Use this page to view and configure the Multiple Spanning Tree (MST) settings for each interface on the device.

To display the Spanning Tree MST Port Summary page, click **Switching > Spanning Tree > MST Port** in the navigation menu.

- 4 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 [Spanning Tree MST Port Summary](#).




Figure 309: Spanning Tree MST Port Summary

Table 293: Spanning Tree MST Port Summary 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: - a Root – A port on the non-root bridge that has the least-cost path to the root bridge. - a Designated – A port that has the least-cost path to the root bridge on its segment. - a Alternate – A blocked port that has an alternate path to the root bridge. - a Backup – A blocked port that has a redundant path to the same network segment as another port on the bridge. - a Master – The port on a bridge within an MST instance that links the MST instance to other STP regions. - a Disabled – The port is administratively disabled and is not part of the spanning tree.
Port Forwarding State	- a 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. - a 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. - a 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. - a Forwarding – The port sends and receives user traffic. - a Disabled – The port is administratively disabled and is not part of the spanning tree.

4 Configuring Switching Information

Field	Description
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a 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**.

Edit MST Port Entry

MST ID

5

Interface

1/0/1

Port Priority

80

(0 to F0 hex)

Port Path Cost

0

(0 to 200000000)

Auto-calculate Port Path Cost

Enabled

Port ID

80:01

Port Up Time Since Counters Last Cleared

0d:00:31:37

Port Mode

Disable

Enable

Port Forwarding State

Disabled

Port Role

Disabled

Designated Root

80:05:00:A0:57:53:6F:C5

Designated Cost

0

Designated Bridge

80:05:00:A0:57:53:6F:C5

Designated Port

00:00

Loop Inconsistent State

False

Transitions Into LoopInconsistent State

0

Transitions Out Of LoopInconsistent State

0

Submit

Cancel

Figure 310: Edit MST Port Entry

Table 294: Edit MST Port Entry 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 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.

Field	Description
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.
Port Up Time Since Counters Last Cleared	The amount of time that the port has been up since the counters were cleared. It is displayed in the format days:hours:minutes:seconds .
Port Mode	The administrative mode of spanning tree on the port (Enable or Disable).
Port Forwarding State	- a Blocking - a Listening - a Learning - a Forwarding - a Disabled
Port Role	The role of the port within the MST, which is one of the following: - a Root - a Designated - a Alternate - a Backup - a Master – - a Disabled
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.

- a Click the **Details** button to view additional information for the selected interface in the MST topology.

4.26.6 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.

4 Configuring Switching Information

To display the Spanning Tree Statistics page, click **Switching > Spanning Tree > Statistics** in the navigation menu.

Interface	STP BPDUs Rx	STP BPDUs Tx	RSTP BPDUs Rx	RSTP BPDUs Tx	MSTP BPDUs Rx	MSTP BPDUs Tx	Sstp BPDUs Rx	Sstp BPDUs Tx
1/0/1	0	0	0	0	0	5852	0	0
1/0/2	0	0	0	0	0	0	0	0
1/0/3	0	0	0	0	0	0	0	0
1/0/4	0	0	0	0	0	75752	0	0
1/0/5	0	0	0	0	0	0	0	0
1/0/6	0	0	0	0	0	0	0	0
1/0/7	0	0	0	0	0	0	0	0
1/0/8	0	0	0	0	0	0	0	0
1/0/9	0	0	0	0	0	0	0	0
1/0/10	0	0	0	0	0	0	0	0

Figure 311: Spanning Tree Statistics

Table 295: Spanning Tree Statistics Fields

Convention	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.
Sstp BPDUs Rx	The number of classic Sstp BPDUs received by the interface.
Sstp BPDUs Tx	The number of classic Sstp BPDUs sent by the interface.

Click **Refresh** to update the screen with most recent data.

4.26.7 PVST/RPVST Global Configuration

Use this page to view and configure Per VLAN Spanning Tree Protocol (PVST)/Per VLAN Rapid Spanning Tree Protocol (RPVST) Global settings for the device.

To display the PVST/RPVST Global Configuration page, click **Switching** > **Spanning Tree** > **PVST Global** in the navigation menu.

Figure 312: Management Access Menu

Table 296: PVST/RPVST Global Configuration 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. When enabled, the switch accelerates the selection of a new root port, when a link or a switch fails or the spanning tree configuration is reconfigured. This is achieved by transitioning the root port directly to forwarding state instead of going through the listening and learning states first.
Max Update Rate (pps)	Configures Fast Uplink's Maximum Update Rate.

Use the buttons to perform the following tasks:

- a If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- a Click **Refresh** to update the screen with most recent data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.26.8 PVST/RPVST VLAN Configuration

Use this page to view and configure Per VLAN Spanning Tree Protocol (PVST)/Per VLAN Rapid Spanning Tree Protocol (RPVST) VLAN settings for the device.

To display the PVST/RPVST VLAN Configuration page, click **Switching** > **Spanning Tree** > **PVST VLAN** in the navigation menu.

Figure 313: PVST/RPVST VLAN Configuration

Table 297: PVST/RPVST VLAN Configuration 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 configure The PVST/RPVST VLANs, use the following buttons to perform the tasks:

- a Click **Refresh** to update the screen with most recent data.
- a Click **Add** to add a new row to the VLAN configuration.
- a To view details of any VLAN, select the entry and click the **Details** button.

Details of PVST/RPVST VLAN Configuration					
VLAN ID		1			
Root ID					
Priority	32769				
Address	00:A0:57:53:6F:C5				
Cost	0				
Port	This switch is the root				
Hello Time (Seconds)	2				
Max Age (Seconds)	20				
Forward Delay (Seconds)	15				
Bridge ID					
Priority	32769 (priority 32768 sys-id-ext 1)				
Address	00:A0:57:53:6F:C5				
Hello Time (Seconds)	2				
Max Age (Seconds)	20				
Forward Delay (Seconds)	15				
Aging Time (Seconds)	300				
Interface Details					
Display 10 rows		Showing 1 to 10 of 74 entries		Filter:	
Interface	Role	Status	Cost	Prio.Nbr	
1/0/1	Disabled	Disabled	0	128.1	
1/0/2	Disabled	Disabled	0	128.2	
1/0/3	Disabled	Disabled	0	128.3	
1/0/4	Disabled	Disabled	0	128.4	
1/0/5	Disabled	Disabled	0	128.5	
1/0/6	Disabled	Disabled	0	128.6	
1/0/7	Disabled	Disabled	0	128.7	
1/0/8	Disabled	Disabled	0	128.8	
1/0/9	Designated	Forwarding	20000	128.9	
1/0/10	Disabled	Disabled	0	128.10	
First Previous 1 2 3 4 5 Next Last					

Figure 314: Details of PVST/RPVST VLAN Configuration

Table 298: Details of PVST/RPVST VLAN Configuration Fields

Field	Description
VLAN ID	The unique VLAN identifier (VID).
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.
Cost	The cost value of the interface.
Prio.Nbr	The priority and neighbor of the interface.

- a Select an entry and then click **Edit** to change the PVST configuration on the VLAN.

The screenshot shows a dialog box titled "Edit PVST/RPVST VLAN Configuration". It contains the following fields and values:

VLAN ID	1
Root	☐ Primary ☐ Secondary ☒ None
Hello Time (Seconds)	2 (1 to 10)
Forward Delay (Seconds)	15 (4 to 30)
Max Age (Seconds)	20 (6 to 40)
Bridge Priority	8000 (0 to F000 hex)

At the bottom right, there are "Submit" and "Cancel" buttons.

Figure 315: Edit PVST/RPVST VLAN Configuration

a **Table 299: Edit PVST/RPVST VLAN Configuration 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.

- a Select an entry and then click **Remove** to remove the PVST row from the VLAN configuration.

4.26.9 PVST/RPVST Interface Configuration

Use this page to view and configure Per VLAN Spanning Tree Protocol (PVST)/Per VLAN Rapid Spanning Tree Protocol (RPVST) Interface settings for the device.

To display the PVST/RPVST Interface Configuration page, click **Switching** > **Spanning Tree** > **PVST Interface** in the navigation menu.

Figure 316: PVST/RPVST Interface Configuration

Table 300: PVST/RPVST Interface Configuration Fields

Field	Description
Interface	Select a physical or LAG interface.
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. Lower numbers mark a port as preferred for forwarding of frames.
Per VLAN Configuration - Configuration of each VLAN.	

Field	Description
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.

To configure the PVST/RPVST Interfaces, use the following buttons to perform the tasks:

- a Click **Refresh** to update the screen with most recent data.
- a Select an entry and then click **Edit** to change the PVST interface configuration.

Figure 317: Edit PVST/RPVST Interface Configuration

Table 301: Edit PVST/RPVST Interface Configuration Fields

Field	Description
VLAN ID	Shows the selected unique VLAN identifier (VID).
Interface	Shows the selected physical or LAG interface.
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.

4.26.10 PVST/RPVST Statistics

Use this page to view and configure Per VLAN Spanning Tree Protocol (PVST)/Per VLAN Rapid Spanning Tree Protocol (RPVST) Statistics settings for the device.

To display the PVST/RPVST Statistics page, click **Switching > Spanning Tree > PVST Statistics** in the navigation menu.




Figure 318: PVST/RPVST Statistics

Table 302: PVST/RPVST 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.

4.27 802.1p Priority Mapping

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.

802.1p 802.1p Priority Mapping

Display 10 rows Showing 1 to 10 of 119 entries Filter:

	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
☐	1/0/1	1	0	0	1	2	2	3	3
☐	1/0/2	1	0	0	1	2	2	3	3
☐	1/0/3	1	0	0	1	2	2	3	3
☐	1/0/4	1	0	0	1	2	2	3	3
☐	1/0/5	1	0	0	1	2	2	3	3
☐	1/0/6	1	0	0	1	2	2	3	3
☐	1/0/7	1	0	0	1	2	2	3	3
☐	1/0/8	1	0	0	1	2	2	3	3
☐	1/0/9	1	0	0	1	2	2	3	3

First Previous 1 2 3 4 5 Next Last

Refresh Edit

Figure 319: 802.1p Priority Mapping

Table 303: 802.1p Priority Mapping Fields

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.

Use the buttons to perform the following tasks:

- Click **Refresh** to update the information on the screen with the most current data.

a To configure 802.1p priority mapping for a specific interface click **Edit**.

802.1p Priority	Traffic Class
0	1 (0 to 6, 1 = Default)
1	0 (0 to 6, 0 = Default)
2	0 (0 to 6, 0 = Default)
3	1 (0 to 6, 1 = Default)
4	2 (0 to 6, 2 = Default)
5	2 (0 to 6, 2 = Default)
6	3 (0 to 6, 3 = Default)
7	3 (0 to 6, 3 = Default)

Figure 320: Edit 802.1p Priority Mapping

Table 304: Edit 802.1p Priority Mapping Fields

Field	Description
Interface	Shows the selected interface(s). The Global entry represents the common settings for all interfaces, unless specifically overridden individually.
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.

4.28 Configuring Port Security

Port Security can be enabled on a per-port basis. When a port is locked, only packets with whitelisted source MAC addresses can be forwarded. All other packets are discarded. A MAC address can be defined as whitelisted by one of two methods: dynamically or statically.

4 Both methods are used concurrently when a port is locked.

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. When 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 a whitelisted source MAC address can be forwarded.

To see the MAC addresses learned on a specific port, see [Configuring and Searching the Forwarding Database](#) on page 88. Disabled ports can only be activated from the Configuring Ports page.

4.28.1 Port Security Global Administration

Use the Port Security Global Administration page to enable or disable the port security feature on your switch.

To access the Port Security Global Administration page, click **Switching** > **Port Security** > **Global** in the navigation menu.

Figure 321: Port Security Global Administration

Select **Enable** or **Disable** from the **Port Security Mode** list.

Use the buttons to perform the following tasks:

- a If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.28.2 Port Security Interface Status

Use this page to configure the Port Security feature on a selected interface.

To access the Port Security Interface Status page, click **Switching** > **Port Security** > **Interface** in the navigation menu.

Interface	Port Security Mode	Max Dynamic Addresses Allowed	Max Static Addresses Allowed	Sticky Mode	Violation Trap Mode	Violation Shutdown Mode	Last Violation MAC/VLAN
1/0/1	Disable	600	20	Disable	Disable	Disable	
1/0/2	Disable	600	20	Disable	Disable	Disable	
1/0/3	Disable	600	20	Disable	Disable	Disable	
1/0/4	Disable	600	20	Disable	Disable	Disable	
1/0/5	Disable	600	20	Disable	Disable	Disable	
1/0/6	Disable	600	20	Disable	Disable	Disable	
1/0/7	Disable	600	20	Disable	Disable	Disable	
1/0/8	Disable	600	20	Disable	Disable	Disable	
1/0/9	Disable	600	20	Disable	Disable	Disable	
1/0/10	Disable	600	20	Disable	Disable	Disable	

Figure 322: Port Security Interface Status

Table 305: Port Security Interface Status Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.

4 Configuring Switching Information

Field	Description
Port Security Mode	Determines whether port security is enabled. The default mode is Disable. - a Enable – Locks the port so that only packets with allowable source MAC addresses can be forwarded. All other packets are discarded. - a Disable – The port is not locked, so no port security restrictions are applied.
Max Dynamic Addresses Allowed	Sets the maximum number of dynamically learned MAC addresses on the selected interface. When 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 Static Addresses Allowed	Sets the maximum number of statically locked MAC addresses on the selected interface.
Sticky Mode	The sticky MAC address learning mode, which is one of the following: - a Enabled – MAC addresses learned or manually configured on this interface are learned in sticky mode. A sticky-mode MAC address is a MAC address that does not age out and is added to the running configuration. If the running configuration is saved to the startup configuration, the sticky addresses are saved to persistent storage and do not need to be relearned when the device restarts. Upon enabling sticky mode on an interface, all dynamically learned MAC addresses in the MAC address table for that interface are converted to sticky mode. Additionally, new addresses dynamically learned on the interface will also become sticky. - a Disabled – When a link goes down on a port, all of the dynamically learned addresses are cleared from the source MAC address table the feature maintains. When the link is restored, the interface can once again learn addresses up to the specified limit. If sticky mode is disabled after being enabled on an interface, the sticky-mode addresses learned or manually configured on the interface are converted to dynamic entries and are automatically removed from persistent storage.
Violation Trap Mode	Indicates whether the port security feature sends a trap to the SNMP agent when a port is locked and a frame with a MAC address not currently in the table arrives on the port. A port is considered to be locked once it has reached the maximum number of allowed dynamic or static MAC address entries in the port security MAC address table.
Violation Shutdown Mode	Indicates whether the port security feature shuts down the port after MAC limit is reached.
Last Violation MAC/VLAN	The source MAC address and, if applicable, associated VLAN ID of the last frame that was discarded at a locked port.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.

- a To configure the settings for one or more interfaces, select each entry to modify and click **Edit**.

Edit Port Configuration	
Interface	1/0/1
Port Security Mode	☒ Enable ☐ Disable
Max Dynamic Addresses Allowed	600 (0 to 600)
Max Static Addresses Allowed	20 (0 to 20)
Sticky Mode	☐ Enable ☒ Disable
Violation Trap Mode	☐ Enable ☒ Disable
Violation Shutdown Mode	☐ Enable ☒ Disable
Submit Cancel	

Figure 323: Edit Port Security Interface Configuration

Table 306: Edit Port Security Interface Configuration Fields

Field	Description
Interface	Shows the selected interface(s).
Port Security Mode	Determines whether port security is enabled. a Enable a Disable
Max Dynamic Addresses Allowed	Sets the maximum number of dynamically learned MAC addresses on the selected interface. When 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 Static Addresses Allowed	Sets the maximum number of statically locked MAC addresses on the selected interface.
Sticky Mode	The sticky MAC address learning mode, which is one of the following: a Enabled a Disabled
Violation Trap Mode	Indicates whether the port security feature sends a trap to the SNMP agent when a port is locked and a frame with a MAC address not currently in the table arrives on the port. A port is considered to be locked once it has reached the maximum number of allowed dynamic or static MAC address entries in the port security MAC address table.
Violation Shutdown Mode	Indicates whether the port security feature shuts down the port after MAC limit is reached (Enable).

- a To apply the same settings to all interfaces, click **Edit All**.

4.28.3 VLAN MAC Locking Status

Use this page to configure VLAN MAC Locking. VLAN MAC locking allows you to secure the network by locking down whitelisted 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.

4 Configuring Switching Information

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 324: VLAN MAC Locking Status


Table 307: 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.
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 reboots. 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. When the value 0 is entered, no dynamic addresses can be learned.
Operational MAC Limit	The current number of dynamically learned source MAC addresses.
Violation Shutdown Mode	After the configured MAC limit has been reached, action will shut down the ports participating in the VLAN (Enable).
Violation Trap Mode	After the configured MAC limit has been reached, a log message will be generated with violation MAC address details (Enable).

To configure the VLAN MAC Locking, use the following buttons to perform the tasks:

- a Click **Refresh** to update the information on the screen with the most current data.

- a Use **Add** to configure VLAN MAC Locking.



The 'Add VLAN' dialog box contains the following fields and options:

VLAN ID	1
Max Dynamic Addresses Allowed	0 (0 to 600)
Violation Shutdown Mode	☐ Enable ☒ Disable
Violation Trap Mode	☐ Enable ☒ Disable

Buttons: Submit, Cancel

Figure 325: Add MAC Locking VLAN

- a **Table 308: Add MAC Locking VLAN Fields**


Field	Description
VLAN ID	Select the VLAN ID to be used for MAC Locking.
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. When the value 0 is entered, no dynamic addresses can be learned.
Violation Shutdown Mode	After the configured MAC limit has been reached, action will shut down the ports participating in the VLAN (Enable).
Violation Trap Mode	After the configured MAC limit has been reached, a log message will be generated with violation MAC address details (Enable).

- a Use **Edit** to modify configuration parameters of VLAN MAC Locking.
- a Use **Remove** to remove configured VLANs.

4.28.4 Port Security Static MAC Addresses

Use the Port Security Static MAC Addresses page to add and remove static MAC addresses configured on an interface.

To access the Port Security Static MAC Addresses page, click **Switching > Port Security > Static MAC** in the navigation menu.



The 'Port Security Static MAC Addresses' page shows a navigation bar with tabs: Global, Interface, VLAN, Static MAC (selected), and Dynamic MAC. Below the tabs, the page title 'Port Security Static MAC Addresses' is displayed. A table is shown with the following columns: Interface, Static MAC Address, VLAN ID, and Sticky Mode. The table is currently empty, displaying 'Table is Empty'. Above the table, there is a 'Display' dropdown set to 'All' rows, a 'Showing 0 to 0 of 0 entries' indicator, and a 'Filter' input field. Below the table, there are 'First', 'Previous', 'Next', and 'Last' navigation buttons. At the bottom of the page, there are 'Refresh', 'Add', and 'Remove' buttons.

Figure 326: Port Security Static MAC Addresses

Table 309: Port Security Static MAC Addresses Fields

Field	Description
Interface	Displays the physical interface or the LAG on which to view the dynamically learned MAC addresses.
Static 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.
Sticky Mode	Indicates whether the static MAC address entry is added in sticky mode. When adding a static MAC address entry, the Sticky Mode field can be selected only if it is enabled on the interface. If a static MAC address is added in sticky mode, and sticky mode is disabled on the interface, the MAC address entry is converted to a dynamic entry and will age out and be removed from the running (and saved) configuration if it is not relearned.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a To associate a static MAC address with an interface, click **Add** and configure the settings in the available fields.

Add Static MAC Entry

Interface

1/0/1

Static MAC Address

(xx:xx:xx:xx:xx:xx)

VLAN ID

(1 to 4093)

Sticky Mode

☐

Submit

Cancel

Figure 327: Add Static MAC Entry

Table 310: Add Static MAC Entry Fields

Field	Description
Interface	Select the physical interface or the LAG to associate with the permitted MAC address.
Static MAC Address	Enter a MAC address to be whitelisted.
VLAN ID	Enter the VLAN ID corresponding to the statically configured MAC address.
Sticky Mode	Indicates whether the static MAC address entry is added in sticky mode. When adding a static MAC address entry, the Sticky Mode field can be selected only if it is enabled on the interface. If a static MAC address is added in sticky mode, and sticky mode is disabled on the interface, the MAC address entry is converted to a dynamic entry and will age out and be removed from the running (and saved) configuration if it is not relearned.

- a To remove one or more configured static MAC address entries, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

4.28.5 Port Security Dynamic MAC Addresses

Use the Port Security Dynamic 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 MAC Addresses page, click **Switching** > **Port Security** > **Dynamic MAC** in the navigation menu.



Figure 328: Port Security Dynamic MAC Addresses

Table 311: Port Security Dynamic MAC Addresses Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When converting dynamic addresses to static addresses, use the Interface menu to select the interface to associate with the MAC addresses.
Dynamic 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a Click the **Convert to Static** button to convert all MAC addresses learned on an interface to static MAC address entries. After you click the button, a window opens and allows you to select the interface associated with the MAC address entries to convert. A static MAC address entry is written to the running configuration file and does not age out.

4.29 Managing LLDP

The IEEE 802.1AB defined standard, Link Layer Discovery Protocol (LLDP), allows stations residing in a local network to advertise major capabilities and physical descriptions. This information is viewed by a network manager to identify system topology and detect bad configurations in the network.

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.

LCOS SX 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 an LLDP port. If all the remote entries on the switch are filled up, 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.

4.29.1 LLDP 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.

GlobalInterfaceLocal DevicesRemote DevicesStatistics

LLDP Global Configuration?

Transmit Interval (Seconds)	30	(5 to 32768)
Transmit Hold Multiplier (Seconds)	4	(2 to 10)
Re-Initialization Delay (Seconds)	2	(1 to 10)
Notification Interval (Seconds)	5	(5 to 3600)

SubmitRefreshCancel

Figure 329: LLDP Global Configuration

Table 312: 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 5-32768 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.

Use the buttons to perform the following tasks:

- a If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.29.2 LLDP Interface Summary

Use the LLDP Interface Summary page to specify LLDP parameters that are applied to a specific interface. To display the LLDP Interface Summary page, click **Switching > LLDP > Interface** in the navigation menu.

Figure 330: LLDP Interface Summary

- 4** 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 313: 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.
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: - a Port Description – To include port description TLV in LLDP frames. To configure the Port Description, see Port Description on page 116. - a System Name – To include system name TLV in LLDP frames. To configure the System Name, see System Description on page 37. - a System Description – To include system description TLV in LLDP frames. - a System Capabilities – To include system capability TLV in LLDP frames.
Transmit Management Information	Tick the checkbox to enable the transmission of management address instance. Clear the checkbox to disable management information transmission. The default is disabled.

Use the buttons to perform the following tasks:

4 Configuring Switching Information

- a Click **Refresh** to update the information on the screen with the most current data.
- a To configure LLDP settings on an interface that does not have any LLDP settings enabled, click **Add**.

Figure 331: Add LLDP Interface

Table 314: Add LLDP Interface Fields

Field	Description
Interface	Select the interface to associate with the desired LLDP settings. 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 .
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)	
Port Description	Select this option to include the user-configured port description in the LLDPDU the interface transmits.
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 functions of the device in the LLDPDU the interface transmits.

- a 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.
- a To clear (disable) all LLDP settings from one or more interfaces, select each entry to clear and click **Remove**.

4.29.3 LLDP Local Device Summary

Use the LLDP Local Device Summary 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 332: LLDP Local Device Summary

Table 315: LLDP Local Device Summary Fields

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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a 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 316: LLDP Local Device 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.

Field	Description
System Description	The device description, which includes information about the product model and platform.
System Capabilities Supported	The primary functions the device supports.
System Capabilities Enabled	The primary functions 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.

4.29.4 LLDP Remote Device Summary

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 333: LLDP Remote Device Summary

Table 317: LLDP Remote Device Summary Fields

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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a 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 318: LLDP Remote Device Summary Fields

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 functions the remote system supports. The possible capabilities include the following options: - a Other - a Repeater - a Bridge - a WLAN AP - a Router - a Telephone - a DOCSIS®cable device - a Station
System Capabilities Enabled	The primary functions of the remote system that are both supported and enabled. The possible capabilities include the following options: - a Other - a Repeater - a Bridge - a WLAN AP - a Router - a Telephone - a DOCSIS®cable device - a 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.

4.29.5 LLDP 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.

LLDP Statistics										
Last Update	0d:00:00:00									
Total Inserts	0									
Total Deletes	0									
Total Drops	0									
Total Ageouts	0									
Display All rows Showing 0 to 0 of 0 entries Filter:										
Interface	Transmit Total	Receive Total	Discards	Errors	Ageouts	TLV Discards	TLV Unknowns	TLV MED	TLV 802.1	TLV 802.3
Table is Empty										
First Previous Next Last										
Refresh Clear										

Figure 334: LLDP Statistics

Table 319: LLDP Statistics Fields

Field	Description
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.
Port Statistics	
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.

Field	Description
TLV Discards	Displays the number of LLDP TLVs (Type, Length, Value sets) discarded for any reason by the LLDP agent on the corresponding port.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the page with the most current information.
- a Click **Clear** to clear the LLDP statistics of all the interfaces.

4.29.6 LLDP-MED

The Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) is an enhancement to LLDP that features:

- a Auto-discovery of LAN policies (such as VLAN, Layer 2 Priority, and DiffServ settings), enabling plug and play networking.
- a Device location discovery for creation of location databases.
- a Extended and automated power management of Power over Ethernet endpoints.
- a Inventory management, enabling network administrators to track their network devices and determine their characteristics (manufacturer, software and hardware versions, serial/asset number).

4.29.6.1 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 335: LLDP-MED Global Configuration

Table 320: LLDP-MED 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: a Class I Generic [IP Communication Controller etc.]

4 Configuring Switching Information

Field	Description
	- a Class II Media [Conference Bridge etc.] - a Class III Communication [IP Telephone etc.] 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.

Use the buttons to perform the following tasks:

- a If you make any configuration changes, click **Submit** to apply the new settings to the switch.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last saved state.

4.29.6.2 LLDP-MED Interface Summary

Use this page to enable LLDP-MED mode on an interface and to configure its properties.

To display this page, click **Switching > LLDP-MED > Interface** in the navigation menu.

Interface	Link Status	MED Status	Notification Status	Operational Status	Transmit TLVs
1/0/1	Up	Disable	Disable	Disable	0, 1
1/0/2	Down	Disable	Disable	Disable	0, 1
1/0/3	Down	Disable	Disable	Disable	0, 1
1/0/4	Up	Disable	Disable	Disable	0, 1
1/0/5	Down	Disable	Disable	Disable	0, 1
1/0/6	Down	Disable	Disable	Disable	0, 1
1/0/7	Down	Disable	Disable	Disable	0, 1
1/0/8	Down	Disable	Disable	Disable	0, 1
1/0/9	Down	Disable	Disable	Disable	0, 1
1/0/10	Down	Disable	Disable	Disable	0, 1

Figure 336: LLDP-MED Interface Summary

Table 321: LLDP-MED Interface Summary Fields

Field	Description
Interface	Select the port that you want to configure LLDP-MED–802.1AB on. You can select the checkbox next to Interface to configure all interfaces with the same properties.
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	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	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 TLVs that the interface transmits: a MED Capabilities: 0

Field	Description
	a Network Policy: 1

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a To configure LLDP-MED settings on an interface that does not have any LLDP-MED settings enabled, click **Add**.

5 The **Add** button can only be used when an existing interface is removed beforehand.

- a 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.

Figure 337: Edit LLDP-MED Interface

Table 322: Edit LLDP-MED Interface Fields

Field	Description
Interface	The interface associated with the rest of the data in the row. When configuring LLDP-MED settings, this field identifies the interfaces that are being configured.
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.
Configuration Notification Mode	Indicates whether LLDP-MED topology change notifications are enabled or disabled on the interface.
Transmit TLVs	
Capabilities	Enable this setting to activate LLDP-MED Capabilities transmission.
Network Policy	Enable this setting to activate LLDP-MED Network Policy transmission.

- a To remove all LLDP-MED settings from one or more interfaces, select each entry to remove and click **Remove**.

4.29.6.3 LLDP-MED Local Device Summary

This page displays a summary of LLPD-MED information advertised on the selected local interface. To display this page, click **Switching > LLDP-MED > Local Devices** in the navigation menu.




Figure 338: LLDP-MED Local Device Summary

Table 323: LLDP-MED Local Device Summary

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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a To view additional LLDP-MED information for a local interface, select the interface with the information to view and click **Details**.




Table 324: LLDP-MED Local Device Information Fields

Field	Description
Interface	This field identifies the interface that is being viewed.

Field	Description
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 following application types exist: - a unknown - a voicesignaling - a guestvoice - a guestvoicesignalling - a softphonevoice - a videoconferencing - a streamingvideo - a 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: - a Coordinate Based – The location map coordinates (latitude, longitude and altitude) of the device. - a Civic Address – The civic or street address location of the device. - a ELIN – The Emergency Call Service (ECS) Emergency Location Identification Number (ELIN) of the device.
Information	This column displays the information related to the coordinates, civic address, and ELIN for the device.

Figure 339: LLDP-MED Local Device Information

4.29.6.4 LLDP-MED Remote Device Summary

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. The following information is organized according to the order in which the fields appear in the LLDP-MED Remote Device Summary window.

To display this page, click **Switching > LLDP-MED > Remote Devices** in the navigation menu.




Figure 340: LLDP-MED Remote Device Summary

Table 325: LLDP-MED Remote Device Summary

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.
Device Class	The MED Classification advertised by the TLV from the remote device. The following three classifications represent the actual endpoints: - a Class I Generic (for example, IP Communication Controller) - a Class II Media (for example, Conference Bridge) - a Class III Communication (for example, IP Telephone) 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a To view additional LLDP-MED information for a local interface, select the interface with the information to view and click **Details**.

Table 326: LLDP-MED Remote Device Information

Field	Description
Interface	The local interface that has received LLDP-MED data units from remote devices.
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: - a Class I Generic (for example, IP Communication Controller) - a Class II Media (for example, Conference Bridge) - a Class III Communication (for example, IP Telephone)

Field	Description
	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 following application types exist: - a unknown - a voicesignaling - a guestvoice - a guestvoicesignalling - a softphonevoice - a videoconferencing - a streamingvideo - a 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.
Inventory Information	
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.
Location Information	
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.

5 Configuring Routing

LCOS SX 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 switch searches the ARP 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.

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.

4 LCOS SX supports the Border Gateway Protocol (BGP). BGP is only available on XS-6128QF switches. The BGP features can be configured only by using the CLI. No web-based administrative pages are available for BGP configuration.

5.1 Configuring ARP

The ARP protocol associates a Layer 2 MAC address with a Layer 3 IPv4 address. LCOS SX 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. When 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.

5.1.1 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 341: ARP Table

The ARP Table displays at the bottom of the page, and contains the following fields:

Table 327: ARP Table Fields

Field	Description
IP Address	The IP address of a network host on a subnet attached to one of the device's routing interfaces.
MAC Address	The unicast MAC address (hardware address) associated with the network host.
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. An entry can have multiple types, but only the type with the highest priority is shown (e.g. an entry can have the types Gateway and Local, in this case only the type Local is shown). - a Dynamic – An ARP entry that has been learned by the router - a Gateway – A dynamic ARP entry that has the IP address of a routing interface - a Local – An ARP entry associated with the MAC address of a routing interface on the device - a Static – An ARP entry configured by the user
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).

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the page with the most current information.

- a 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.

A dialog box titled "Add Static ARP Entry" with a close button (X) in the top right corner. It contains two input fields: "IP Address" with a placeholder "(x.x.x.x)" and "MAC Address" with a placeholder "(xx:xx:xx:xx:xx:xx)". At the bottom right, there are "Submit" and "Cancel" buttons.

Figure 342: Add Static ARP Entry

a Table 328: Add Static ARP Entry 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.
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.

- a 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.
- a 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.

5.1.2 ARP Table Configuration

Use this page to change the configuration parameters for the Address Resolution Protocol Table. To display the page, click **Routing > ARP Table > Configuration** in the navigation menu.


A web page titled "ARP Table Configuration" with tabs for "Summary", "Configuration", and "Statistics". The "Configuration" tab is active. It contains a table with the following fields: "Age Time (Seconds)" with value 1200 (range 15 to 21600), "Response Time (Seconds)" with value 1 (range 1 to 10), "Retries" with value 4 (range 0 to 10), "Cache Size" with value 6144 (range 384 to 6144), and "Dynamic Renew" with an unchecked checkbox. At the bottom, there are "Submit", "Refresh", and "Cancel" buttons.

Figure 343: ARP Table Configuration

Table 329: 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.

Field	Description
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.

5.2 Configuring Global IP Settings

This menu allows you to configure and display IP routing data.

5.2.1 Routing IP Configuration

Use the Routing IP 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.

Field	Description
Routing Mode	☒ Disable ☐ Enable
ICMP Echo Replies	☒
ICMP Redirects	☒
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	10922
Global Default Gateway	

Figure 344: Routing IP Configuration

Table 330: Routing IP Configuration Fields

Field	Description
Routing Mode	The administrative mode of routing on the device. The options are as follows: - Enable – The device can act as a Layer 3 device by routing packets between interfaces configured for IP routing. - Disable – The device acts as a Layer 2 bridge and switches traffic between interfaces. The device does not perform any internetwork routing.

5 Configuring Routing

Field	Description
ICMP Echo Replies	Select the Enable or Disable check box. If you select Enable , 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	The maximum burst interval for ICMP error messages transmitted by the device. 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.
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: - a To configure the default gateway, click the Edit icon and specify the IP address of the default gateway in the available field. - a To reset the IP address of the default gateway to the factory default value, click the Reset icon associated with this field.

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

5.2.2 Routing IP 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.

Configuration **Interface Summary** Interface Configuration Loopback Configuration Statistics

Routing IP Interface Summary

Display **10** rows Showing 1 to 10 of 28 entries Filter:

	Interface	Status	IP Address	Subnet Mask	Admin Mode	State	MAC Address	Proxy ARP	IP MTU
☐	1/0/1	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:A0:57:60:A2:4F	Enabled	1500
☐	1/0/2	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:A0:57:60:A2:4F	Enabled	1500
☐	1/0/3	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:A0:57:60:A2:4F	Enabled	1500
☐	1/0/4	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:A0:57:60:A2:4F	Enabled	1500
☐	1/0/5	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:A0:57:60:A2:4F	Enabled	1500
☐	1/0/6	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:A0:57:60:A2:4F	Enabled	1500
☐	1/0/7	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:A0:57:60:A2:4F	Enabled	1500
☐	1/0/8	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:A0:57:60:A2:4F	Enabled	1500
☐	1/0/9	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:A0:57:60:A2:4F	Enabled	1500
☐	1/0/10	Down	0.0.0.0	0.0.0.0	Enabled	Inactive	00:A0:57:60:A2:4F	Enabled	1500

First Previous 1 2 3 Next Last

Refresh Edit Details Add Loopback Remove Loopback

Figure 345: Routing IP Interface Summary

Table 331: Routing IP 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the page with the most current information.
- a To edit any interface, select the interface and click **Edit** and you are redirected to the [Routing IP Interface Configuration](#).

5 Configuring Routing

- a To view additional routing configuration information for an interface, select the interface with the settings to view and click **Details**.



Details	
Interface	1/0/1
Status	Down
Routing Mode	Disabled
Admin Mode	Enabled
State	Inactive
Link Speed Data Rate	10000
IP Address Configuration Method	None
IP Address	
Subnet Mask	
MAC Address	00:A0:57:60:A2:4F
IP MTU	1500
Bandwidth	10000000
Encapsulation Type	Ethernet
Forward Net Directed Broadcasts	Disabled
Proxy ARP	Disabled
Local Proxy ARP	Disabled
Destination Unreachables	Enabled
ICMP Redirects	Enabled

Close

Figure 346: Routing IP Interface Summary Details

Table 332: Routing IP Interface Summary Details 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).
Routing Mode	Indicates whether routing is administratively enabled or disabled on the interface.
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.
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: - a None – The interface does not have an IP address. - a Manual – The IP address has been statically configured by an administrator. - a 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.
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.

Field	Description
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	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: a Enabled – Network directed broadcasts are forwarded. a Disabled – Network directed broadcasts are dropped.
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.
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.

- a After you click **Add Loopback**, the next available loopback interface is added and you are redirected to the [Routing IP Loopback Configuration](#). The button is disabled, if maximum no of loopback interfaces are configured.
- a After you click **Remove Loopback**, the selected entries are deleted on confirmation.

5.2.3 Routing IP Interface Configuration

Use the Routing IP 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 347: Routing IP Interface Configuration

Table 333: Routing IP Interface Configuration Fields

Field	Description
Type	The type of interface that can be configured for routing: a VLAN – Enables list of all VLANs that can be configured for routing. a Interface – Enables list of all non-loopback interfaces that can be configured for routing.
VLAN	The menu contains all VLANs that can be configured for routing. To configure routing settings for a VLAN, select it from the menu and then configure the rest of the settings on the page.
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).

Field	Description
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: - a None – No address is to be configured. - a 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. - a DHCP – The interface will attempt to acquire an IP address from a network DHCP server.
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 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

Field	Description
	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	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.

5.2.4 Routing IP Loopback Configuration

Use this page to configure the IP routing settings for each loopback interface.

To display the Routing IP Loopback Configuration page, click **Routing** > **IP** > **Loopback Configuration** in the navigation menu.

Figure 348: Routing IP Loopback Configuration

Table 334: Routing 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.

Field	Description
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.

Use the buttons to perform the following tasks:

- a Click **Submit** to send the updated configuration to the switch.
- a Click **Refresh** to update the page with the most current information.
- a After clicking **Add Loopback**, the next available loopback interface will be added. If the maximum number of loopback interfaces are configured this button will be disabled.
- a After you click **Remove Loopback**, the selected entry is deleted on confirmation.

5.2.5 Routing IP Statistics

The statistics reported on the Routing 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.

Configuration	Interface Summary	Interface Configuration	Loopback Configuration	Statistics
Routing IP Statistics				
IpInReceives	0			
IpInHdrErrors	0			
IpAddrErrors	0			
IpFwdDatagrams	2			
IpInUnknownProtos	128			
IpInDiscards	0			
IpInDelivers	418132			
IpOutRequests	277798			
IpOutDiscards	18			
IpOutNoRoutes	0			
IpReasmTimeout	0			
IpReasmReqds	0			
IpReasmOKs	0			
IpReasmFails	0			
IpFragOKs	0			
IpFragFails	0			
IpFragCreates	0			

Figure 349: Routing IP Statistics

Table 335: Routing 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.

5 Configuring Routing

Field	Description
IpAddrErrors	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.
IpFwdDatagrams	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 .
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 IpFwdDatagrams 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 IpFwdDatagrams 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.

Field	Description
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.
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.

5.3 Router

In this menu you can configure and display route tables.

5.3.1 Route Table Summary

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 350: Route Table Summary

Table 336: Route Table Summary 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.
Protocol	Identifies which protocol created the route. A route can be created one of the following ways: - a Dynamically learned through a supported routing protocol - a Dynamically learned by being a directly-attached local route - a Statically configured by an administrator - a Configured as a default route by an administrator The following protocols can occur: - a Local - a Static - a Default - a OSPF Intra - a OSPF Inter - a OSPF Type-1 - a OSPF Type-2
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.

5.3.2 Configured Route Summary

Use the Configured Route Summary page to create and display static routes.

To display the page, click **Routing** > **Router** > **Configured Routes** in the navigation menu.

Figure 351: Configured Route Summary

Table 337: Configured Route Summary 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 Address	The next hop router address to use when forwarding traffic to the destination.
Next Hop Interface	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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a To configure a route, click **Add** and specify the desired settings in the available fields.

Figure 352: Add Route

Table 338: Add Route Fields

Field	Description
Route Type	The type of route to configure, which is one of the following: - a Default – The route the device uses to send a packet if the routing table does not contain a longer matching prefix for the packet's destination. The routing table can contain only one default route. - a Static – A route that is manually added to the routing table by an administrator. - a Static Reject – A route where packets that match the route are discarded instead of forwarded. The device might send an ICMP Destination Unreachable message. 5 The route type you select determines the fields available on the page. Some of the fields that describes are not available when configuring certain types of routes.
Network Address	Specify the IP route prefix for the destination from the 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.
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.
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.

- a To remove a configured route, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

5.3.3 IP Route Summary

The IP Route 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.

Route Table

Configured Routes

Summary

ECMP Group

IP Route Summary

?

Route Types

Connected Routes	0
Static Routes	0
RIP Routes	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

Refresh

Clear Counters

Figure 353: IP Route Summary

Table 339: IP Route 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.
	5 BGP is only supported by XS-6128QF switches.
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.
OSPF Routes	The total number of routes installed by the OSPF protocol.

5 Configuring Routing

Field	Description
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a Click the **Clear Counters** button to reset to zero the 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.

5.3.4 ECMP Groups Summary

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 354: ECMP Groups Summary

Table 340: 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.

5.4 Configuring IPv6 Settings

In this menu you can configure and display IPv6 routing parameters.

5.4.1 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.

5 Configuring Routing

To display the IPv6 Global Configuration page, click **Routing > IPv6 > Configuration** in the navigation menu.

The screenshot shows the 'IPv6 Global Configuration' page. At the top, there are tabs for 'Configuration', 'Interface Summary', 'Interface Configuration', 'Loopback Configuration', 'Global Addresses', 'Address Configuration', 'Statistics', 'Detailed Statistics', and 'Neighbor Table'. The 'Configuration' tab is active. Below the tabs, the page title is 'IPv6 Global Configuration'. The configuration fields are as follows:

IPv6 Unicast Routing Mode	☒ Disable ☐ Enable
IPv6 Neighbors Dynamic Renew	☐
IPv6 Hop Limit	64 (1 to 255)
IPv6 Unresolved Packets Rate Limit (pps)	1024 (5 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

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

Figure 355: IPv6 Global Configuration

Table 341: IPv6 Global Configuration Fields

Field	Description
IPv6 Unicast Routing Mode	The administrative mode of IPv6 routing on the device. The options are as follows: a Disable – The device does not support IPv6 routing. a Enable – The device can act as a Layer 3 device by routing IPv6 packets between interfaces configured for IPv6 routing.
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 Unicast Solicits	The maximum number of unicast neighbor solicitations sent during NUD before switching to multicast neighbor solicitations.
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

Field	Description
	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.

5.4.2 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.

Interface	Operational Status	IPv6 Mode	Routing Mode	Admin Mode	IPv6 Prefix	Prefix Length	State
1/0/1	Disabled	Disabled	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			
1/0/9	Disabled	Disabled	Disabled	Enabled			
1/0/10	Disabled	Disabled	Disabled	Enabled			

Figure 356: IPv6 Interface Summary

Table 342: 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 : - a The IPv6 mode is enabled on the interface. - a The routing mode is enabled on the interface. - a The administrative mode is enabled on the interface. - a The link is up.
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.

5 Configuring Routing

Field	Description
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a 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.
- a To view additional routing configuration information for an interface, select the interface with the settings to view and click **Details**.



Interface	1/0/1
Operational Status	Disabled
Link Local Prefix	
Link Local Prefix Length	
Link Local Status	
Routing Mode	Disabled
IPv6 Mode	Disabled
Admin Mode	Enabled
DHCPv6 Client Mode	Disabled
Stateless Address AutoConfig	0
Interface Maximum Transmit Unit	1500
Router Duplicate Address Detection Transmits	1
Router Advertisement NS Interval	0
Router Lifetime Interval	1800
Router Advertisement Reachable Time	0
Router Advertisement Interval	600
Router Advertisement Managed Config	Disabled
Router Advertisement Other Config	Disabled
Router Advertisement Suppress	Disabled
IPv6 Destination Unreachable Messages	Enabled
IPv6 Hop Limit Unspecified	Disabled

Figure 357: IPv6 Interface Details

Table 343: IPv6 Interface Details 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.

Field	Description
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: - a The IPv6 mode is enabled on the interface. - a The routing mode is enabled on the interface. - a The administrative mode is enabled on the interface. - a The link is up.
Link Local Prefix	The autoconfigured link-local address which is: - a Allocated from part of the IPv6 unicast address space - a Not visible off the local link - a Not globally unique
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 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.
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.
Admin Mode	The administrative mode on the interface.
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

5 Configuring Routing

Field	Description
	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).

- a To add the next available loopback interface, click **Add Loopback**. You are redirected to the [IPv6 Loopback Configuration](#) page.

5.4.3 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.

The screenshot displays the 'IPv6 Interface Configuration' page. The 'Interface Configuration' tab is active. The configuration fields are as follows:

- Type:** ☒ VLAN ☐ Interface
- VLAN:** VLAN 1 (dropdown)
- Interface:** 1/0/1 (dropdown)
- Operational Status:** Disabled
- Link Local Prefix:**
- Link Local Prefix Length:**
- Link Local Status:**
- Routing Mode:** ☒ Disable ☐ Enable
- IPv6 Mode:** ☒ Disable ☐ Enable
- Admin Mode:** ☒ Disable ☐ Enable
- DHCPv6 Client Mode:** ☒ Disable ☐ Enable
- Stateless Address AutoConfig:** ☐
- Interface Maximum Transmit Unit:** 1500 (range: 1280 to 12288)
- Router Duplicate Address Detection Transmits:** 1 (range: 0 to 600)
- Router Advertisement NS Interval:** (range: 1000 to 4294967295)
- Router Lifetime Interval:** 1800 (range: 0 to 9000)
- Router Advertisement Reachable Time:** 0 (range: 0 to 3600000)
- Router Advertisement Interval:** 600 (range: 4 to 1800)
- Router Advertisement Managed Config:** ☐
- Router Advertisement Other Config:** ☐
- Router Advertisement Suppress:** ☐
- IPv6 Destination Unreachable Messages:** ☐
- ICMPv6 Redirects:** ☐
- IPv6 Hop Limit Unspecified:** ☐

Buttons at the bottom: Submit, Refresh, Cancel.

Figure 358: IPv6 Interface Configuration

Table 344: IPv6 Interface Configuration Fields

Field	Description
Type	The type of interface that can be configured for IPv6 routing: - a VLAN – Enables list of all VLANs that can be configured for IPv6 routing. - a Interface – Enables list of all non-loopback interfaces that can be configured for IPv6 routing.
VLAN	The menu contains all VLANs that can be configured for IPv6 routing. To configure routing settings for a VLAN, select it from the menu and then configure the rest of the settings on the page.
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 : - a The IPv6 mode is enabled on the interface. - a The routing mode is enabled on the interface. - a The administrative mode is enabled on the interface.

5 Configuring Routing

Field	Description
	a The link is up.
Link Local Prefix	The autoconfigured link-local address which is: a Allocated from part of the IPv6 unicast address space a Not visible off the local link a Not globally unique
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.
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. 5 This parameter can only be changed, when the Type is set to Interface .
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. 5 This parameter can only be changed, when the Type is set to Interface .
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.

Field	Description
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.

5.4.4 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 359: IPv6 Loopback Configuration

Table 345: 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 Interface Summary 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: a Allocated from part of the IPv6 unicast address space

Field	Description
	- a Not visible off the local link - a Not globally unique
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.

5.4.5 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 360: IPv6 Global Address Table

Table 346: 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.

Field	Description
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).

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a To edit any interface, select the interface and click **Edit**. You are redirected to the [IPv6 Global Address Configuration](#) page for the selected interface.
- a To delete the IPv6 address configuration from one or more interfaces, select each entry to remove and click **Remove**. You must confirm the action.

5.4.6 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 361: IPv6 Global Address Configuration

Table 347: 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

Field	Description
	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).

Use the buttons to perform the following tasks:

- a If you change any of the settings, click **Submit** to apply the changes to the switch. To preserve the changes across a switch reboot, you must perform a save.
- a To update the information on the screen, click **Refresh**.
- a Click **Cancel** to discard changes and revert to the last saved state.
- a To configure an IPv6 address on an interface that already has an IPv6 address, click **Add**.

5.4.7 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 display the IPv6 Statistics page, click **Routing > IPv6 > Statistics** in the navigation menu.

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

Refresh Details

Figure 362: IPv6 Statistics

Table 348: IPv6 Statistics Fields

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 ICMPv6 Messages With Errors Received . Note that this interface is the interface to which the ICMPv6 messages were addressed, which may not necessarily be 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.)
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 ICMPv6 Messages Not Transmitted Due To Error .
ICMPv6 Duplicate Address Detects	The number of duplicate IPv6 addresses detected by the interface.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a 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.

5.4.8 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.

5 Configuring Routing

To display the IPv6 Detailed Statistics page, click **Routing > IPv6 > Statistics** in the navigation menu. Only a portion of the menu is shown in the following image.

IPv6 Detailed Statistics	
Interface	1/0/1
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

Figure 363: IPv6 Detailed Statistics

Table 349: IPv6 Detailed Statistics Fields

Field	Description
IPv6 Statistics	
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.

Field	Description
Received Datagrams Discarded Due To Truncated Data	The number of input datagrams discarded because datagram frame didn't carry enough data.
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.
ICMPv6 Statistics	
Total ICMPv6 Messages Received	The total number of ICMPv6 messages received by the interface, which includes all those counted by ICMPv6 Messages With Errors Received . 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.)

5 Configuring Routing

Field	Description
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.
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 ICMPv6 Messages Not Transmitted Due To Error .
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.

Field	Description
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.
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.

5.4.9 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.

5 Configuring Routing

To display the IPv6 Neighbor Table page, click **Routing > IPv6 > Neighbor Table** in the navigation menu.




Figure 364: IPv6 Neighbor Table

Table 350: IPv6 Neighbor Table Fields

Field	Description
Interface	The local interface where 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.
Type	The type of the neighbor entry, which is one of the following: - a Static – The neighbor entry is manually configured. - a Dynamic – The neighbor entry is dynamically resolved. - a Local – The neighbor entry is a local entry. - a Other – The neighbor entry is an unknown entry.
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: - a 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. - a Reach - Positive confirmation was received within the last ReachableTime 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. - a 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. - a 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. - a Probe - A reachability confirmation is actively sought by resending neighbor solicitation messages every RetransTimer milliseconds until a reachability confirmation is received.
Last Updated	The amount of time that has passed since the address was confirmed to be reachable.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a To configure a neighbor entry, click **Add**.

Figure 365: Add IPv6 Neighbors

Table 351: Add IPv6 Neighbors Fields

Field	Description
Interface	Select a local interface for the IPv6 neighbor.
IPv6 Address	Enter the IPv6 prefix and prefix length of the neighbor interface.
MAC Address	Enter a MAC address to be 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.

- a To remove one or more configured neighbor entries, select each entry to remove and click **Remove**.

5.5 Configuring IPv6 Routes

In this menu you can configure and display IPv6 routing parameters.

5.5.1 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.

5 Configuring Routing

To display the IPv6 Route Table page, click **Routing > IPv6 Routes > IPv6 Route Table** in the navigation menu.




Figure 366: IPv6 Route Table

Table 352: 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: - a Dynamically learned through a supported routing protocol - a Dynamically learned by being a directly-attached local route - a Statically configured by an administrator
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.

5.5.2 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 display the IPv6 Configured Routes page, click **Routing > IPv6 Routes > IPv6 Configured Routes** in the navigation menu.




Figure 367: IPv6 Configured Routes

Table 353: IPv6 Configured Routes Fields

Field	Description
IPv6 Prefix/Length	The IPv6 address, including the prefix and prefix length, for the destination network.
Next Hop IPv6 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.

Use the buttons to perform the following tasks:

- Click **Refresh** to update the information on the screen.
- To configure a route, click **Add** and specify the desired settings in the available fields.

Figure 368: Add Route

Table 354: Add Route Fields

Field	Description
Route Type	The type of route to configure, which is one of the following: - a Global – A route with an address that is globally routable and is recognized outside of the local network. - a 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. - a Static Reject – A route where packets that match the route are discarded instead of forwarded. The device might send an ICMPv6 Destination Unreachable message.
IPv6 Prefix/Length	Enter an IPv6 address, including the prefix and prefix length, for the destination 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. 5 This parameter can only be changed, when the Route Type is set to Link Local.
Next Hop IPv6 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. 5 This parameter can only be changed, when the Route Type is set to Global or Link Local.
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.

- a To remove a configured route, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.
- a If you make any changes to the page, click **Submit** to apply the changes to the system.

5.5.3 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 369: IPv6 ECMP Groups Summary**

Table 355: 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.
Next Hops	The IPv6 address of each next hop in the group.

Click **Refresh** to update the information on the screen.

5.5.4 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.

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
Reject 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	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
Number of Prefixes	

Figure 370: IPv6 Route Summary**Table 356: IPv6 Route Summary Fields**

Field	Description
Route Types	
Connected Routes	The total number of connected routes in the IPv6 routing table.

5 Configuring Routing

Field	Description
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.
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.
Route Table Counters	
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.
Number of Prefixes	The unique IPv6 prefixes in the IPv6 routing table.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a Click the **Clear Counters** button to reset all IPv6 routing table event counters on this page to zero. Note 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.

5.6 Configuring DHCPv4

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. You can configure and display DHCP parameters and information in the DHCP Server menu.

5.6.1 DHCP Server Global Configuration

Use the DHCP Server Global Configuration page to configure DHCP global parameters.

To display the page, click **Routing > DHCP Server > Global** in the navigation menu.

Figure 371: DHCP Server Global Configuration

Table 357: 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.
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.

Use the buttons to perform the following tasks:

- a If you change any settings click **Submit** to apply the changes to the system.
- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Cancel** to discard changes and revert to the last state.

5.6.2 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 **Routing** > **DHCP Server** > **Excluded Addresses** in the navigation menu.




Figure 372: DHCP Server Excluded Addresses

Table 358: 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 name that identifies the VRF (VPN Routing and Forwarding) instance associated with the excluded IP address. 5 This parameter is only available on XS-6128QF switches.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a To add one or more IP addresses to exclude, click **Add** and specify the IPv4 address or range of addresses in the available fields.



Figure 373: Add Exclusion

Table 359: Add Exclusion 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 name that identifies the VRF instance associated with the excluded IP address. 5 This parameter is only available on XS-6128QF switches.

- a To remove an excluded address or range of addresses, select each entry to delete and click **Remove**.

5.6.3 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 **Routing > DHCP Server > Pool Summary** in the navigation menu.




Figure 374: DHCP Server Pool Summary

Table 360: 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. - a 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. - a Dynamic – The DHCP server can assign the client any available IP address within the pool. This type is also known as Automatic. - a Undefined – The pool has been created by using the CLI, but the pool information has not been configured.
Network	- a For a Manual pool, indicates the host IP address to assign the client. - a 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.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.

5 Configuring Routing

- a To add a pool, click **Add** and configure the pool information in the available fields.

The screenshot shows a 'Add DHCP Server Pool' window with the following fields and options:

- Pool Name: (1 to 31 characters)
- Type of Binding: ☒ Dynamic, ☐ Manual
- Network Base Address: (x.x.x.x)
- Network Mask: (x.x.x.x)
- Client Name: (0 to 255 characters)
- Hardware Address Type: Ethernet (dropdown)
- Hardware Address: (xx:xx:xx:xx:xx:xx)
- Client ID Type: ☒ Char, ☐ 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: ☐ Enable, ☒ Disable
- Lease Duration: Days 0, Hours 0, Minutes 0 (dropdowns)
- VRF Name: (0 to 64 characters)
- Default Router Address: (x.x.x.x)
- DNS Server Address 1: (x.x.x.x)
- DNS Server Address 2: (x.x.x.x)

Buttons: Submit, Cancel

Figure 375: Add DHCP Server Pool

Table 361: 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 Create . When the user selects Create , 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. a Dynamic – The DHCP server can assign the client any available IP address within the pool. This type is also known as Automatic. a 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. 4 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.
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 Ethernet and IEEE802 . The default value is Ethernet . 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.

Field	Description
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. - a 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. - a 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.
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 DHCP Server Pool Configuration on page 421 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 DHCP Server Pool Configuration on page 421 page.

- a To remove a pool, select each entry to delete and click **Remove**. You must confirm the action before the pool is deleted.

5.6.4 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.

5 Configuring Routing

To access the DHCP Server Pool Configuration page, click **Routing** > **DHCP Server** > **Pool Configuration** in the navigation menu.

Figure 376: DHCP Server Pool Configuration

Table 362: DHCP Server Pool Configuration Fields

Field	Description
Pool Name	Select the pool to configure. The menu includes all pools that have been configured on the device.
Type of Binding	Specifies the type of binding for the pool. 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. Dynamic – The DHCP server can assign the client any available IP address within the pool. This type is also known as Automatic. 5 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.
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.
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 Ethernet and IEEE802 . The default value is Ethernet . 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.

Field	Description
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	If Client ID type is HEX, specify the client in xx:xx:xx:xx:xx:xx format.
Host IP Address	(Manual pools only). The IP address to offer the client.
Host Mask	(Manual pools only). The subnet mask to offer the client.
Lease Expiration	Indicates whether the information the server provides to the client should expire. - a 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. - a 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.
Lease Duration	a The number of Days, Hours, and Minutes the lease is valid. This field cannot be configured if the Lease Expiration 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.
Next Server Address	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.
Default Router	Lists the IP address of each router to which the clients in the pool should send traffic. The default router should be in the same subnet as the client. To configure settings for one or more default routers that can be used by the clients in the pool, use the buttons available in the appropriate table to perform the following tasks: - a To add an entry to the Default router list, click the + (plus) button and enter the IP address of the server to add. - a To edit the address of a configured Default router, click the Edit icon associated with the entry to edit and update the address. - a To delete an entry from the list, click the – (minus) button associated with the entry to remove. - a To delete all entries from the list, click the – (minus) button in the heading row.
DNS Server	Lists the IP address of each DNS server the clients in the pool can contact to perform address resolution. To configure settings for one or more DNS servers that can be used by the clients in the pool, use the buttons available in the appropriate table to perform the following tasks: - a To add an entry to the DNS server list, click the + (plus) button and enter the IP address of the server to add. - a To edit the address of a configured DNS server, click the Edit icon associated with the entry to edit and update the address.

5 Configuring Routing

Field	Description
	- a To delete an entry from the list, click the – (minus) button associated with the entry to remove. - a To delete all entries from the list, click the – (minus) button in the heading row.
NetBIOS Server	Lists the IP address of each NetBIOS Windows Internet Naming Service (WINS) name server that is available for the selected pool. To configure settings for one or more NetBIOS servers that can be used by the clients in the pool, use the buttons available in the appropriate table to perform the following tasks: - a To add an entry to the NetBIOS server list, click the + (plus) button and enter the IP address of the server to add. - a To edit the address of a configured NetBIOS server, click the Edit icon associated with the entry to edit and update the address. - a To delete an entry from the list, click the – (minus) button associated with the entry to remove. - a To delete all entries from the list, click the – (minus) button in the heading row.

Use the buttons to perform the following tasks:

- a After you configure values for the DHCP address pool, click **Submit** to create the pool and apply the changes to the system.
- a To update the information on the screen with the latest information, click **Refresh**.
- a Click **Cancel** to discard changes and revert to the last saved state.

5.6.5 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 broadcasts 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 **Routing > DHCP Server > Pool Options** in the navigation menu.

If no DHCP pools exist, the DHCP Server Pool Options page does not display the fields.

Figure 377: DHCP Server Pool Options

If any DHCP pools are configured on the system, the DHCP Server Pool Options page contains the following fields.

Table 363: 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 Edit icon in the row. To reset the field to the default value, click the Reset icon in the row. The options are: - a B-Node Broadcast – Broadcast only. - a H-Node Hybrid – NetBIOS name server, then broadcast. - a M-Node Mixed – Broadcast, then NetBIOS name server. - a P-Node Peer-to-Peer – NetBIOS name server only.
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.
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	The type of data to associate with the Option Code, which can be one of the following: - a ASCII - a HEX - a IP Address
Option Value	The data associated with the Option Code. When adding or editing a vendor option, the field(s) 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.

The option table shows the Vendor Options that have been added to the selected pool. Use the buttons to perform the following tasks:

5 Configuring Routing

- a Click **Refresh** to update the information on the screen with the most current data.
- a To add a vendor option, click **Add Vendor Option** and configure the desired information in the available fields.

The screenshot shows a window titled "Add Vendor Option". Inside, there are several input fields: "Pool Name" with the value "DHCP.Pool", "Option Code" with a range "(1 to 254)", "Option Type" with radio buttons for "ASCII" (selected), "HEX", and "IP Address". Below these are fields for "ASCII", "HEX" (with a format "(xx:xx:xx:xx:xx:xx)"), and "IP Address" (with a table-like structure for multiple entries). At the bottom right are "Submit" and "Cancel" buttons.

Figure 378: Add Vendor Option

Table 364: Add Vendor Options Fields

Field	Description
Pool Name	Shows the selected DHCP pool the DHCP vendor options should be added to.
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: - a ASCII – The option type is a text string. - a HEX – The option type is a hexadecimal number. - a IP Address – The option type is an IP address.

- a To edit a vendor option, select the entry to change and click **Edit**.
- a To remove a vendor option, select each entry to delete and click **Remove**. You must confirm the action before the entry is deleted.

5.6.6 DHCP Server Bindings

Use the DHCP Server Bindings page to view and delete entries in the DHCP Bindings table. After a client leases an IP address from the DHCP server, the server adds an entry to its database. The entry is called a binding.

To access the DHCP Server Bindings page, click **Routing** > **DHCP Server** > **Bindings** in the navigation menu.

The screenshot shows the "DHCP Server Bindings" page. At the top, there are tabs: "Global", "Excluded Addresses", "Pool Summary", "Pool Configuration", "Pool Options", "Bindings" (selected), "Statistics", and "Conflicts". Below the tabs is a blue header "DHCP Server Bindings" with a help icon. The main area contains a table with columns: "IP Address", "Hardware Address", "Lease Time Left", "Pool Allocation Type", "Pool Name", and "VRF Name". The table is empty, with a message "Table is Empty" in the center. Above the table, there is a "Display" dropdown set to "All" and "rows", and a "Showing 0 to 0 of 0 entries" indicator. Below the table are "First", "Previous", "Next", and "Last" navigation links. At the bottom are "Refresh" and "Clear Entries" buttons.

Figure 379: DHCP Server Bindings

Table 365: 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: - a Dynamic – The address was allocated dynamically from a pool that includes a range of IP addresses. - a Manual – A static IP address was assigned based on the MAC address of the client. - a Inactive – The pool is not in use.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a 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.

5.6.7 DHCP Server Statistics

Use the DHCP Server Statistics page to view information about the DHCP server bindings and messages. The values on this page indicate the various counts that have accumulated since they were last cleared. To access the DHCP Server Statistics page, click **Routing** > **DHCP Server** > **Statistics** in the navigation menu.

Global Excluded Addresses Pool Summary Pool Configuration Pool Options Bindings Statistics Conflicts	
DHCP Server Statistics ?	
Automatic Bindings	0
Expired Bindings	0
Malformed Messages	0
DHCP DISCOVER packets discarded	0
Messages Received	
DHCPDISCOVER	0
DHCPREQUEST	23
DHCPDECLINE	0
DHCPRELEASE	0
DHCPINFORM	0
Messages Sent	
DHCPOFFER	0
DHCPACK	0
DHCPNAK	23
Refresh Clear Counters	

Figure 380: DHCP Server Statistics

Table 366: 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.
Messages Received	
This table shows statistical information about the messages received from DHCP clients on the network.	
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.
Messages Sent	
This table shows statistical information about messages the DHCP server has sent to DHCP clients on the network.	
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a Click **Clear Counters** to reset all DHCP statistics counters.

5.6.8 DHCP Server Conflicts Information

Use the DHCP Server Conflicts Information page to view information about IP address conflicts detected during the DHCP message exchange process between server and client. An address conflict occurs when two hosts on the same network use the same IP address. Any address detected as a duplicate is removed from the pool and will not be offered to any DHCP clients until the conflict is resolved.

To access the DHCP Server Conflicts Information page, click **Routing > DHCP Server > Conflicts** in the navigation menu.




Figure 381: DHCP Server Conflicts Information

Table 367: DHCP Server Conflicts InformationFields

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: - a 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. - a 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. - a 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.
Detection Time	The time when the conflict was detected in days, hours, minutes and seconds since the system was last reset (that is, system up time).

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen with the most current data.
- a Click **Clear Entries** to clear all of the address conflict entries.

5.7 Configuring DHCPv6

In this menu you can configure and view DHCPv6 parameters.

5.7.1 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 382: DHCPv6 Global Configuration

Table 368: 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.

Use the buttons to perform the following tasks:

- a If you change any of the settings, click **Submit** to apply the changes to the switch. To preserve the changes across a switch reboot, you must perform a save.
- a To update the information on the screen, click **Refresh**.
- a Click **Cancel** to discard changes and revert to the last saved state.

5.7.2 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 383: DHCPv6 Pool Summary

Table 369: 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a To add a pool, click **Add** and configure the pool information in the available fields.

Figure 384: Add DHCPv6 Server Pool**Table 370: Add DHCPv6 Server Pool Fields**

Field	Description
Pool Name	The name that identifies the DHCPv6 server pool.
Delegated Prefix	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.
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.

- a 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.
- a To remove a pool, select each entry to delete and click **Remove**. You must confirm the action before the pool is deleted.

5.7.3 DHCPv6 Pool Configuration

Use this page to edit pool settings or to configure additional settings for existing DHCPv6 pools.

5 Configuring Routing

To display the DHCPv6 Pool Configuration page, click **Routing > DHCPv6 > Pool Configuration** in the navigation menu.

Figure 385: 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:

- a To add a new binding to the selected DHCPv6 pool, click the + (plus) icon in the header row above the binding entries.
- a To remove all bindings from the selected pool, click the – (minus) icon in the header row above the binding entries.
- a To update the information for a binding, click the **Edit** icon associated with the binding.
- a To remove a binding from the selected pool, click the – (minus) icon associated with the binding.
- a 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.
- a 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.
- a To remove a single DNS or domain name entry, click the – (minus) icon associated with the entry to remove.

Table 371: 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.

5.7.4 DHCPv6 Interface Summary

Use this page to view the per-interface settings for DHCPv6.

To display the DHCPv6 Interface Summary page, click **Routing** > **DHCPv6** > **Interface** in the navigation menu.

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

Figure 386: DHCPv6 Interface Summary

Table 372: 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: - a None – The interface is not configured as a DHCPv6 server or DHCPv6 relay agent. - a Server – The interface responds to requests from DHCPv6 clients. - a 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.
Pool Name	(DHCPv6 server interface only) The name of the DHCPv6 pool the server uses to assign client information.
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a 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.

5.7.5 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 DHCPv6 Interface Configuration page, click **Routing** > **DHCPv6** > **Interface Configuration** in the navigation menu.

Figure 387: DHCPv6 Interface Configuration

Table 373: DHCPv6 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: - a None – The interface is not configured as a DHCPv6 server or DHCPv6 relay agent. - a Server – The interface responds to requests from DHCPv6 clients. - a Client – The interface initiates requests on a link to obtain configuration parameters from one or more DHCPv6 servers. - a 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.
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.
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.
Client Options	

Field	Description
(DHCPv6 client interface only) The information in this section can be configured only if the selected Interface Mode is Client.	
Prefix Delegation Client	When enabled, the interface can receive a general prefix for assignment to local router interfaces.
Rapid Commit	The mode of the rapid commit message exchange on the DHCPv6 client interface. The DHCPv6 client can obtain configuration parameters from a server either through a rapid two-message exchange (solicit, reply) or through a four-message exchange (solicit, advertise, request, and reply). When this option is disabled on either the client or server, the four-message exchange is used. When this option is enabled on both the client and the server, the two-message exchange is used.
Relay Options	
(DHCPv6 relay interface only) 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.

Use the buttons to perform the following tasks:

- a Click **Submit** to send the updated configuration to the switch.
- a Click **Refresh** to update the information on the screen.
- a Click **Cancel** to discard changes and revert to the last saved state.

5.7.6 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 DHCPv6 Binding Summary page, click **Routing** > **DHCPv6** > **Bindings** in the navigation menu.




Figure 388: DHCPv6 Binding Summary

Table 374: DHCPv6 Binding Summary Fields

Field	Description
Client IP Address	The IPv6 address associated with the client.

Field	Description
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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a 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.

5.7.7 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.

Global	Pool Summary	Pool Configuration	Interface	Interface Configuration	Bindings	Statistics	Conflicts

DHCPv6 Statistics						
Display 10 rows	Showing 1 to 10 of 55 entries				Filter:	
☐	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

First Previous 1 2 3 4 5 Next Last

Refresh Details Clear

Figure 389: DHCPv6 Statistics

Table 375: 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.

Field	Description
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 Transmitted	The number of DHCPv6 messages sent on 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.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the information on the screen.
- a To view detailed DHCPv6 statistics for an interface, select the entry with the information to view and click **Details**.

Interface	1/0/1
DHCPv6 Solicit Packets Received	0
DHCPv6 Request Packets Received	0
DHCPv6 Confirm Packets Received	0
DHCPv6 Renew Packets Received	0
DHCPv6 Rebind Packets Received	0
DHCPv6 Release Packets Received	0
DHCPv6 Decline Packets Received	0
DHCPv6 Inform Packets Received	0
DHCPv6 Relay-forward Packets Received	0
DHCPv6 Relay-reply Packets Received	0
DHCPv6 Malformed Packets Received	0
Received DHCPv6 Packets Discarded	0
Total DHCPv6 Packets Received	0
DHCPv6 Advertisement Packets Transmitted	0
DHCPv6 Reply Packets Transmitted	0
DHCPv6 Reconfig Packets Transmitted	0
DHCPv6 Relay-forward Packets Transmitted	0
DHCPv6 Relay-reply Packets Transmitted	0
Total DHCPv6 Packets Transmitted	0

Figure 390: Details

Table 376: Details 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.
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.
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.
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.

5 Configuring Routing

Field	Description
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.
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 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 Advertisement Packets Transmitted	The number of DHCPv6 Advertise 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 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.
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.

Field	Description
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.

- a To reset the DHCPv6 counters for one or more interfaces, select each interface with the statistics to reset and click **Clear**.

5.7.8 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 391: DHCPv6 Server Conflicts Information

Table 377: 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:

- a To update the information on the screen, click **Refresh**.
- a 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.
- a To remove all entries from the table, click **Clear All Entries**. You must confirm the action before all bindings are deleted.

6 Managing Device Security

Use the features in the Security menu to set management security parameters for port, user, and server security.

6.1 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:

- a **Authenticators** – Specifies the port that is authenticated before permitting system access.
- a **Supplicants** – Specifies host connected to the authenticated port requesting access to the system services.
- a **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 menu allows you to configure and view 802.1X features on the system.

6.1.1 Port Access Control Configuration

Use the Port Access Control Configuration page to enable or disable port access control on the system.

To display the Port Access Control Configuration page, click **Security > Port Access Control > Configuration** in the navigation menu.

Figure 392: Port Access Control Configuration

Table 378: Port Access Control Configuration Fields

Field	Description
Admin Mode	Select Enable or Disable 802.1x mode on the switch. The default is Disable . This feature permits port-based authentication on the switch.
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 is disabled.

Field	Description
Software Version	The version of 802.1x software running on the switch. It is not the 802.1x protocol version, but the software implementation version. The software version is set when the Admin Mode is enabled.

Use the buttons to perform the following tasks:

- a If you change the mode, click **Submit** to apply the new settings to the system.
- a Click **Refresh** to update the information on the screen.
- a Click **Cancel** to discard changes and revert to the last saved state.

6.1.2 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.

Interface	Interface Status
☐ 1/0/1	N/A
☐ 1/0/2	N/A
☐ 1/0/3	N/A
☐ 1/0/4	N/A
☐ 1/0/5	N/A
☐ 1/0/6	N/A
☐ 1/0/7	N/A
☐ 1/0/8	N/A
☐ 1/0/9	N/A
☐ 1/0/10	N/A

Figure 393: Port Access Control Port Summary

Table 379: Port Access Control Port Summary Fields

Field	Description
Interface	The interface associated with the rest of the data in the row.
Interface Status	The authorization status of the port, which is one of the following: - a Unauthorized - a Authorized - a N/A

Use the buttons at the bottom of the page to perform the following actions:

- a Click **Refresh** to update the page with the most current information.
- a 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.

- a 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.

6.1.3 Port Access Control 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 Access Control Port Configuration page, click **Security > Port Access Control > Port Configuration** in the navigation menu.

Figure 394: Port Access Control Port Configuration

Table 380: 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.
Protocol Version	The protocol version associated with this port. The only possible value is 1, corresponding to the first version of the Dot1x specification.
PAE Capabilities	The Port Access Entity (PAE) role, which is one of the following: - a 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. - a 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.

Field	Description
	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.
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).	
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.
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.
Maximum Request Identity	The maximum number of times the port will retransmit an EAP Request-Identity frame before timing out the supplicant.
Key Transmission	Indicates if the key is transmitted to the supplicant for the specified port.
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: - a Force Unauthorized – The port ignores supplicant authentication attempts and does not provide authentication services to the client. - a Force Authorized – The port sends and receives normal traffic without client port-based authentication. - a Auto – The port is unauthorized until a successful authentication exchange has taken place.
Supplicant PACP State	Current state of the supplicant PACP state machine, which is one of the following: - a Initialize - a Logoff - a Held - a Unauthenticated - a Authenticating - a Authenticated
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.

Field	Description
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.

Use the buttons to perform the following tasks:

- a To save your changes to the changed interface, click **Submit**.
- a Click **Refresh** to update the information on the screen.
- a Click **Cancel** to discard changes and revert to the last saved state.

6.1.4 Port Access Control 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.

The screenshot shows the 'Port Access Control Port Details' page. At the top, there are tabs: Configuration, Port Summary, Port Configuration, Port Details (selected), Statistics, and Privileges Summary. Below the tabs, the page title 'Port Access Control Port Details' is displayed. The main content area shows the following configuration:

Interface	1/0/1
PAE Capabilities	Authenticator
Authenticator Options	
Quiet Period (Seconds)	60
Transmit Period (Seconds)	30
Supplicant Timeout (Seconds)	30
Server Timeout (Seconds)	30
Maximum Requests	2
Maximum Request Identity	2
Key Transmission	Disabled

At the bottom right of the configuration area, there is a 'Refresh' button.

Figure 395: Port Access Control Port Details

Table 381: 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: - a 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. - a 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.
Authenticator Options	

Field	Description
The fields in this section provide information about the settings that apply to the port when it is configured as an 802.1X authenticator.	
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.
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.
Maximum Request Identity	The maximum number of times the port will retransmit an EAP Request-Identity frame before timing out the supplicant.
Key Transmission	Indicates if the key is transmitted to the supplicant for the specified port.
Supplicant Options	
The fields in this section are displayed 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: - a 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. - a Force Unauthorized – The port is placed into an unauthorized state and is automatically denied system access. - a 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.
Supplicant PACP State	Current state of the supplicant PACP state machine, which is one of the following: - a Initialize - a Logoff - a Held - a Unauthenticated - a Authenticating - a Authenticated
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.

Click **Refresh** to update the information on the screen.

6.1.5 Port Access Control 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 access the Port Access Control Statistics page, click **Security > Port Access Control > Statistics** in the navigation menu.

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

Figure 396: Port Access Control Statistics

Table 382: 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: - a 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. - a Supplicant – The port must be granted permission by the authentication server before it can access the remote authenticator port.
EAPOL Frames Received	The total number of valid EAPOL frames received on the interface.
EAPOL Frames Transmitted	The total number of EAPOL frames sent by the interface.
Last EAPOL Frame Version	The protocol version number attached to the most recently received EAPOL frame.
Last EAPOL Frame Source	The source MAC address attached to the most recently received EAPOL frame.

Use the buttons at the bottom of the page to perform the following actions:

- a Click **Refresh** to update the information on the screen.

- a To view additional per-interface EAPOL and EAP message statistics, select the interface with the information to view and click **Details**.



Details	
Interface	1/0/1
PAE Capabilities	Authenticator
EAPOL Frames Received	0
EAPOL Frames Transmitted	0
EAPOL Start Frames Received	0
EAPOL Logoff Frames Received	0
Last EAPOL Frame Version	0
Last EAPOL Frame Source	00:00:00:00:00:00
EAP Response/ID Frames Received	0
EAP Response Frames Received	0
EAP Request/ID Frames Transmitted	0
EAP Request Frames Transmitted	0
Invalid EAPOL Frames Received	0
EAPOL Length Error Frames Received	0

Figure 397: Details

Table 383: Port Access Control Statistics Fields (Details View)

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: a Authenticator a Supplicant
EAPOL Frames Received	The total number of valid EAPOL frames received on the interface.
EAPOL Frames Transmitted	The total number of EAPOL frames sent by the interface.
EAPOL Start Frames Received	(Only for PAE capability Authenticator) 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.
EAPOL Logoff Frames Received	(Only for PAE capability Authenticator) 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.
EAP Response/ID Frames Received	(Only for PAE capability Authenticator) 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.
EAP Response Frames Received	(Only for PAE capability Authenticator) 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.
EAP Request/ID Frames Transmitted	(Only for PAE capability Authenticator) 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.
EAP Request Frames Transmitted	(Only for PAE capability Authenticator) The total number of EAP-Request frames the interface has sent. EAP-Request frames are sent from an

Field	Description
	authentication server to a supplicant (and translated by the authenticator) during the authentication process.
EAPOL Start Frames Transmitted	(Only for PAE capability Supplicant) 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.
EAPOL Logoff Frames Transmitted	(Only for PAE capability Supplicant) 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.
EAP Response/ID Frames Transmitted	(Only for PAE capability Supplicant) 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.
EAP Response Frames Transmitted	(Only for PAE capability Supplicant) The total number of EAP-Response frames the interface has sent. EAP-Response frames are sent from a supplicant to an authentication server during the authentication process.
EAP Request/ID Frames Received	(Only for PAE capability Supplicant) 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.
EAP Request Frames Received	(Only for PAE capability Supplicant) 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.
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.

- a Click **Clear** to reset all statistics counters to 0 for the selected interface or interfaces.

6.1.6 Port Access Control Privileges Summary

Use this page to grant or deny port access to users configured on the system (see [User Accounts](#) on page 61).

To access the Port Access Control Privileges Summary page, click **Security > Port Access Control > Privileges Summary**

in the navigation menu.

Configuration Port Summary Port Configuration Port Details Statistics **Privileges Summary**

Port Access Control Privileges Summary

Display **10** rows Showing 1 to 10 of 54 entries Filter:

Interface	Users
☐ 1/0/1	admin
☐ 1/0/2	admin
☐ 1/0/3	admin
☐ 1/0/4	admin
☐ 1/0/5	admin
☐ 1/0/6	admin
☐ 1/0/7	admin
☐ 1/0/8	admin
☐ 1/0/9	admin
☐ 1/0/10	admin

First Previous 1 2 3 4 5 Next Last

Refresh Edit

Figure 398: Port Access Control Privileges Summary

Table 384: Port Access Control Privileges Summary Fields

Field	Description
Interface	The local interface associated with the rest of the data in the row.
Users	The users that are allowed access to the system through the associated port.

Use the buttons at the bottom of the page to perform the following actions:

- a Click **Refresh** to update the information on the screen.
- a 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.

Privileges Configuration

Configurable Interfaces

Interface 1/0/1

Users

Available Users

Selected Users

admin

Submit Cancel

Figure 399: Privileges Configuration

Table 385: Privileges Configuration Fields

Field	Description
Interface	When configuring access information for one or more interfaces, this field identifies each interface being configured.

Field	Description
Users	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 CTRL + click to select multiple users) and click the appropriate arrow.

6.2 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:

- a Telnet Access
- a Web Access
- a Console to Switch Access
- a Port Access Control (802.1X)

The RADIUS menu allows you to configure and view RADIUS settings.

6.2.1 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 400: RADIUS Configuration

Table 386: 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

Field	Description
	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.
MAB Attribute	The RADIUS attribute 1 (User-Name) for sending MAC-based Authentication Bypass (MAB) requests from the client to the RADIUS server. 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. 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. After you click Edit, the Set MAB Attribute window appears and includes the following fields: - a 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. - a In the following example, the group size is 1: 0:0:1:0:1:8:9:9:F:2:B:3 - a In the following example, the group size is 2: 00:10:18:99:F2:B3 - a In the following example, the group size is 4: 0010:1899:F2B3 - a In the following example, the group size is 12: 00101899F2B3 - a Separator – The separator used by the switch to format the RADIUS attribute 1 (User-Name) of the MAB request. - a In the following example, the separator is - (hyphen): 00-10-18-99-F2-B3 - a In the following example, the separator is : (colon): 00:10:18:99:F2:B3 - a Case – The case of any letters used by the switch to format the RADIUS attribute 1 (User-Name) of the MAB request. - a In the following example, the case is lowercase: 00:d0:18:99:f2:b3 - a In the following example, the case is uppercase: 00:D0:18:99:F2:B3
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)	The network access server (NAS) IP address for the RADIUS server. 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.

Field	Description
CALLED-STATION-ID (Attribute 30)	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: - a Legacy Lowercase – Format the MAC address as xx:xx:xx:xx:xx:xx - a Legacy Uppercase – Format the MAC address as XX:XX:XX:XX:XX:XX - a IETF Lowercase – Format the MAC address as xx-xx-xx-xx-xx-xx - a IETF Uppercase – Format the MAC address as XX-XX-XX-XX-XX-XX - a Unformatted Lowercase – Format the MAC address as aaaabbbbcccc - a Unformatted Uppercase – Format the MAC address as AAAABBBBCCCC
CALLING-STATION-ID (Attribute 31)	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: - a Legacy Lowercase – Format the MAC address as xx:xx:xx:xx:xx:xx - a Legacy Uppercase – Format the MAC address as XX:XX:XX:XX:XX:XX - a IETF Lowercase – Format the MAC address as xx-xx-xx-xx-xx-xx - a IETF Uppercase – Format the MAC address as XX-XX-XX-XX-XX-XX - a Unformatted Lowercase – Format the MAC address as aaaabbbbcccc - a Unformatted Uppercase – Format the MAC address as AAAABBBBCCCC
NAS-IDENTIFIER (Attribute 32 MAC Format)	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: - a Legacy Lowercase – Format the MAC address as xx:xx:xx:xx:xx:xx - a Legacy Uppercase – Format the MAC address as XX:XX:XX:XX:XX:XX - a IETF Lowercase – Format the MAC address as xx-xx-xx-xx-xx-xx - a IETF Uppercase – Format the MAC address as XX-XX-XX-XX-XX-XX - a Unformatted Lowercase – Format the MAC address as aaaabbbbcccc - a Unformatted Uppercase – Format the MAC address as AAAABBBBCCCC
NAS-IDENTIFIER (Attribute 32 Include in Access/Accounting Request)	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: - a Include in Access/Accounting Request – When selected, the attribute is sent to the RADIUS server in access-request and accounting-request messages. - a 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: - a %m – MAC address - a %i – IP address - a %h – Host name - a %d – Domain name - a Any String – A string including any or all of the above formatting options

Field	Description
	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.
ACCT-SESSION-ID (Attribute 44)	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.
NAS-IPV6-ADDRESS (Attribute 95)	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.

Use the buttons at the bottom of the page to perform the following actions:

- a If you make changes to the page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to update the page with the most current information.
- a Click **Cancel** to discard changes and revert to the last saved state.

6.2.2 RADIUS 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 401: RADIUS Named Server Status

Table 387: RADIUS Named Server Status Fields

Field	Description
Current	Indicates whether the RADIUS server is the current server (True) or a backup server (False) within its group. If more than one RADIUS server is configured with the same Server Name, the switch selects one of the servers to be the current server in the named server group. 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. If no server is configured as the primary server, the current server is the RADIUS server that is added to the group first.

Field	Description
IP Address/Host Name	The IP address or host name of the RADIUS server. Host names must be resolvable by DNS and are composed of a series of labels separated by dots.
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	Indicates whether the server is the Primary or a Secondary RADIUS authentication server. When multiple RADIUS servers have the same Server Name value, the RADIUS client attempts to use the primary server first. If the primary server does not respond, the RADIUS client attempts to use one of the backup servers within the same named server group.
Secret Configured	Indicates whether the shared secret for this server has been configured.
Message Authenticator	Indicates whether the RADIUS server requires the Message Authenticator attribute to be present. The Message Authenticator adds protection to RADIUS messages by using an MD5 hash to encrypt each message. The shared secret is used as the key, and if the message fails to be verified by the RADIUS server, it is discarded.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the page with the most current information.
- a To add a RADIUS authentication server to the list of servers the RADIUS client can contact, click **Add**.

Figure 402: Add RADIUS Server

Table 388: Add RADIUS Server Fields

Field	Description
IP Address/Host Name	Enter an IP address or host name of the RADIUS server. Host names must be resolvable by DNS and are composed of a series of labels separated by dots.
Server Name	Optionally you can modify 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	Optionally you can modify the default RADIUS port 1812 if your RADIUS server uses a different port.
Secret	Enter a shared secret to be used for authenticating and encrypting all RADIUS communications between the RADIUS client on the switch and the RADIUS server. The secret specified in this field must match the shared secret configured on the RADIUS server.

Field	Description
Server Type	Choose if the server should be the Primary or a Secondary RADIUS authentication server. When multiple RADIUS servers have the same Server Name value, the RADIUS client attempts to use the primary server first. If the primary server does not respond, the RADIUS client attempts to use one of the backup servers within the same named server group.
Message Authenticator	Choose if the RADIUS server should require the Message Authenticator attribute to be present. The Message Authenticator adds protection to RADIUS messages by using an MD5 hash to encrypt each message. The shared secret is used as the key, and if the message fails to be verified by the RADIUS server, it is discarded.

- a 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.
- a 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.

6.2.3 RADIUS 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 403: RADIUS Server Statistics


Table 389: 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.

Field	Description
Packets Dropped	The number of RADIUS packets received from the server on the authentication port and dropped for some other reason.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the page with the most current information.
- a After you click **Details**, a window opens and displays additional statistics about the number and type of messages sent between the selected RADIUS server and the RADIUS client on the device.



IP Address/Host Name	192.168.45.254
Round Trip Time	0
Access Requests	0
Access Retransmissions	0
Access Accepts	0
Access Rejects	0
Access Challenges	0
Malformed Access Responses	0
Bad Authenticators	0
Pending Requests	0
Timeouts	0
Unknown Types	0
Packets Dropped	0

Figure 404: RADIUS Server Detailed Statistics

Table 390: RADIUS Server Detailed 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 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 Rejects	The number of RADIUS Access-Reject 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.

Field	Description
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.
Unknown Types	The number of RADIUS packets of unknown type which were received from the server on the authentication port.
Packets Dropped	The number of RADIUS packets received from the server on the authentication port and dropped for some other reason.

6.2.4 RADIUS Accounting Server Status

The RADIUS 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 405: RADIUS Accounting Server Status

Table 391: 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.

Use the buttons to perform the following tasks:

- Click **Refresh** to update the page with the most current information.

- a To add a RADIUS accounting server to the list of servers the RADIUS client can contact, click **Add**.




Figure 406: Add RADIUS Accounting Server

Table 392: Add RADIUS Accounting Server Fields

Field	Description
IP Address/Host Name	Enter an 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	Optionally you can modify 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	Optionally you can modify the default RAIDUS accounting server port 1813 if your server uses a different port.
Secret	Enter a shared secret to be used for authenticating and encrypting all RADIUS communications between the RADIUS client on the switch and the RADIUS accounting server. The secret specified in this field must match the shared secret configured on the RADIUS accounting server.

- a 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.
- a 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.

6.2.5 RADIUS Accounting Server Statistics

Use the RADIUS Accounting Server Statistics page to view statistical information for each RADIUS server configured on the system.

To access the RADIUS Accounting Server Statistics page, click **Security > RADIUS > Accounting Statistics** in the navigation menu.




Figure 407: RADIUS Accounting Server Statistics

Table 393: RADIUS Accounting Server 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.

Use the buttons to perform the following tasks:

- a To update the information on the screen, click **Refresh**.
- a To view additional statistics, select the RADIUS accounting server with the statistics to view and click **Details**.

IP Address/Host Name	192.168.45.254
Round Trip Time	0
Accounting Requests	0
Accounting Retransmissions	0
Accounting Responses	0
Malformed Access Responses	0
Bad Authenticators	0
Pending Requests	0
Timeouts	0
Unknown Types	0
Packets Dropped	0

Figure 408: RADIUS Accounting Server Detailed Statistics


Table 394: RADIUS Accounting Server Detailed 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.
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.
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.
Timeouts	The number of times a response was not received from the server within the configured timeout value.
Unknown Types	The number of RADIUS packets of unknown type which were received from the server on the accounting port.
Packets Dropped	The number of RADIUS packets received from the server on the accounting port and dropped for some other reason.

6.2.6 RADIUS 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 409: 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.

6.2.7 RADIUS 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 410: RADIUS Source Interface Configuration

Table 395: RADIUS Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: - a None – The primary IP address of the originating (outbound) interface is used as the source address. - a Interface – The primary IP address of a physical port is used as the source address. - a VLAN – The primary IP address of a VLAN routing interface is used as the source address. - a 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. - a Network – The network source IP is used as the source address. - a Service Port – The management port source IP is used as the source address. 5 The option Service Port is only available on switches which have a Service Port.
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.
Loopback Interface	When the selected Type is Loopback , select the loopback interface to use as the source interface.

Use the buttons to perform the following tasks:

- a Click **Submit** to send the updated configuration to the switch.
- a Click **Refresh** to update the page with the most current information.
- a Click **Cancel** to discard changes and revert to the last saved state.

6.3 TACACS+ Configuration

To access the TACACS+ Configuration page, click **Security > TACACS+ > Configuration** in the navigation menu.




Figure 411: TACACS+ Configuration

Table 396: 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.

Use the buttons to perform the following tasks:

- a If you make any changes to the page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to update the page with the most current information.
- a Click **Cancel** to discard changes and revert to the last saved state.

6.3.1 TACACS+ Server Summary

Use this page to view and configure information about the TACACS+ Servers.

To access the TACACS+ Server Summary page, click **Security > TACACS+ > Server Summary** in the navigation menu.




Figure 412: TACACS+ Server Summary

Table 397: TACACS+ Server Summary Fields

Field	Description
Server	Specifies the TACACS+ Server IP address or Hostname.

Field	Description
Priority	Specifies the order in which the TACACS+ servers are used. The value 0 has the highest priority. If two or more entries have the same priority the first entry is used.
Port	Specifies the authentication port.
Connection Timeout	The amount of time that passes before the connection between the device and the TACACS+ server timeout.

Use the buttons to perform the following tasks:

- a Click **Refresh** to update the page with the most current information.
- a To add a TACACS+ Server to the list of servers the TACACS+ client can contact, click **Add**. If the maximum number of 5 servers has been reached, the button will be disabled.

Figure 413: Add TACACS+ Server

Table 398: Add TACACS+ Server Fields

Field	Description
Server	Specifies the TACACS+ Server IP address or Hostname.
Priority	Specifies the order in which the TACACS+ servers are used. The value 0 has the highest priority. If two or more entries have the same priority the first entry is 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 timeout.

- a To edit a configured TACACS+ server from the list, select the entry and click **Edit**. You will be forwarded to the Server Configuration menu.
- a 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.

6.3.2 TACACS+ Server Configuration

Use this page to view and configure information about the TACACS+ Servers.

To access the TACACS+ Server Configuration page, click **Security > TACACS+ > Server Configuration** in the navigation menu.

Field	Value	Range/Notes
Server	172.23.54.80	
Priority	0	(0 to 65535)
Port	49	(0 to 65535)
Key String		
Connection Timeout	5	(1 to 30 secs)

Buttons: Submit, Remove, Refresh, Cancel

Figure 414: TACACS+ Server Configuration

Table 399: 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 (the default port is TCP 49).
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 timeout.

Use the buttons to perform the following tasks:

- a If you make any changes to the page, click **Submit** to apply the changes to the system.
- a To remove a TACACS+ server, select each entry to delete and click **Remove**. You must confirm the action before the server is deleted.
- a To update the information on the screen, click **Refresh**.
- a Click **Cancel** to discard changes and revert to the last saved state.

6.3.3 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 415: TACACS+ Source Interface Configuration

Table 400: TACACS+ Source Interface Configuration Fields

Field	Description
Type	The type of interface to use as the source interface: - a None – The primary IP address of the originating (outbound) interface is used as the source address. - a Interface – The primary IP address of a physical port is used as the source address. - a VLAN – The primary IP address of a VLAN routing interface is used as the source address. - a 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. - a Network – The network source IP is used as the source address. - a Service Port – The management port source IP is used as the source address. 5 The option Service Port is only available on switches with a Service Port.
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.
Loopback Interface	When the selected Type is Loopback , select the loopback interface to use as the source interface.

Use the buttons to perform the following tasks:

- a If you make any changes to the page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to update the page with the most current information.
- a Click **Cancel** to discard changes and revert to the last saved state.

6.4 Authentication Manager

The Authentication Manager feature allows you to configure the authentication methods used on the individual interface.

6.4.1 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:

- a Change of the client VLAN
- a Client re-authentication
- a Change of the Filter-ID for client
- a Change of the Downloadable Access Control List (DACL) for client

The following CoA requests result in termination of all sessions on the matching port:

- a Bounce port
- a Disable port

To access the **Authentication Manager Configuration** page, click **Security > Authentication Manager > Configuration** in the navigation menu.

Field	Value
Admin Mode	☒ Disable ☐ Enable
Dynamic VLAN Creation Mode	☒ Disable ☐ Enable
VLAN Assignment Mode	☒ Disable ☐ Enable
Authentication Monitor Mode	☒ Disable ☐ Enable
Critical Recovery Max Re-Authentication	10 (1 to 50)
CoA Bounce Host Port	☒ Accept ☐ Reject
CoA Disable Host Port	☒ Accept ☐ Reject
Authenticated Clients	0
Clients in Monitor Mode	0

Submit Refresh Cancel

Figure 416: Authentication Manager Configuration

Table 401: 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. 5 As the Control Mode in the Authentication Manager Interface Configuration on page 467 is set to Auto for all interfaces, authentication is active for each interface. Thus you will lose access to the switch if the Admin Mode is set to Enable without configuring the interfaces and authentication parameters first. It is therefore recommended to set the management interface to Force Authorized.
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.

Field	Description
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.
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 conditions. 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 <i>CoA disabled</i> . The disabled port can be re-enabled using one of the following methods: - a 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 Port Auto Recovery Configuration on page 253. - a The administrator manually re-enables the port on the Port Summary page. See Port Summary on page 113.
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:

- a 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.
- a Click **Refresh** to display the latest information from the switch.
- a Click **Cancel** to cancel the change.

6.4.2 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 417: Authentication Manager Interface Configuration

Table 402: 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: - a Force Unauthorized – The port ignores supplicant authentication attempts and does not provide authentication services to the client. - a Force Authorized – The port sends and receives normal traffic without client port-based authentication. - a Auto – The port is unauthorized until a successful authentication exchange has taken place.
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: a Single Authentication – Only one data client can be authenticated on a port and the client is granted access to the port.

Field	Description
	- a 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. - a Multiple Domain – One data client and one voice client can be authenticated on a port and both clients are granted access to the port. - a Multiple Authentication – One voice client and multiple data clients can be authenticated on a port and these clients are granted access to the port. - a 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.
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	When the checkbox is ticked, the Re-Authentication Period can be set. It specifies 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: - a Protect - a Restrict - a Shutdown
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: - a None – No action is configured. - a Reinitialize – Dot1x triggers the re-authentication of clients authenticated on the critical VLAN.
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: - a None – No action is configured. - a Authorize – Allows port access on the voice VLAN when all RADIUS servers are dead.

6 Managing Device Security

Field	Description
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: - a None – No action is configured. - a 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. - a 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.
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.
MAB Authentication Type	The authentication type to be used for MAB access requests sent to the RADIUS server, which is one of the following: - a 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. - a 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. - a 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.
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:

- a 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.
- a Click **Refresh** to display the latest information from the switch.
- a Click **Cancel** to cancel the change.

6.4.3 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 and MAC Authentication Bypass (MAB).

To access the **Authentication Tiering** page, click **Security > Authentication Manager > Authentication Tiering** in the navigation menu.




Figure 418: Authentication Tiering

Table 403: 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: - a Dot1x – The port-based authentication method. - a MAB – MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide. - a Captive Portal – The authentication method that prevents clients from accessing the network until user verification has been established. 5 The method 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.

Field	Description
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:

- a Click **Refresh** to display the latest information from the switch.
- a Click **Edit** to configure the settings for one or more interfaces, select each entry to modify. The settings are applied to all selected interfaces.

Table 404: Edit 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.
Re-Authentication Timer	Interval, in seconds, after which an attempt is made to authenticate an unauthorized port.
Configured Method Order	Modify the order by moving the methods from the Methods field to the Order field by using the appropriate arrow buttons.
Configured Method Priority	Modify the priority by moving the methods from the Methods field to the Order field by using the appropriate arrow buttons.

Figure 419: Edit Authentication Tiering

6.4.4 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 420: Authenticated Clients

Table 405: 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: - a Dot1x – The port-based authentication method. - a MAB – MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide. - a None – The authentication method is undefined. - a Captive Portal – The authentication method that prevents clients from accessing the network until user verification has been established.
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: - a 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. - a 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. - a Multiple Domain – One data client and one voice client can be authenticated on a port, and both clients are granted access to the port. - a Multiple Authentication – One voice client and multiple data clients can be authenticated on a port, and these clients are granted access to the port. - a 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.
Control Mode	The authentication control mode on the port, which is one of the following: - a Force Unauthorized – The port ignores client authentication attempts and does not provide authentication services to the client. - a Force Authorized – The port sends and receives normal traffic without client port-based authentication. - a Auto – The port is unauthorized until a successful authentication exchange has taken place.

Field	Description
VLAN ID	The VLAN ID 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: - a Default Assigned VLAN – The client is authenticated on the port in the default VLAN, and the authentication server is not RADIUS. - a RADIUS Assigned VLAN – RADIUS is used for authenticating the client. - a Unauthenticated VLAN – The client is authenticated on the unauthenticated VLAN. - a Guest VLAN – The client is authenticated on the guest VLAN. - a Voice VLAN – The client is authenticated on the voice VLAN. - a Monitor Mode VLAN – The client is authenticated via monitor mode. - a Critical VLAN – The client is authenticated on the critical VLAN. - a Not Assigned – No VLAN is assigned to the port.

Use the buttons to perform the following tasks:

- a Click **Refresh** to display the latest information from the switch.
- a 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 406: Authenticated Clients Details Fields

Field	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 when 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.
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.

6.4.5 Authentication 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.

Interface	Dot1x Attempts	Dot1x Failures	MAB Attempts	MAB Failures
1/0/1	0	0	0	0
1/0/2	0	0	0	0
1/0/3	0	0	0	0
1/0/4	0	0	0	0
1/0/5	0	0	0	0
1/0/6	0	0	0	0
1/0/7	0	0	0	0
1/0/8	0	0	0	0
1/0/9	0	0	0	0
1/0/10	0	0	0	0

Figure 421: Authentication Statistics

Table 407: 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.

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** and acknowledge the prompt to reset all statistics counters to 0 for the selected interfaces.

6.4.6 Authentication 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 422: Authentication History

Table 408: 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: - a Authorized – Indicates client is authorized on the port. - a Unauthorized – Indicates client is not authorized on the port.
Authenticated Method	The authentication method used to authenticate a client connected to an interface, which can be one of the following: - a Dot1x – The port-based authentication method. - a MAB – MAC Authentication Bypass method that uses the MAC address of the client to determine the kind of network access to provide. - a Captive Portal – The authentication method that prevents clients from accessing the network until user verification has been established. - a None - The authentication method is undefined.
VLAN Assigned	The VLAN ID where the client is placed as a result of the authentication process.
VLAN Assigned Reason	The reason why the authenticator placed in the client in the VLAN, which can be one of the following: - a Default Assigned VLAN – The client is authenticated on the port in the default VLAN, and the authentication server is not RADIUS. - a RADIUS Assigned VLAN – RADIUS is used for authenticating the client. - a Unauthenticated VLAN – The client is authenticated on the unauthenticated VLAN. - a Guest VLAN – The client is authenticated on the guest VLAN. - a Voice VLAN – The client is authenticated on the voice VLAN. - a Monitor Mode VLAN – The client is authenticated via monitor mode. - a Critical VLAN – The client is authenticated on the critical VLAN. - a Not Assigned – No VLAN is assigned to the port.

Field	Description
Reason	The reason for the successful or unsuccessful authentication.

Use the buttons at the bottom of the page to perform the following actions:

- a Click **Refresh** to display the latest information from the switch.
- a Click **Clear** to clear the Authentication Manager history log on the selected interface.

7 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.

4

Some of the features described in this section may not be supported in LCOS SX software releases for particular hardware platforms.

7.1 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. LCOS SX software supports IPv4, IPv6, and MAC ACLs. The total number of MAC and IP ACLs supported by LCOS SX 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.

7.1.1 IP Access Control Lists

IP ACLs allow network managers to define classification actions and rules for specific ports. ACLs are composed of access control entries (ACEs), 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.

You can configure and view IP ACLs in the menu IP Access Control List. To configure an IP ACL:

1. Use the [Access Control List Summary](#) page to define the IP ACL type and assign an ID to it.
2. Use the [Access Control List Interface Summary](#) page to create rules for the ACL.

7.1.1.1 Access Control List Summary

Use the Access Control List Summary 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 and created using the [Access Control List Interface Summary](#) page.

The Summary page displays the dynamic ACLs with an additional tag #d appended to the ACL identifier.

To display the Access Control List Summary page, click **QoS > Access Control Lists > Summary** in the navigation menu.

Figure 423: Access Control List Summary

Use the buttons at the bottom of the page to perform the following tasks:

- a To add an ACL, click **Add** and configure the ACL type and ACL Identifier.
- a 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.
- a 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 409: Access Control 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.
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: a IPv4 Standard – Match criteria is based on the source address of IPv4 packets.

Field	Description
	- a 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. - a IPv4 Named – Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. - a 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. - a 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.
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 interfaces to which the ACL has been applied.
VLAN	Each VLAN to which the ACL has been applied.

7.1.1.2 Access Control List Configuration

Use this page to configure rules for the existing 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 for the specified rule action (Permit/Deny) to take place.

The Access Control List Configuration page displays the dynamic ACLs with an additional tag **#d** appended to the ACL Identifier in the **ACL Identifier** list.

To display the Access Control List Configuration page, click **QoS > Access Control Lists > Configuration** in the navigation menu.

Figure 424: Access Control List Configuration

Use the buttons to perform the following tasks:

- a 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.
- a To remove the most recently configured rule for an ACL, select the rule in the list and click **Remove Rule**. You must confirm the action before the entry is deleted.
- a To resequence rules for an ACL, select the ID of the appropriate ACL from the ACL Identifier menu and click Resequence Rules.

Table 410: Access Control List Configuration Fields

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 425 shows the 'Edit Access Control List' dialog box for an IPv4 Named ACL. The 'ACL Type' is 'IPv4 Named'. The 'ACL Identifier' is 'TEST', with a note indicating it must be 1 to 255 alphanumeric characters. The dialog includes 'Submit' and 'Cancel' buttons.

Figure 425: Edit Access Control List

Figure 426 shows the 'Edit Access Control List' dialog box for an IPv6 Named ACL. The 'ACL Type' is 'IPv6 Named'. The 'ACL Identifier' is 'TEST', with a note indicating it must be 1 to 255 alphanumeric characters. The dialog includes 'Submit' and 'Cancel' buttons.

Figure 426: 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 427: Edit Access Control List - Dynamic ACLs

Table 411: 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: - a IPv4 Standard – Match criteria is based on the source address of IPv4 packets. - a 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. - a IPv4 Named – Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. - a 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. - a 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.
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: - a Permit – The packet or frame is forwarded. - a Deny – The packet or frame is dropped.

Field	Description
	4 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.
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.

Use the buttons available in the **ACL Remarks** table to perform the following tasks:

- a To add a remark, click the + (plus) button and enter the remark to add.
- a 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 412: 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 IPv4 ACL 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 information describes the fields in this window. The Match Criteria tables that apply to IPv6 ACLs, and MAC ACLs are described separately.

Add IPv4 ACL Rule

Sequence Number

(1 to 2147483647)

Action

Permit

Deny

Match Criteria

Every

Protocol

(0 to 255, or keyword) ?

Fragments

Source IP Address / Wildcard Mask

/

(x.x.x.x)

Source L4 Port

Equal

Not Equal

Less Than

Greater Than

Range

-

(0 to 65535, or keyword) ?

Destination IP Address / Wildcard Mask

/

(x.x.x.x)

Destination L4 Port

Equal

Not Equal

Less Than

Greater Than

Range

-

(0 to 65535, or keyword) ?

TTL Field Value

(0 to 255)

IGMP Type

(0 to 255)

ICMP Type

(0 to 255)

ICMP Code

(0 to 255)

ICMP Message

TCP Flags

+FIN

-FIN

+SYN

-SYN

+RST

-RST

+PSH

-PSH

Service Type

IP DSCP

(0 to 63, or keyword) ?

IP Precedence

(0 to 7)

IP TOS Bits / Wildcard Mask

/

(0 to FF hex)

Rule Attributes

Assign Queue

(0 to 6)

Interface

Redirect

Mirror

Log

Redirect External Agent

(1 to 100)

Time Range Name

(1 to 31 alphanumeric characters)

Committed Rate / Burst Size

(1 to 4294967295)

(1 to 128)

Submit

Cancel

Figure 428: Add IPv4 ACL Rule

Table 413: 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. See Figure 426: Edit Access Control List on page 481.
Action	Select the option to permit or deny associating the sequence number to the ACL rule.

Field	Description
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.

7 Configuring Quality of Service

Field	Description
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.
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 services types are as follows: - a 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. - a 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. - a 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. - a TOS Bits – Requires the bits in a packet's TOS field to match the two-digit hexadecimal number entered in this field. - a TOS Mask – The bit positions that are used for comparison against the IP TOS field in a packet.
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

Field	Description
	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.
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.

7 Configuring Quality of Service

Field	Description
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).
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: - a 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. - a 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.
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.

Field	Description
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 414: 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 429: Add ACL Remark

7.1.1.3 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 430: Access Control List Interface Summary

Use the buttons to perform the following tasks:

- a To apply an ACL to an interface, click **Add** and configure the settings in the available fields.
- a 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 415: 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.

7.1.1.4 Access Control List VLAN Summary

Use this page to associate one or more ACLs with one or more VLANs on the device.

4 You can also associate an ACL with a VLAN routing interface.

To display the Access Control List VLAN Summary page, click **QoS > Access Control Lists > VLANs** in the navigation menu.




Figure 431: Access Control List VLAN Summary

Use the buttons to perform the following tasks:

- a To associate an ACL with a VLAN, click **Add** and configure the settings in the available fields.
- a 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 416: Access Control List VLAN Summary Fields

Field	Description
VLAN ID	The VLAN ID 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.

7.1.1.5 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 432: Access Control List Control Plane Configuration

Use the buttons to perform the following tasks:

- a To apply an ACL to the CPU interface, click **Add** and configure the settings in the available fields.
- a 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 417: 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: - a IPv4 Standard – Match criteria is based on the source address of IPv4 packets. - a 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. - a IPv4 Named – Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. - a 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. - a 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.
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.

7.1.1.6 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 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 433: Access Control List Statistics

Use the buttons to perform the following tasks:

- a 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 rules.
- a 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.
- a 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 418: Access Control List Statistics Fields

Field	Description
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: - a IPv4 Standard – Match criteria is based on the source address of the IPv4 packets. - a 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. - a IPv4 Named – Match criteria is the same as IPv4 Extended ACLs, but the ACL ID can be an alphanumeric name instead of a number. - a 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. - a 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.

Field	Description
ACL Identifier	A list of ACL IDs that exist on the system for a given ACL type. To view the rules 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.
Action	The action to take when a packet or frame matches the criteria in the rule: a Permit – The packet or frame is forwarded. a Deny – The packet or frame is dropped.
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.

7.1.1.7 IPv6 ACL Rules

The maximum number of IPv6 rules depends on the following factors:

- a 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.
- a 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.
- a 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.

7.1.2 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:

- a Session Initiation Protocol (SIP)
- a H.323
- a 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.

7.1.2.1 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 434: Auto VoIP Global Configuration

Table 419: 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 after you confirmed the reset procedure.

7.1.2.2 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 435: OUI Table Summary

Use the buttons to perform the following tasks:

- a To add an OUI, click **Add** and specify an OUI and its description in the available fields.
- a 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 420: 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.

7.1.2.3 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.

Global OUI Table **OUI Based Auto VoIP** Protocol Based Auto VoIP

OUI Based Auto VoIP

Auto VoIP VLAN: Not Configured

Priority: 7 (0 to 7)

Display 10 rows Showing 1 to 10 of 118 entries Filter:

	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
☐	1/0/9	Disable	Down
☐	1/0/10	Disable	Down

First Previous 1 2 3 4 5 Next Last

Submit Refresh Edit Edit All Cancel

Figure 436: OUI Based Auto VoIP

Use the buttons to perform the following tasks:

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

Table 421: 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 interfaces 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.

7.1.2.4 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.

7 Configuring Quality of Service

To display the Protocol-Based Auto VoIP page, click **Quality of Service > Auto VoIP > Protocol Based Auto VoIP** in the navigation menu.

The screenshot shows the 'Protocol Based Auto VoIP' configuration page. The top section contains configuration fields: 'Auto VoIP VLAN' is 'Not Configured'; 'Prioritization Type' has 'Remark' selected and 'Traffic Class' unselected; '802.1p Priority' is set to 0 (range 0 to 7); 'Traffic Class' is set to 6 (range 0 to 6). Below the configuration fields is a table with 10 columns: 'Interface', 'Auto VoIP Mode', and 'Operational Status'. The table lists 10 interfaces (1/0/1 to 1/0/10) with 'Auto VoIP Mode' set to 'Disable' and 'Operational Status' set to 'Down'. At the bottom of the page are buttons for 'Submit', 'Refresh', 'Edit', 'Edit All', and 'Cancel'.

Figure 437: Protocol Based Auto VoIP

Use the buttons to perform the following tasks:

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

Table 422: 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: - a Remark – Remark the voice traffic with the specified 802.1p priority value at the ingress interface. - a Traffic Class – Assign VoIP traffic to the specified traffic class when egressing the interface.
802.1p Priority	The 802.1p priority used for protocol-based VoIP traffic. This field can be configured if the Prioritization Type is Remark . 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.
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

Field	Description
	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 interfaces being configured.
Auto VoIP Mode	The administrative mode of the Auto VoIP feature on the interface: - a Enable – The interface scans incoming traffic for the following call-control protocols: - a Session Initiation Protocol (SIP) - a H.323 - a Skinny Client Control Protocol (SCCP) - a Disable – The interface does not use the Auto VoIP feature to scan for call-control protocols.
Operational Status	The operational status of an interface. To be up, an interface must be administratively enabled and have a link.

- a 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.
- a Click **Refresh** to update the page with the most current data from the switch.

7.1.3 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.

7.1.3.1 CoS IP DSCP Mapping Configuration

Use the CoS IP DSCP Mapping Configuration page to map an IP DSCP value to an internal traffic class.

To display the CoS IP DSCP Mapping Configuration page, click **QoS > Class of Service > IP DSCP** in the navigation menu.

Figure 438: CoS IP DSCP Mapping Configuration

Table 423: CoS 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	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.

7.1.3.2 CoS Interface Configuration

Use the CoS Interface Configuration page to apply an interface shaping rate to all ports or to a specific port.

To display the CoS Interface Configuration page, click **QoS > Class of Service > Interface** in the navigation menu.

Figure 439: CoS Interface Configuration

Table 424: CoS 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.

Field	Description
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.

7.1.3.3 CoS Interface Queue Configuration

Use the CoS 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 CoS Interface Queue Configuration page, click **QoS > Class of Service > Queue** in the navigation menu.

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

Figure 440: CoS Interface Queue Configuration

Table 425: CoS Interface Queue Configuration Fields

Field	Description
Interface	Specifies the interface (physical, LAG, or Global) to configure.
Total Minimum Bandwidth Allocation (%)	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

Field	Description
	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 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. - a <i>Weighted</i>: Weighted round robin associates a weight to each queue. This is the default. - a <i>Strict</i>: Strict priority services traffic with the highest priority on a queue first
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.

- a If you make changes to the page, click **Submit** to apply the changes to the system.
- a Click **Restore Default** to reset the settings for the selected interface.
- a To reset the defaults for all interfaces, select Global from the **Slot/Port** menu before you click the button **Restore Default**.

7.1.3.4 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 on their drop precedence level.

- a To display the CoS Interface Queue Drop Precedence Configuration page, click **QoS > Class of Service > Drop Precedence** in the navigation menu.

The screenshot displays the 'CoS Interface Queue Drop Precedence Configuration' page. At the top, there are tabs for 'IP DSCP', 'Interface', 'Queue', 'Drop Precedence', and 'Statistics'. The 'Drop Precedence' tab is active. Below the tabs, the configuration is for Interface '1/0/1' and Queue ID '0'. The Threshold Unit is set to 'Percent' and ECN Capability is 'Disabled'. A table shows the WRED settings for four drop precedence levels:

Drop Precedence Level	WRED Minimum Threshold	WRED Maximum Threshold	WRED Drop Probability Scale
1	40	100	10
2	30	90	10
3	20	80	10
4	99	100	10

At the bottom of the page, there are three buttons: 'Refresh', 'Edit', and 'Restore Default'.

Figure 441: CoS Interface Queue Drop Precedence Configuration

Table 426: CoS Interface Queue Drop Precedence Configuration Fields

Field	Description
Interface	The interface where 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.

Field	Description
Threshold Unit	The unit for the WRED minimum and maximum threshold values. Threshold unit can be in percentage or kilobytes (KBs).
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.
Drop Precedence Level	The four drop precedence levels as follows: a Green: Low drop level 1 for classified TCP packets. a Yellow: Medium drop level 2 for classified TCP packets. a Red: High drop level 3 for classified TCP packets. a Non-TCP: Drop level 4 for non-TCP classified packets.
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.

- a Click **Refresh** to refresh the settings displayed on the page.
- a Click **Restore Default** 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.

7.1.3.4.1 Edit CoS Interface Queue Drop Precedence Configuration

- a 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:
 - a Threshold Unit in in percentage or KBs.
 - a Enable or disable ECN Capability.
 - a WRED minimum and maximum thresholds in percentage or KBs.

7 Configuring Quality of Service

a WRED drop probability scale.

Edit CoS Interface Queue Drop Precedence Configuration

Interface

1/0/1

Queue ID

0

Threshold Unit

☒ Percent ☐ Kbytes

ECN Capability

☒ Disable ☐ 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

Submit

Cancel

Figure 442: Edit CoS Interface Queue Drop Precedence Configuration

- a If you make changes to the page, click **Submit** to apply the changes to the system.
- a Click **Cancel** to cancel the changes.

7.1.3.5 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.

To display the CoS Statistics page, click **QoS > Class of Service > Statistics** in the navigation menu.

QoS > Class of Service > Statistics

Save Configuration

Hide Device View

Log Out

System

Switching

Routing

Security

QoS

Stacking

IP DSCP

Interface

Queue

Drop Precedence

Statistics

CoS Statistics

Display 10 rows

Showing 1 to 10 of 54 entries

Filter:

☐	Interface	Congestion Drop Packets	Tx Queue Kbytes	Rx Queue Kbytes	Green Drop Packets	Yellow Drop Packets	Red Drop Packets	Tx WRED Queue Kbytes	Tx ECN Packets
☐	1/0/1	0	-	-	-	0	0	-	-
☐	1/0/2	0	-	-	-	0	0	-	-
☐	1/0/3	0	-	-	-	0	0	-	-
☐	1/0/4	0	-	-	-	0	0	-	-
☐	1/0/5	0	-	-	-	0	0	-	-
☐	1/0/6	0	-	-	-	0	0	-	-
☐	1/0/7	0	-	-	-	0	0	-	-
☐	1/0/8	0	-	-	-	0	0	-	-
☐	1/0/9	0	-	-	-	0	0	-	-
☐	1/0/10	0	-	-	-	0	0	-	-

First

Previous

1

2

3

4

5

Next

Last

Refresh

Details


Clear

Figure 443: CoS Statistics

Table 427: 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 kilobytes transmitted on the interface.
Rx Queue Kbytes	The total number of kilobytes 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.
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 in the following figure displays.


**Figure 444: Clear CoS Queue Statistics**

Click **OK** to confirm the action to reset the counter values for the selected interfaces.

- 4** 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.



The screenshot shows a window titled "CoS Queue Statistics" with a close button (X) in the top right corner. Below the title bar, there is a field labeled "Interface" with the value "1/0/1". Below this is a table with the following data:

Queue ID	Total Drop Packets	Total Kbytes	Peak Kbytes	Current Kbytes	Average Kbytes
0	0	0	-	-	-
1	0	1172	-	-	-
2	0	0	-	-	-
3	0	0	-	-	-
4	0	2	-	-	-
5	0	0	-	-	-
6	0	0	-	-	-

At the bottom right of the window is a "Close" button.

Figure 445: CoS Queue Statistics Details

Table 428: 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 kilobytes enqueued to the associated queue.
Peak Kbytes	The total number of peak kilobytes 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 kilobytes enqueued to the associated queue.
Average Kbytes	The total number of average kilobytes enqueued to the associated queue.

7.1.4 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.

7.1.4.1 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.

MIB Table	Current Number / Maximum Number
Class Table	0 / 50
Class Rule Table	0 / 1300
Policy Table	0 / 64
Policy Instance Table	0 / 768
Policy Attribute Table	0 / 2304
Service Table	0 / 528

Figure 446: DiffServ Global Configuration and Status

Table 429: 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.

- a If you make changes to the page, click **Submit** to apply the changes to the system.
- a Click **Refresh** to update the page with the most current data from the switch.

7.1.4.2 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 447: DiffServ Class Summary

Use the buttons to perform the following tasks:

- a To add a DiffServ class, click **Add** and complete the fields in the **Add Class** window.
- a To change the name of an existing class, select the entry to modify and click **Rename**.
- a 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 430: 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: - a 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. - a Any–Any: of various match criteria defined for the class can be satisfied for a packet match.
Protocol	The Layer 3 protocol to use for filtering class types, which can be IPv4, IPv6 or None. None implies that no protocol is configured for the DiffServ class.
Match Criteria	The criteria used to match packets.

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

7.1.4.3 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:

- a To define criteria for matching packets within a class, click **Add Match Criteria**. When 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.
- a 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 page, click **QoS > DiffServ > Class Configuration** in the navigation menu.

Figure 448: DiffServ Class Configuration

Table 431: 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: - a 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. - a Any – Any of various match criteria defined for the class can be satisfied for a packet match.
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: - a The menu lists the name or number that identifies the ACL for all configured ACLs that are valid for the class type and protocol. - a Standard and Extended IPv4 ACLs use numbers in the range 1 to 199. All other ACL types use names. - a 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.

Field	Description
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.
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: - a EtherType Keyword – The menu includes several common protocols that are mapped to their EtherType values. - a EtherType Value – This field accepts custom EtherType values.
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: - a VLAN ID Start – The VLAN ID to match or the VLAN ID with the lowest value within a range of VLANs. - a 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.
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: - a Secondary VLAN ID Start – The secondary VLAN ID to match or the secondary VLAN ID with the lowest value within a range of VLANs. - a 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.
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: - a MAC Address – The source MAC address to match. - a 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.
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: a MAC Address – The destination MAC address to match.

Field	Description
	a 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.
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: - a Source Prefix – The source IPv6 prefix to match. - a Source Prefix Length – The IPv6 prefix length.
Destination IPv6 Address	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: - a Destination Prefix – The destination IPv6 prefix to match. - a Destination Prefix Length – The IPv6 prefix length.
Source L4 Port	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: - a 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. - a Port – The source port number to match.
Destination L4 Port	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: - a 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. - a Port – The destination port number to match.
IP DSCP	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: - a 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. - a IP DSCP Value – The IP DSCP value to match.
IP Precedence	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

Field	Description
	a packet is defined as the high-order three bits of the Service Type octet in the IP header.
IP TOS	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: - a IP TOS Bits – Enter a two-digit hexadecimal number to match the bits in a packet's ToS field. - a IP TOS Mask – Specify the bit positions that are used for comparison against the IP ToS field in a packet.
Protocol	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: - a 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. - a 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. - a Protocol Value – The IANA L4 protocol number value to match.

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

7.1.4.3.1 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 449: Add Match Criteria

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.

- 4 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.

7.1.4.4 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 450: DiffServ Policy Summary

Use the buttons to perform the following tasks:

- a To add a DiffServ policy, click **Add**.
- a To change the name of an existing policy, select the entry to modify and click **Rename**.
- a 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 432: 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: - a In – The policy is specific to inbound traffic. - a Out – The policy is specific to outbound traffic direction.
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.

7.1.4.5 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:

- a To add a class to the policy, click **Add Class**.
- a To add attributes to a policy or to change the policy attributes, select the policy with the attributes to configure and click **Add Attribute**.
- a 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 451: DiffServ Policy Configuration

Table 433: 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 steam 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: - a 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. - a IP DSCP Value – The IP DSCP value.
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

Field	Description
	Value field to select the IP Precedence value to mark in packets that match the policy-class.
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: - a Color Mode – The type of color policing used in DiffServ traffic conditioning. - a 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. - a Committed Rate (Kbps) – The maximum allowed arrival rate of incoming packets for this class. - a Committed Burst Size (KBs) – The amount of conforming traffic allowed in a burst. - a Conform Action – The action taken on packets that are considered conforming (below the police rate). - a Violate Action – The action taken on packets that are considered non-conforming (above the police rate).
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: - a Color Mode – The type of color policing used in DiffServ traffic conditioning. - a Committed Rate (Kbps) – The maximum allowed arrival rate of incoming packets for this class. - a Committed Burst Size (KBs) – The amount of conforming traffic allowed in a burst. - a Excess Burst Size (KBs) – The amount of conforming traffic allowed to accumulate beyond the Committed Burst Size (KBs) value during longer-than-normal idle times. This value allows for occasional bursting. - a Conform Action – The action taken on packets that are considered conforming (below the police rate). - a Exceed Action – The action taken on packets that are considered to exceed the committed burst size but are within the excessive burst size. - a Violate Action – The action taken on packets that are considered non-conforming (above the police rate).
Police Two Rate	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: a Color Mode – The type of color policing used in DiffServ traffic conditioning.

Field	Description
	- a Committed Rate (Kbps) – The maximum allowed arrival rate of incoming packets for this class. - a Committed Burst Size (KBs) – The amount of conforming traffic allowed in a burst. - a Peak Rate (Kbps) – The maximum peak information rate for the arrival of incoming packets for this class. - a Excess Burst Size (KBs) – The maximum size of the packet burst that can be accepted to maintain the Peak Rate (Kbps). - a Conform Action – The action taken on packets that are considered conforming (below the police rate). - a Exceed Action – The action taken on packets that are considered to exceed the committed burst size but are within the excessive burst size. - a Violate Action – The action taken on packets that are considered non-conforming (above the police rate).
Redirect Interface	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.

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

7.1.4.6 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 452: DiffServ Service Summary

Use the buttons to perform the following tasks:

- a To add a policy to an interface, click **Add**.
- a To edit a configured interface-policy association, select the entry to modify and click **Edit**.
- a 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 434: 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.

Field	Description
Direction	The traffic flow direction to which the policy is applied: a Inbound – The policy is applied to traffic as it enters the interface. a Outbound – The policy is applied to traffic as it exits the interface.
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.
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.

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.

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

8 Getting Started with Stacking

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

4 For complete syntax and usage information for the commands used in this chapter, refer to the *LCOS SX CLI Command Reference* for this release.

8.1 Understanding Switch Stacks

A *switch stack* is a set of up to 8 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.

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

- a System-level (global) features that apply to all stack members
- a 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* aka *Switch ID* or *Stack Unit 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:

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

8.1.1 Switch Stack Membership

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, LCOS SX configures the new member.

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

8.1.2 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:

- a The switch that is currently the stack manager
- a The switch with the highest stack member priority value (Admin Management Preference).

4 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.

- a The switch with the higher MAC address

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

- a The stack manager is removed from the switch stack
- a The stack manager is reset or powered off
- a The stack manager has failed
- a 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.

8.1.3 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 [Renumbering Stack Members](#) on page 524 and [Merging Two Operational Stacks](#) on page 526.

8.1.4 Stack Member Priority Values

You can set the stack member's priority under **Unit Configuration > Admin Management Preference** in the range 0 to 15.

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

8.2 Switch Stack Software Compatibility Recommendations

All stack members, including the stack manager, must run the same LCOS SX software version to ensure compatibility between stack members. 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 firmware on the stack manager is automatically rolled out to the stack member when the option **Stack Firmware Auto Upgrade** in the menu **Stacking > Base > Firmware Synchronizartion** is active (default setting). With the exception of the XS-6128QF (stacking only possible between multiple XS-6128QF) this also applies to different switch models (e.g. XS-5110F and GS-4530X) as the firmware is the same.

-
- 4 You can view status information regarding the stack with the CLI command `show switch`.

8.3 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.

8.4 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. When 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 (in Privileged EXEC):

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

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.

-
- 4 After downloading a configuration file to a stack, you must perform a configuration save operation from the LCOS SX 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.

8.5 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.

8.5.1 Connectivity to the Switch Stack through Console Ports

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

8.5.2 Connectivity to the Switch Stack through Telnet

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

8.6 General Practices

The following practices are recommended:

- a 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.
- a When physically removing or relocating a unit, always power down the unit before disconnecting stack cables.
- a 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.

8.7 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, make sure there is the same LCOS SX software version on each switch.

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 [Preconfiguration](#) on page 526.
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 [Renumbering Stack Members](#) on page 524 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.

8.8 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 > Base > Summary** page. To delete a member, select the unit number on the **Switch ID** menu and click the **Delete** button.

8.9 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. If the LCOS SX software version of the newly added member is not the same as the existing stack, update the software image.
2. Make sure that the redundant stack link is in place and functional. All stack members should be connected in a logical ring.
3. Preconfigure the new unit, if desired.
4. Install the new unit in the rack. (Assumes installation below the bottom-most unit, or above the top-most unit).
5. 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.

6. 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.
7. 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 (type `enable` to activate the EXEC mode first). The new unit should always join as a *member* (never as manager; the existing manager of the stack should not change).

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. [Merging Two Operational Stacks](#) on page 526 for more details.

Using the Web Interface, you create a new member for the stack through the **Stacking > Base > Summary** page. To create a new member, select the **Add** option from the **Switch ID** menu.

8.10 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 [Removing a Unit from the Stack](#) on page 523 to remove the desired stack member.
2. Follow the process in [Adding a Unit to an Operating Stack](#) on page 523 to add a new member to the stack with the following exceptions:
 - a Insert the new member in the same position in the stack as the one removed.
 - a The preconfiguration described in Step 2 of [Adding a Unit to an Operating Stack](#) on page 523 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 [Removing a Unit from the Stack](#) on page 523 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 [Adding a Unit to an Operating Stack](#) on page 523. 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.

8.11 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 [Adding a Unit to an Operating Stack](#) on page 523.

Using the Web Interface, you renumber a switch through the Stacking - Base - Summary 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.

8.12 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><tonunit-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.

8.13 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 [Moving a Manager to a Different Unit in the Stack](#) on page 525 procedure on this page.
2. Using the procedure [Removing a Unit from the Stack](#) on page 523, remove the unit from the stack.

8.14 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.

8.15 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.

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 [Stack Manager Election and Re-Election](#) on page 520. 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.

8.16 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.
 - a To display supported switches:
Use Privileged EXEC mode
Enter the command `show supported switchtype <x>` where *x* is the SID.
 - a To add a new member (see [Adding a Unit to an Operating Stack](#) on page 523):
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.

- 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 [Renumbering Stack Members](#) on page 524, or delete the preconfigured unit type using the `no member <unit-id>` command from the config stack mode.

8.17 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.

Stacking > Base > Summary

Save Configuration Hide Device View Log Out

System Switching Routing Security QoS Stacking

Summary Unit Configuration Firmware Update Firmware Synchronization

Stack Summary

Stack Trunk Hash Mode Source/Destination MAC

Display All rows Showing 1 to 1 of 1 entries Filter:

	Switch ID	Status	Management Status	Standby Switch	Preconfigured Model Identifier	Plugged-in Model Identifier
☐	1	OK	Management Switch		GS-4554XP	GS-4554XP

First Previous 1 Next Last

Submit Refresh Add Edit Remove

Figure 453: Stack Summary

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.

8.18 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 and loading of aggregate members when assigning a new flow.
- Account for existing flow assignment when changing loading across members.
- Identify instances where active flows can be moved to another aggregate member while avoiding re-ordering.

8 Getting Started with Stacking

DLB only works for known unicast traffic. It should not be used for multicast, broadcast, unknown unicast, and mirrored packets.

9 Configuration Examples

This appendix contains examples of how to configure selected features available in LCOS SX. Each example contains procedures on how to configure the feature by using the Web interface, and/or CLI, and/or SNMP.

- 4** Each configuration example starts from a factory-default configuration unless otherwise noted.

9.1 VLAN Configuration Examples

The following pages show how to configure VLAN, via the webinterface, the CLI and via SNMP.

9.1.1 Using the Web Interface to Configure VLANs

Scenario:

The default VLAN 1 should be operated along with the VLAN 20. The ports 1/0/1 to 1/0/3 are assigned the following tagging modes:

- a 1/0/1: Access with the port VLAN ID 20
- a 1/0/2: Hybrid with the port VLAN ID 1
- a 1/0/3: Trunk




Figure 454: VLAN Configuration Scenario

Procedure:

Changing the management VLAN: A switch is administered via the management VLAN. The management VLAN is set in the menu **System > Connectivity > IPv4**.

- 4 An XS or GS-45xx series switch can be administered from the management VLAN and also from any other network if the switch has an IP address in this network. It is not possible to set up routing for the management VLAN.

The screenshot shows the 'IPv4 Network Connectivity' configuration page. The 'Management VLAN ID' field is highlighted with a red box and contains the value '1'. The page includes tabs for IPv4, IPv6, IPv6 Neighbors, and DHCP Client Options. The configuration fields include Network Configuration Protocol (None, Bootp, DHCP), DHCP Client Identifier, IP Address (192.168.10.10), Subnet Mask (255.255.255.0), Default Gateway (192.168.10.1), MAC Address Type (Burned In, Locally Administered), Burned In MAC Address (00:A0:57:53:6F:59), Locally Administered MAC Address (00:00:00:00:00:00), and Management VLAN ID (1).

Figure 455: IPv4 Network Connectivity

Setting up the VLAN function:

The tagging modes for the VLAN are configured in two parts in the menus **Switching > VLAN > Port Configuration** and **Switching > VLAN > Port Summary**. You have to complete the configuration in both menus for the VLAN to work correctly. Please also refer to the Knowledge Base article [VLAN tagging modes explained](#).

- 4 VLAN tagging modes can also be set in the menu **Switching > VLAN > Switchport Summary** using the option **Switchport Mode**. Be sure to leave this at the default setting **General**, otherwise conflicts will occur and communication will fail!

1. Connect to the switch via the web interface and navigate to the menu **Switching > VLAN > Overview**. The default VLAN 1 is displayed:

The screenshot shows the 'VLAN Overview' page. It displays a table with one entry: VLAN ID 1, Name default, Type Default, and RSPAN. The 'Add' button is highlighted with a red box. The page includes tabs for Overview, Port Configuration, Port Summary, Switchport Summary, Internal Usage, Reset, and RSPAN.

Figure 456: VLAN Overview

2. Click **Add** to create a new VLAN.
3. Enter the desired VLAN ID (in this example VLAN 20) in the Individual/Range field and click on **Submit**.

The screenshot shows the 'Add VLAN' dialog box. The 'VLAN ID or Range' field is highlighted with a red box and contains the value '20'. The 'Submit' button is also highlighted with a red box. The dialog box includes a text area with instructions: '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.'

Figure 457: Add VLAN-ID

Access tagging mode:

1. Change to the Port Configuration tab and make sure that VLAN ID 1 is selected. Then select the interface 1/0/1 and click **Edit**.

- 4 You can edit all of the interfaces at the same time by clicking the button **Edit All**. It is also possible to check multiple interfaces and click on **Edit** to edit all of them at the same time. The same settings are saved for all of them.




Figure 458: Edit VLAN Port Configuration

2. Set the Participation parameter to **Exclude** so that VLAN 1 cannot communicate through interface 1/0/1. Then click on **Submit**.




Figure 459: Edit VLAN Port Configuration

3. Select the VLAN ID 20 created in [Add VLAN-ID](#). Then select the interface 1/0/1 and click **Edit**.




Figure 460: Edit VLAN Port Configuration

4. Modify the following parameters and click **Submit**:
- Participation: Select the option **Include** so that VLAN 20 may communicate via the interface.

9 Configuration Examples

- a Tagging: Select the option **Untagged**.

Edit VLAN Port Configuration

NOTE: Tagging will only be enabled for VLAN member ports.

VLAN ID	20
Interface	1/0/1
Participation	☒ Include ☐ Auto Detect ☐ Exclude
Tagging	☒ Untagged ☐ Tagged

Submit **Cancel**

Figure 461: Edit VLAN Port Configuration

5. Change to the tab **Port Summary**, mark the interface 1/0/1 and click on **Edit**.

VLAN Port Summary

Display 10 rows Showing 1 to 10 of 74 entries

Interface	Port VLAN ID	Port VLAN ID Current	Acceptable Frame Type	Ingress Filtering	Untagged VLANs	Tagged VLANs	Forbidden VLANs
☒ 1/0/1	1	1	Admit All	Disabled	20		1
☐ 1/0/2	1	1	Admit All	Disabled	1		
☐ 1/0/3	1	1	Admit All	Disabled	1		
☐ 1/0/4	1	1	Admit All	Disabled	1		
☐ 1/0/5	1	1	Admit All	Disabled	1		
☐ 1/0/6	1	1	Admit All	Disabled	1		
☐ 1/0/7	1	1	Admit All	Disabled	1		
☐ 1/0/8	1	1	Admit All	Disabled	1		
☐ 1/0/9	1	1	Admit All	Disabled	1		
☐ 1/0/10	1	1	Admit All	Disabled	1		

First Previous 1 2 3 4 5 Next Last

Edit **Edit All**

Figure 462: VLAN Port Summary

6. Modify the following parameters and click **Submit**:
 - a Port VLAN ID: Enter the VLAN ID that is appended to an incoming packet that has no VLAN tag (in this example the PVID 20).
 - a Acceptable Frame Type: Select the option **Only Untagged**.

Edit Port Configuration

Interface	1/0/1
Port VLAN ID	20 (1 to 4093)
Acceptable Frame Type	☐ Admit All ☐ Only Tagged ☒ Only Untagged
Ingress Filtering	☒ Disable ☐ Enable
Priority	0 (0 to 7)

Submit **Cancel**

Figure 463: Edit VLAN Port Configuration

Hybrid tagging mode:

1. Change to the **Port Configuration** tab and make sure that VLAN ID 1 is selected. Then select the interface 1/0/2 and click **Edit**.

- 4 You can edit all of the interfaces at the same time by clicking the button **Edit All**. It is also possible to check multiple interfaces and click on **Edit** to edit all of them at the same time. The same settings are saved for all of them.




Figure 464: Edit VLAN Port Configuration

2. Modify the following parameters and click **Submit**:
 - a Participation: Select the option **Include** so that VLAN 1 may communicate via the interface.
 - a Tagging: Select the option **Untagged**.




Figure 465: Edit VLAN Port Configuration

3. Select the VLAN ID 20 created in [Add VLAN-ID](#). Then select the interface 1/0/2 and click **Edit**.





Figure 466: Edit VLAN Port Configuration

- 4 You can edit all of the interfaces at the same time by clicking the button **Edit All**. It is also possible to check multiple interfaces and click on **Edit** to edit all of them at the same time. The same settings are saved for all of them.
4. Modify the following parameters and click **Submit**:
 - a Participation: Select the option **Include** so that VLAN 20 may communicate via the interface.


- a Tagging: Select the option **Tagged**.



The dialog box titled "Edit VLAN Port Configuration" contains a note: "NOTE: Tagging will only be enabled for VLAN member ports." Below the note are four fields: "VLAN ID" with value 20, "Interface" with value 1/0/2, "Participation" with radio buttons for "Include" (selected), "Auto Detect", and "Exclude", and "Tagging" with radio buttons for "Untagged" and "Tagged" (selected). At the bottom right are "Submit" and "Cancel" buttons.

Figure 467: Edit VLAN Port Configuration

5. Change to the tab **Port Summary**, mark the interface 1/0/2 and click on **Edit**.




The "VLAN Port Summary" table shows a list of interfaces and their VLAN configurations. The interface 1/0/2 is selected. The table has columns for Interface, Port VLAN ID, Port VLAN ID Current, Acceptable Frame Type, Ingress Filtering, Untagged VLANs, Tagged VLANs, and Forbidden VLANs. The interface 1/0/2 has Port VLAN ID 1, Acceptable Frame Type Admit All, Ingress Filtering Disabled, Untagged VLANs 20, and Tagged VLANs 20.

Interface	Port VLAN ID	Port VLAN ID Current	Acceptable Frame Type	Ingress Filtering	Untagged VLANs	Tagged VLANs	Forbidden VLANs
1/0/1	20	20	Only Untagged	Disabled	20		1
1/0/2	1	1	Admit All	Disabled	1	20	
1/0/3	1	1	Admit All	Disabled	1		
1/0/4	1	1	Admit All	Disabled	1		
1/0/5	1	1	Admit All	Disabled	1		
1/0/6	1	1	Admit All	Disabled	1		
1/0/7	1	1	Admit All	Disabled	1		
1/0/8	1	1	Admit All	Disabled	1		
1/0/9	1	1	Admit All	Disabled	1		
1/0/10	1	1	Admit All	Disabled	1		

Figure 468: VLAN Port Summary

4 You can edit all of the interfaces at the same time by clicking the button **Edit All**. It is also possible to check multiple interfaces and click on **Edit** to edit all of them at the same time. The same settings are saved for all of them.

6. Modify the following parameters and click **Submit**:
- a Port VLAN ID: Enter the VLAN ID that is appended to an incoming packet that has no VLAN tag (in this example the PVID 1).
- a Acceptable Frame Type: Select the option **Admit All** so untagged and tagged packets can be transmitted.



The dialog box titled "Edit Port Configuration" contains five fields: "Interface" with value 1/0/2, "Port VLAN ID" with value 1 (range 1 to 4093), "Acceptable Frame Type" with radio buttons for "Admit All" (selected), "Only Tagged", and "Only Untagged", "Ingress Filtering" with radio buttons for "Disable" (selected) and "Enable", and "Priority" with value 0 (range 0 to 7). At the bottom right are "Submit" and "Cancel" buttons.

Figure 469: Edit VLAN Port Configuration

Trunk tagging mode:

1. Change to the Port Configuration tab and make sure that VLAN ID 1 is selected. Then select the interface 1/0/3 and click **Edit**.




Figure 470: Edit VLAN Port Configuration

2. Modify the following parameters and click **Submit**:
 - a Participation: Select the option **Include** so that VLAN 1 may communicate via the interface.
 - a Tagging: Select the option **Tagged**.




Figure 471: Edit VLAN Port Configuration

3. Select the VLAN ID 20 created in [Add VLAN-ID](#). Then select the interface 1/0/3 and click **Edit**.





Figure 472: Edit VLAN Port Configuration

4. You can edit all of the interfaces at the same time by clicking the button **Edit All**. It is also possible to check multiple interfaces and click on **Edit** to edit all of them at the same time. The same settings are saved for all of them.
4. Modify the following parameters and click Submit:
 - a Participation: Select the option **Include** so that VLAN 20 may communicate via the interface.

a Tagging: Select the option **Tagged**.




NOTE: Tagging will only be enabled for VLAN member ports.

VLAN ID	20
Interface	1/0/3
Participation	☒ Include ☐ Auto Detect ☐ Exclude
Tagging	☐ Untagged ☒ Tagged

Submit Cancel

Figure 473: Edit VLAN Port Configuration

5. Change to the tab Port Summary, mark the interface 1/0/3 and click on **Edit**.



Overview Port Configuration **Port Summary** Switchport Summary Internal Usage Reset RSPAN

VLAN Port Summary

Display 10 rows Showing 1 to 10 of 74 entries

	Interface	Port VLAN ID	Port VLAN ID Current	Acceptable Frame Type	Ingress Filtering	Untagged VLANs	Tagged VLANs	Forbidden VLANs
☐	1/0/1	20	20	Only Untagged	Disabled	20		
☐	1/0/2	1	1	Admit All	Disabled	1	20	1
☒	1/0/3	1	1	Admit All	Disabled	1	1, 20	
☐	1/0/4	1	1	Admit All	Disabled	1		
☐	1/0/5	1	1	Admit All	Disabled	1		
☐	1/0/6	1	1	Admit All	Disabled	1		
☐	1/0/7	1	1	Admit All	Disabled	1		
☐	1/0/8	1	1	Admit All	Disabled	1		
☐	1/0/9	1	1	Admit All	Disabled	1		
☐	1/0/10	1	1	Admit All	Disabled	1		


First Previous 1 2 3 4 5 Next Last

Refresh Edit Edit All

Figure 474: VLAN Port Summary

4 You can edit all of the interfaces at the same time by clicking the button **Edit All**. It is also possible to check multiple interfaces and click on **Edit** to edit all of them at the same time. The same settings are saved for all of them.

6. For the Acceptable Frame Type select the option **Only Tagged** and click **Submit**.



Interface	1/0/3
Port VLAN ID	1 (1 to 4093)
Acceptable Frame Type	☐ Admit All ☒ Only Tagged ☐ Only Untagged
Ingress Filtering	☒ Disable ☐ Enable
Priority	0 (0 to 7)

Submit Cancel

Figure 475: Edit VLAN Port Configuration

Save the switch configuration as the boot configuration:

With the configuration complete, click on **Save Configuration** in the top right-hand corner to save the configuration as the **boot configuration**.

- 4 The start configuration is retained even if the device is restarted or there is a power failure.



Figure 476: Save as Boot Configuration

9.1.2 Using the CLI to Configure VLANs

Scenario:

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.




Figure 477: VLAN Example Network Diagram

Procedure:

1. Create VLAN 2 and VLAN 3.

```
(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.

```
(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.

```
(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.

```
(Routing) (Config) #vlan port tagging all 2
exit
```

5. Assign the ports that will belong to VLAN 3.

4 Port 1/0/2 belongs to both VLANs, and port 1/0/1 can never belong to VLAN 3.

```
(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.

```
(Routing) #Config
interface 1/0/4
  vlan acceptframe all
  exit
exit
```

9.1.3 Using SNMP to Configure VLANs

Scenario:

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.





Figure 478: VLAN Example Network Diagram

Procedure:

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).



2. 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:

- a 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.
- a 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.

3. 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.
4. 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.
5. 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.
6. 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.

- 4** 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 > Overview** 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**.

Spanning Tree Switch Configuration	
Spanning Tree Admin Mode	☐ Disable ☒ Enable
Force Protocol Version	IEEE 802.1s
Configuration Name	00-A0-57-53-6F-C5 (1 to 32 characters)
Configuration Revision Level	0 (0 to 65535)
Configuration Digest Key	0xAC36177F50283CD4883821D8A826DE62
Configuration Format Selector	0

Figure 479: Spanning Tree Switch Configuration

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.

Figure 480: Add MST Entry

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.

```
(Routing) #vlan database
vlan 10
vlan 20
exit
```

2. Enable spanning tree Globally

```
(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
```

9 Configuration Examples

- Change the priority of MST ID 20 to ensure the other bridge is the root bridge.

```
spanning-tree mst priority 20 61440
```

- Enable STP on interface 1/0/1

```
interface 1/0/1
 spanning-tree port mode
exit
```

- Enable STP on interface 1/0/2

```
interface 1/0/2
 spanning-tree port mode
```

- 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

- Use the objects in dot1qVlanStaticTable (in dot1qVlan in the QBRIDGE-MIB module) to create VLANs 10 and 20.
- To enable spanning tree globally, set the agentStpAdminMode object in the LCOS SX-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).LCOSSX(1).LCOSSXSwitching(1).agentConfigGroup(2).agentStpSwitchConfigGroup(15).agentStpAdminMode(6).

- 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.
- Use the agentStpMstRowStatus object in the agentStpMstTable to create MST instances 10 and 20.



- 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.
- 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 1.3.6.1.4.1.4413.1.1.2.15.14.1.1.10.10 (the final .10 is the VLAN ID)

a 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 LCOS SX 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 LCOS SX software to provide the VLAN routing support shown in the diagram.




Figure 481: VLAN Routing Example Network Diagram

9.3.1 Using the CLI to Configure VLAN Routing

1. Create VLAN 10 and VLAN 20.

```
(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
```

9 Configuration Examples

```
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:

```
(Routing) #vlan database
vlan routing 10
vlan routing 20
exit
```

6. View the logical interface IDs assigned to the VLAN routing interfaces.

```
(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


1. Use the `dot1qVlanStaticRowStatus` object in the `dot1qVlanStaticTable` to create VLAN 10 and VLAN 20.
2. To configure VLAN membership, retrieve the current `dot1qStaticEgressPorts` mask and append the desired interfaces to the mask.
 - a VLAN 10: 1/0/1 and 1/0/2
 - a VLAN 20: 1/0/3
3. To assign the PVID for an interface, use the `dot1qPvid` object.
 - a 1/0/1: PVID 10
 - a 1/0/2: PVID 10

a 1/0/3: PVID 20

4. 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.

5. To enable routing for the VLANs, use the `agentSwitchIpVlanRoutingStatus` object in the `agentSwitchIpVlanTable` under `agentSwitchIpGroup` in LCOS SX Routing to set the value for VLAN 10 and VLAN 20 to `CreateAndGo` (4).



6. Walk the `agentSwitchIpVlanIfIndex` 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.

4 While setting the ip address for the VLAN interface, the `agentSwitchIpInterfaceIpAddress` and `agentSwitchIpInterfaceNetMask` should be set together.

a VLAN index 482 (VLAN 10): 192.150.3.1 255.255.255.0

a 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.

PBR 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.

- a Physical Interface 1/0/2 – Member of VLAN 10
- a Physical Interface 1/0/4 – Member of VLAN 20
- a Physical Interface 1/0/22 – Member of VLAN 30
- a Physical Interface 1/0/24 – Member of VLAN 40




Figure 482: 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 482: Policy-Based Routing Example](#) on page 546. Traffic sent to VLAN Interface 10 is destined for VLAN Interface 20. 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.

```
(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

```

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.

```

config
    interface vlan 10
    ip policy pbr_test
    exit
exit

```

After this step, traffic mentioned in [5](#) on page 547 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.

```

(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.

9.5.1 Using the CLI to Configure OSPF

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 LCOS SX.




Figure 483: OSPF Example Network Diagram: Border Router

1. Enable routing on the switch.

```
(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

```

5. **4** 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:

```

(Routing) #config
router ospf
network 192.164.4.0 0.0.0.255 area 2

```

9.5.2 Configuring Stub and NSSA Areas

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.

- 4** 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 484: OSPF Configuration – Stub Area and NSSA Area on page 550 illustrates this example OSPF configuration.




Figure 484: OSPF Configuration – Stub Area and NSSA Area

9.5.3 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:

```
(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:

```

(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

```

9.5.4 Using the CLI to Configure OSPFv3 Enhancements

To configure OSPFv3 enhancements using the CLI, perform the following steps:

9 Configuration Examples

1. Enable the IPv6 router admin mode on the switch.

```
(Routing) #config
ipv6 unicast-routing
exit
```

2. On port 1/0/1:

- a Enable routing
- a Assign an IPv6 address
- a Enable OSPFv3 and specify its area
- a Set the OSPFv3 interface type
- a 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 485: 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.

```
(Routing) #config
radius server host auth 10.10.10.10
```

2. 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 LCOS SX-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. User 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 LCOS SX-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 LCOS SX-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, perform the following steps:

1. Access the **Security > Authentication Manager > Configuration** page.

The **Authentication Manager Configuration** page is displayed (see [Figure 486: Authentication Manager Configuration](#) on page 554).




Figure 486: Authentication Manager Configuration

2. Select **Enable**.
3. Click **Submit**.
4. Access the **Security > Authentication Manager > Authentication Tiering** page.

The **Authentication Tiering** page is displayed (see [Figure 487: Authentication Tiering](#) on page 555).

Interface	Configured Order	Enabled Order	Configured Priority	Enabled Priority	Re-Authentication Timer
1/0/1	Dot1x, MAB		Dot1x, MAB		30
1/0/2	Dot1x, MAB		Dot1x, MAB		30
1/0/3	Dot1x, MAB		Dot1x, MAB		30
1/0/4	Dot1x, MAB		Dot1x, MAB		30
1/0/5	Dot1x, MAB		Dot1x, MAB		30
1/0/6	Dot1x, MAB		Dot1x, MAB		30
1/0/7	Dot1x, MAB		Dot1x, MAB		30
1/0/8	Dot1x, MAB		Dot1x, MAB		30
1/0/9	Dot1x, MAB		Dot1x, MAB		30
1/0/10	Dot1x, MAB		Dot1x, MAB		30

Figure 487: Authentication Tiering

5. Select interface 1/0/3 check box and click **Edit**.

The **Edit Authentication Tiering** page is displayed (see [Figure 488: Edit Authentication Tiering](#) on page 555).

Interface	1/0/1
Re-Authentication Timer (Seconds)	30 (10 to 65535)
Configured Method Order	Methods Captive Portal Order Dot1x MAB
Configured Method Priority	Methods Captive Portal Priority Dot1x MAB

Figure 488: Edit Authentication Tiering

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 489: 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.

```
(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 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
```

5. 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 `LCOS SXQOSDiffServPrivate` in the `LCOS SX-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 `LCOS SX-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 `agentDiffServPolicyAttrStmtAssignQueueId` 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 `agentDiffServPolicyAttrStmtMarkIpDscpVal` 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 LCOS SX router:

1. Configure IP routing, IP multicast, IGMP, and PIM-SM in the global configuration mode with the following commands:

```
(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 preceding configuration example enabled PIM-SM on the router.

9.9.2 Using SNMP to Configure PIM-SMv4

1. Use the following objects to configure an OSPF router and globally enable IP routing, multicast, IGMP, and PIM-SM.
 - a Enable OSPF: `ospfAdminStat` under `ospfGeneralGroup` in the OSPF-MIB module
 - a Set OSPF router ID: `ospfRouterId` under `ospfGeneralGroup` in the OSPF-MIB module
 - a Enable routing: `agentSwitchIpRoutingMode` object in `agentSwitchIpGroup` under FASTPATH Routing
 - a Enable multicast: `agentMulticastRoutingAdminMode` under `agentMulticastRoutingConfigGroup` in the FASTPATH-MULTICAST-MIB module
 - a Enable IGMP: `agentMulticastIGMPAdminMode` under `agentMulticastIGMPConfigGroup`
 - a Enable PIM-SM: `agentMulticastPIMSMAdminMode` under `agentMulticastPIMSMConfigGroup`
2. Use the `pimSmStaticRPIpAddress` object in the `agentMulticastPIMSMStaticRPTable` 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.
3. Use the following objects to enable routing, IGMP, PIM-SM, and OSPF on one or more interfaces:
 - a Enable routing on the interface: `agentSwitchIpInterfaceRoutingMode` in the `agentSwitchIpInterfaceTable` under the FASTPATH-ROUTING-MIB module.
 - a Enable IGMP on the interface: `mgmdRouterInterfaceStatus` in the `mgmdRouterInterfaceTable` under the MGMD-STD-MIB module.
 - a Enable PIM-SM on an interface: `pimSmInterfaceStatus` in the `pimSmInterfaceTable` under the PIM-STD-MIB module.
 - a Enable OSPF on an interface: `ospfIfStatus` in the `ospfIfTable` in the OSPF-MIB module.
 - a 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 LCOS SX router:

9 Configuration Examples

1. Configure IP routing, IP multicast, IGMP, and PIM-SM in the global configuration mode with the following commands:

```
(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.

9.9.3.2 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 490: Multicast Data Forwarding Example

Configure PIM-SM on all the routers in the above topology in a similar way as was explained in [Configuring PIM-SM](#) on page 559. 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 follows:

```
ip pim rp-address 192.168.102.4 24.0.0.0 255.0.0.0
```

In the topology shown in [Figure 490: Multicast Data Forwarding Example](#) on page 561, 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:

```
(LCOS SX-RTR5)#show ip mcast mroute summary
Multicast Route Table Summary
```

Source IP	Group IP	Protocol	Incoming Interface	Outgoing 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. LCOS SX 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:

```
(LCOS SX-RTR4)#show ip mcast mroute summary
Multicast Route Table Summary
```

Source IP	Group IP	Protocol	Incoming Interface	Outgoing 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:

```
(LCOS SX-RTR5)#show ip mcast mroute summary
Multicast Route Table Summary
```

Source IP	Group IP	Protocol	Incoming Interface	Outgoing 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, LCOS SX 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:

```
(LCOS SX-RTR5)#show ip mcast mroute summary
Multicast Route Table Summary
```

Source IP	Group IP	Protocol	Incoming Interface	Outgoing 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, LCOS SX 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:

```
(LCOS SX-RTR5)#show ip mcast mroute summary
Multicast Route Table Summary
```

Source IP	Group IP	Protocol	Incoming Interface	Outgoing 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:

```
(LCOS SX-RTR5)#show ip mcast mroute summary
Multicast Route Table Summary
```

Source IP	Group IP	Protocol	Incoming Interface	Outgoing Interface List
*	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.

```
(LCOS SX-RTR5)#show ip mcast mroute summary
Multicast Route Table Summary
```

Source IP	Group IP	Protocol	Incoming Interface	Outgoing Interface List
*	224.1.2.1	PIMSM	vlan104	
192.168.107.10	224.1.2.1	PIMSM	vlan104	

9.10 IGMP and MLD Snooping Switches

9.10.1 Snooping Functionality on a LCOS SX 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.

9.10.1.1 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.

9.10.1.2 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.

9.10.1.3 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.

9.10.1.4 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.

9.10.1.5 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.

9.10.1.6 Version Compatibility

The following table shows the IGMP/MLD versions that the LCOS SX snooping switch supports.

Table 435: IGMP/MLD Version Support

Protocol Version	Support
IGMPv1	Yes
IGMPv2	Yes
IGMPv3	Yes
MLDv1	Yes
MLDv2	Yes

9.10.2 Snooping Switch Restrictions

This section describes the IGMP and MLD Snooping implementation on a LCOS SX-based snooping switch.

9.10.2.1 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. LCOS SX snooping switches support the following record types:

- a MODE_IS_INCLUDE
- a MODE_IS_EXCLUDE
- a CHANGE_TO_INCLUDE_MODE
- a CHANGE_TO_EXCLUDE_MODE
- a ALLOW_NEW_SOURCES
- a 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.

9.10.2.2 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.

9.10.2.3 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 LCOS SX 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.

9.10.3 Configuring IGMP and MLD Snooping

9.10.3.1 Configuration Commands

The LCOS SX 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 *LCOS SX CLI Command Reference*.

9.10.3.2 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
```

9.10.3.3 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
```

9.10.3.4 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	Source 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

9.10.3.5 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 Interface
1	01:00:5E:11:22:33	Filter	Static	Mgmt Config	Fwd: 1/0/2	Fwd: 1/0/2

```
console#show mac-address-table static all
```

MAC Address	VLAN ID	Source Port(s)	Destination Port(s)
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 Interface
1	01:00:5E:11:22:33	Filter	Static	Mgmt Config	Fwd: 1/0/2	Fwd: 1/0/2

9.11 Multicast Snooping Example (with IP Multicast Routing)

The examples in this section use the network topology shown in [Figure 491: Network Topology for Multicast Snooping with IP Multicast Routing](#) on page 568.




Figure 491: Network Topology for Multicast Snooping with IP Multicast Routing

The above network topology includes the following elements:

- a Snooping Switches: D1, D2, D3 with snooping on VLAN 10, 20
- a Multicast Router: D4 with PIM-SM and snooping on VLAN10, 20
- a Multicast Listeners: Client A–G
- a Multicast Sources: Server A— 239.20.30.40, Server B—239.20.30.42
- a Subnets: VLAN 10— 192.168.10.70/24, VLAN 20— 192.168.20.70/24
- a Router attached ports: D3—1/0/20, D2—PortChannel1, D1-1/0/15

9.11.1 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.

9.11.1.1 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.

9.11.1.2 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.

9.11.2 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.

9.11.2.1 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.

9.11.2.2 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.

9.11.2.3 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.

9.11.2.4 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.

9.12 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 LCOS SX 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.

9.13 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.

The following figure provides a visual overview of the RSPAN configuration example.




Figure 492: RSPAN Configuration Example


9.13.1 Configuring RSPAN Using the Web Interface

The following sections describe configuring RSPAN using the web interface.

9.13.1.1 Configuration on the Source Switch (SW1)

1. Create vlan 100.

- a. Access the **Switching > VLAN > Overview** page. The **VLAN Overview** page is displayed.




The screenshot shows the 'VLAN Overview' page with a navigation bar at the top containing 'Overview', 'Port Configuration', 'Port Summary', 'Switchport Summary', 'Internal Usage', 'Statistics', 'Reset', and 'RSPAN'. Below the navigation bar is a table with the following columns: 'VLAN ID', 'Name', 'Type', 'RSPAN', and 'Statistics Mode'. The table contains one entry with 'VLAN ID' 1, 'Name' 'default', 'Type' 'Default', 'RSPAN' (empty), and 'Statistics Mode' 'Enabled'. Below the table are navigation buttons: 'First', 'Previous', '1', 'Next', and 'Last'. At the bottom are buttons: 'Refresh', 'Add', 'Edit', and 'Remove'.

VLAN ID	Name	Type	RSPAN	Statistics Mode
1	default	Default		Enabled

Figure 493: VLAN Overview

- b. Click **Add**.

The **Add VLAN** page is displayed.



The screenshot shows the 'Add VLAN' dialog box. It has a title bar 'Add VLAN' with a close button. Inside, there is a text input field labeled 'VLAN ID or Range'. Below the field is a hint: '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.' At the bottom right are 'Submit' and 'Cancel' buttons.


Figure 494: 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 is displayed.



The screenshot shows the 'RSPAN Configuration' page. It has a navigation bar at the top with 'Overview', 'Port Configuration', 'Port Summary', 'Switchport Summary', 'Internal Usage', 'Statistics', 'Reset', and 'RSPAN'. Below the navigation bar is a section titled 'RSPAN Configuration'. It contains two lists: 'VLAN IDs' and 'RSPAN VLAN IDs'. The 'VLAN IDs' list contains '1'. The 'RSPAN VLAN IDs' list contains '100'. Between the lists are arrows for adding and removing items. At the bottom are 'Submit', 'Refresh', and 'Cancel' buttons.

Figure 495: RSPAN Configuration

- 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 is displayed.

Figure 496: Multiple Port Mirroring

- b. In the **Destination** field, click the **Edit** icon, as shown in [Figure 496: Multiple Port Mirroring](#) on page 574.
- 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**, as shown in [Figure 497: Multiple Port Mirroring](#) on page 574. The **Configure Source** page is displayed.

Figure 497: Multiple Port Mirroring

- 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**, as shown in [Figure 498: Multiple Port Mirroring](#) on page 575. The **Configure Session** page is displayed.

The screenshot shows the 'Multiple Port Mirroring' configuration page. At the top, there are tabs: Summary, Description, Cable Test, Mirroring (selected), Mirroring Summary, and SFP Information. Below the tabs, the page title 'Multiple Port Mirroring' is displayed. The main configuration area includes fields for Session ID (set to 1), Mode (Disabled), Destination (None), IP ACL (None), and MAC ACL (None). Below these fields, there is a table with columns 'Source' and 'Direction'. The table is currently empty, with a message 'Table is Empty' and navigation buttons (First, Previous, Next, Last). At the bottom of the page, there are four buttons: Refresh, Configure Session (circled in red), Configure Source, and Remove Source.

Figure 498: Multiple Port Mirroring

- b. In **Mode**, select **Enable**.
- c. Click **Submit**.

9.13.1.2 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 is displayed (see [Figure 493: VLAN Overview](#) on page 573).
 - b. Click **Add**.
The **Add VLAN** page is displayed (see [Figure 494: Add VLAN](#) on page 573).
 - 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 is displayed (see [Figure 495: RSPAN Configuration](#) on page 573).
 - 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 is displayed.




Figure 499: VLAN Port Configuration

- 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 is displayed.




Figure 500: Edit VLAN Port Configuration

- e. In the Participation field, select **Include**.
- f. In Tagging, select **Tagged**.

9.13.1.3 Configuration on the Destination Switch (SW3)

To configure the destination switch:

- 1. Create **vlan 100**.
 - a. Access the **Switching > VLAN > Status** page.
The **VLAN Overview** page is displayed (see [Figure 493: VLAN Overview](#) on page 573).
 - b. Click **Add**.
The **Add VLAN** page is displayed (see [Figure 494: Add VLAN](#) on page 573).
 - 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** [Figure 495: RSPAN Configuration](#) on page 573 page is displayed.
 - 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** [Figure 496: Multiple Port Mirroring](#) on page 574 page is displayed.
 - 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 is displayed (see [Figure 497: Multiple Port Mirroring](#) on page 574).
 - b. From **Type**, select **Remote VLAN**.
 - c. Click **Submit**.
5. Enable the port mirroring session.
 - a. Click **Configure Session**.
The **Configure Session** [Figure 498: Multiple Port Mirroring](#) on page 575 page is displayed.
 - b. In **Mode**, select **Enable**.
 - c. Click **Submit**.

9.13.2 Configuring RSPAN Using the CLI

The following sections describe configuring RSPAN using the CLI interface.

9.13.2.1 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
```

9.13.2.2 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
```

9.13.2.3 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
```

9.14 Configuring Virtual Port Channel

See [Figure 501: VPC Configuration Diagram](#) on page 579 for a visual overview of the Virtual Port Channel (VPC) configuration steps.




Figure 501: VPC Configuration Diagram

4 This feature is only supported by the LANCOM XS-6128QF.

9.14.1 Configuring Virtual Port Channel Using the Web Interface

The following sections describe how to configure VPC using the web interface.

9.14.1.1 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 > Overview** in the navigation menu.





Figure 502: VLAN-Overview

- Click **Add** to create the VLANs.



Add VLAN


VLAN ID or Range

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.

Submit Cancel

Figure 503: Add VLAN

- In the VLAN ID or Range field, enter **10-17** and click **Submit**.
- Click **Add** to create the VLAN routing interface that will be used for the Dual Control Plane Detection Protocol.
- In the VLAN ID or Range field, enter **100** and click **Submit**.
- To modify Port Channels 1, 2, and 3, click **Switching > Port Channel / LAG > Summary** in the navigation menu.



Port Channel Summary

Display 10 rows Showing 1 to 10 of 64 entries Filter:

Name	Type	Admin Mode	STP Mode	Link State	Link Trap	Local Preference Mode	Members	Active Ports	Load Balance
☐ ch1	Static	Enable	Enable	Down	Disable	Disable	1/0/1, 1/0/2		Source/Destination MAC, VLAN, Ethertype, Incoming Port
☐ ch2	Static	Enable	Enable	Down	Disable	Disable	1/0/3, 1/0/4, 1/0/5, 1/0/6		Source/Destination MAC, VLAN, Ethertype, Incoming Port
☐ ch3	Static	Enable	Enable	Down	Disable	Disable	1/0/7, 1/0/8, 1/0/9, 1/0/10		Source/Destination MAC, VLAN, Ethertype, Incoming Port
☐ ch4	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype, Incoming Port
☐ ch5	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype, Incoming Port
☐ ch6	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype, Incoming Port
☐ ch7	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype, Incoming Port
☐ ch8	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype, Incoming Port
☐ ch9	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype, Incoming Port
☐ ch10	Static	Enable	Enable	Down	Disable	Disable			Source/Destination MAC, VLAN, Ethertype, Incoming Port

First Previous 1 2 3 4 5 Next Last

Refresh Edit

Figure 504: Port Channel Summary

- Edit Ch1 to include 1/0/1 and 1/0/2.
- Edit Ch2 to include 1/0/3, 1/0/4, 1/0/5, and 1/0/6.
- Edit Ch3 to include 1/0/7, 1/0/8, 1/0/9, and 1/0/10.

Click **Refresh** to reload the page.

9.14.1.2 Configuring the Virtual Port Channel Global Settings

To configure the VPC global settings:

1. Click **Switching > Virtual Port Channel > Global**.

Virtual Port Channel Global Configuration

Domain ID	1 (1 to 255)
VPC Mode	☐ Disable ☒ Enable
Operational VPC Mode	Disabled
VPC State	Disabled
VPC MAC	(xxxxxx.xxxxxx.xxxx)
Operational VPC MAC	
VPC System Priority	32767 (1 to 65535)
Operational VPC System Priority	
Self Rule	None
Local System MAC	00:40:C7:1D:70:98

Keepalive Parameters

Keepalive Priority	10 (1 to 255)
Keepalive Timeout (Seconds)	5 (2 to 15)
Keepalive Mode	☒ Disable ☐ Enable

Peer

Domain ID	
VPC MAC	
Operational VPC MAC	
VPC System Priority	
Operational VPC System Priority	
Peer Rule	None
System MAC	

Peer Link

Port Channel	0/3/3
Peer Link Status	Down
Peer Link STP Mode	Enabled
Configured VLANs	1
Egress Tagging	None

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	☐ Disable ☒ Enable
Tx Interval (msecs)	1000 (200 to 4000)
Operational Tx Interval (msecs)	1000
Rx Timeout (msecs)	3500 (700 to 14000)
Operational Rx Timeout (msecs)	3500

Submit Refresh Remove Cancel

Figure 505: Virtual Port Channel Global Configuration

- Enter **1** for the Domain ID.
- Select **Enable** for the VPC Mode.
- Enter **10** for the Keepalive Priority.
- Click **Edit** in the Peer Link section to select the Peer Link port.
- 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 1 to 65535 (Default is 60000).
- Enter 192.168.0.2 as the Source IP Address. This is the address used by DCPDP.
- Click **Submit**.

9.14.1.3 Configuring the Virtual Port Channel Interface

To configure the VPC interface settings:

- Click **Switching > Virtual Port Channel > Interface Configuration**.

Virtual Port Channel Interface Configuration

Display All rows Showing 1 to 2 of 2 entries Filter:

	Interface	VPC Identifier	Operational Mode	VPC Interface State	Configured VLANs
☐	0/3/1	1	Disabled	Disabled	1
☐	0/3/2	2	Disabled	Disabled	1

First Previous 1 Next Last

Refresh Add Details Remove

Figure 506: Virtual Port Channel Interface Configuration

2. Click **Add** to add both (0/3/1) and (0/3/2).

9.14.2 Configuring Virtual Port Channel Using the CLI

To configure VPC using the CLI interface:

1. Enter VLAN data base mode and create the VPC VLANs.

```
(Routing) #vlan database
(Routing) (Vlan) #vlan 10-17
```

2. Create the VLAN routing interface that will be used for the Dual Control Plane Detection Protocol.

```
(Routing) (Vlan) #vlan 100
(Routing) (Vlan) #vlan routing 100
(Routing) (Vlan) #exit
```

3. Enable the VPC feature.

```
(Routing) #config
(Routing) (Config) #feature vpc
```

4. Enable the keepalive protocol.

```
(Routing) #config
(Routing) (Config) #vpc domain 1
(Routing) (Config-VPC 1) #peer-keepalive enable
(Routing) (Config-VPC 1) #
```

5. Configure the VPC role priority.

```
(Routing) (Config) #vpc domain 1
(Routing) (config-VPC 1) #role priority 10
```

6. Create LAG 1.

```
(Routing) (Config) #interface lag 1
(Routing) (Interface lag 1) #description "VPC-Peer-Link"
```

7. Disable spanning tree on the LAG.

```
(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.

```
(Routing) (Interface lag 1) #vlan participation include 1-99
(Routing) (Interface lag 1) #vlan tagging 1-99
(Routing) (Interface lag 1) #vlan acceptframe vlanonly
(Routing) (Interface lag 1) #vpc peer-link
(Routing) (Interface lag 1) #exit
```

9. Create the peer link.

```
(Routing) (Config) #interface 1/0/1-1/0/2
(Routing) (Interface 1/0/1-1/0/2) #addport lag 1
(Routing) (Interface 1/0/1-1/0/2) #description "VPC-Peer-Link"
```

10. Enable UDLD (if required).

```
(Routing) (Interface 1/0/1-1/0/2) #udld enable
(Routing) (Interface 1/0/1-1/0/2) #udld port aggressive
(Routing) (Interface 1/0/1-1/0/2) #exit
```

11. Configure Dual Control Plane Detection Protocol Configuration (if required):

- a. Configure a VLAN routing interface and assign a local IP address (independent form the peer address).

```
(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).

```
(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.

```
(Routing) (Interface vlan 100) #exit
```

c. Configure the keepalive source and destination IP address.

```
(Routing) #config
(Routing) #vpc domain 1
(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.

```
(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 63.

```
(Routing) (Config) #interface 1/0/3-1/0/6
(Routing) (Interface 1/0/3-1/0/6) #addport lag 2
(Routing) (Interface 1/0/3-1/0/6) #exit

(Routing) (Config) #interface 1/0/7-1/0/10
(Routing) (Interface 1/0/7-1/0/10) #addport lag 3
(Routing) (Interface 1/0/7-1/0/10) #exit

(Routing) (Config) #interface lag 2
(Routing) (Interface lag 2) #vlan participation include 1-100
(Routing) (Interface lag 2) #vlan tagging 1-100
(Routing) (Interface lag 2) #vlan acceptframe vlanonly
(Routing) (Interface lag 2) #vpc 1
(Routing) (Interface lag 2) #exit

(Routing) (Config) #interface lag 3
(Routing) (Interface lag 3) #vlan participation include 1-100
(Routing) (Interface lag 3) #vlan tagging 1-100
(Routing) (Interface lag 3) #vlan acceptframe vlanonly
(Routing) (Interface lag 3) #vpc 2
(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.

9.15 Bidirectional Forwarding Detection

- 4** 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.

9.15.1 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 LCOS SX 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 LCOS SX implementation works with IPv4 and IPv6 networks and supports IPv4 and IPv6 address-based encapsulations.

9.15.2 Configuring BFD

The following command sequence enables BFD and configures session parameters:

1.

```
(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
```

 - a 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).
 - a 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).
 - a The argument **multiplier** specifies the number of BFD control packets to be missed in a row to declare a session down.
 - a 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
```

9.16 Interactive SSH

The Interactive SSH feature allows a remote client to send `configuration` and `show` commands to the switch and receive the output on the client. Some commands must be run on the switch to configure ssh handling and some commands are needed on the client to generate the ssh public

and private key pair, then install the public key on the switch. The following example shows how this feature can be used with a Linux host.

1. On the Linux host, generate the public/private key pair. There are many options to control how the keys are generated.

The following command generates a 2048-bit RSA key: `ssh-keygen -t rsa -b 2048:`

```
(Localhost):~$ ssh-keygen -t rsa -b 2048
Generating public/private rsa key pair.
Enter file in which to save the key (/home/(user)/.ssh/id_rsa):
/home/(user)/.ssh/id_rsa already exists.
Overwrite (y/n)? y
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/(user)/.ssh/id_rsa.
Your public key has been saved in /home/(user)/.ssh/id_rsa.pub.
The key fingerprint is:
c5:05:66:38:79:5c:d8:16:08:27:2d:2d:2c:e1:c8:31 (user)@i06-35
```

This generates the public and private key pair and puts the files in `/home/(username)/.ssh/` where the username is the name of the Linux user.

2. On the switch, make sure the switch has the rsa/dsa keys configured and the ssh/scp services configured.

```
(Routing)#configure
(Routing)(Config)#crypto key generate rsa
(Routing)(Config)#crypto key generate dsa
(Routing)(Config)#exit
(Routing)#ip ssh server enable
(Routing)#ip ssh pubkey-auth
(Routing)#ip scp server enable
```

Verify ip ssh config:

```
(Routing) #show ip ssh
SSH Configuration
Administrative Mode: ..... Enabled
SSH Port: ..... 22
Protocol Level: ..... Version 2
SSH Sessions Currently Active: ..... 0
Max SSH Sessions Allowed: ..... 5
SSH Timeout (mins): ..... 5
Keys Present: ..... DSA RSA
Key Generation In Progress: ..... None
SSH Public Key Authentication Mode: ..... Enabled
SCP server Administrative Mode: ..... Enabled
```

3. Copy the public key that was generated on the Linux client to the switch (given the user's input of the password) so it can be used for a specific login user. In this case, copy the public key created in `.ssh/id_rsa.pub` to the switch as `toaster.pub`.

```
(user)@i06-35:~$ scp .ssh/id_rsa.pub admin@10.27.21.251:toaster.pub
The authenticity of host '10.27.21.251 (10.27.21.251)' can't be established.
RSA key fingerprint is f7:aa:e5:d4:d3:d7:f8:85:5f:7a:95:57:58:2b:6e:dd.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.27.21.251' (RSA) to the list of known hosts.
admin@10.27.21.251's password:
id_rsa.pub 100% 395 0.4KB/s 00:00
```

4. Configure the public key on the switch (given the user's input of the password) so it can be used for a specific login user.

```
(user)@i06-35:~$ ssh -F ssh_config admin@10.52.135.124 "config; username \"admin\" sshkey file toaster.pub"
admin@10.52.135.124's password:
```

5. Test the connection with a simple command, such as `show version`.

```
(use)@i06-35:~$ ssh admin@10.27.21.251"show version"
Switch: 1
System Description..... x86_64-quanta_common_rglbmc-r0, 3.4.4.2,
Linux 4.4.117-ceeeb99d, 2016.05.00.04
Machine Type..... x86_64-quanta_common_rglbmc-r0
Machine Model..... BES-53248
Serial Number..... QTFCU39030018
Maintenance Level..... A
Manufacturer..... 0xbc00
```

9 Configuration Examples

```
Burned In MAC Address..... D8:C4:97:A5:7C:E8
Software Version..... 3.4.4.2
Operating System..... Linux 4.4.117-ceeeb99d
Network Processing Device..... BCM56873_A0
CPLD version..... 0xff040c03
Additional Packages..... BGP-4
    QOS
    Multicast
    IPv6
    Routing
    Data Center
    OpEN API
    Prototype Open API
```