

GS950 PS V2 Series


Gigabit Ethernet PoE+ Switches

GS950/10PS V2

GS950/18PS V2

GS950/28PS V2

GS950/52PS V2



User Guide

Copyright © 2021 Allied Telesis, Inc.

All rights reserved. No part of this publication may be reproduced without prior written permission from Allied Telesis, Inc.

Allied Telesis and the Allied Telesis logo are trademarks of Allied Telesis, Incorporated. All other product names, company names, logos or other designations mentioned herein are trademarks or registered trademarks of their respective owners.

Allied Telesis, Inc. reserves the right to make changes in specifications and other information contained in this document without prior written notice. The information provided herein is subject to change without notice. In no event shall Allied Telesis, Inc. be liable for any incidental, special, indirect, or consequential damages whatsoever, including but not limited to lost profits, arising out of or related to this manual or the information contained herein, even if Allied Telesis, Inc. has been advised of, known, or should have known, the possibility of such damages.

Contents

Preface	19
Document Conventions	20
Contacting Allied Telesis	21
Section I: Basic Switch Management	23
Chapter 1: Introduction	25
Introduction.....	26
Switch Models.....	26
Front Panels.....	26
Hardware Features.....	28
Copper Ports.....	28
Power Over Ethernet Plus.....	28
SFP Slots	28
System LEDs	29
Port LEDs.....	29
Installation Options.....	29
Power Conservation.....	29
MAC Address Table.....	30
Management Interfaces	30
Command Line Interface	31
Starting the First Management Session	32
Starting the First Management Session with a DHCP Server.....	33
Starting the First Management Session with a Direct Connection.....	33
Starting the First Management Session without a DHCP Server.....	34
Starting a Web Browser Management Session.....	36
Starting a Management Session	36
Main Menu.....	37
Displaying General Switch Information.....	38
Saving Your Changes to a Configuration File	41
Unsupported Special Characters.....	43
Chapter 2: Port LEDs and Basic Port and Switch Settings	45
Viewing the Front Panel of the Switch.....	46
SYSTEM LED	46
POWER LED.....	47
PoE MAX LED.....	47
Copper Port LEDs.....	47
SFP Port LEDs.....	50
Spanning Tree LEDs.....	50
Configuring Basic Port Settings.....	51
Setting the Name, Location, and Contact.....	56

Chapter 3: IPv4 Address	59
IPv4 Address Overview.....	60
Manually Setting the IPv4 Address	61
Setting the IPv4 Address from a DHCP or BOOTP Server.....	64
Chapter 4: IPv6 Address	65
IPv6 Address Overview.....	66
Setting the IPv6 Address	67
Setting the IPv6 Neighbor Settings	71
Chapter 5: DHCP Auto Configuration	73
Enabling or Disabling DHCP Auto Configuration	74
Chapter 6: System Time	77
Setting the Date and Time	78
Manually Setting the Date and Time	78
Setting the Date and Time from an SNTP Server	80
Chapter 7: Management Security	83
Adding or Deleting Manager Accounts	84
Configuring the Management Interfaces	87
Enabling or Disabling HTTP or HTTPS Web Browser Management	89
Enabling or Disabling the Secure Shell Server for the Command Line Interface.....	91
Identifying Authorized Management Workstations.....	93
Chapter 8: Configuration Files	97
Overview to Switch Configuration Files	98
Backing Up Configuration Files from the Switch with HTTP	99
Restoring Configuration Files to the Switch with HTTP	101
Backing Up Configuration Files from the Switch with TFTP.....	102
Restoring Configuration Files to the Switch with TFTP.....	104
Booting the Switch with the Config2 Configuration File	106
Chapter 9: LED ECO Mode and Energy-Efficient Ethernet	107
Enabling or Disabling the LED ECO Mode	108
Enabling or Disabling IEEE 802.3az Energy-Efficient Ethernet	109
Chapter 10: Power over Ethernet (PoE)	111
Overview	112
Maximum PoE Budgets	112
PoE Standards	113
Powered Device Classes.....	113
PoE Port Priorities	113
Configuring PoE on the Ports	115
Chapter 11: Link Layer Discovery Protocol	119
LLDP Overview	120
Configuring LLDP.....	121
Displaying Neighbor Information.....	124

Section II: Spanning Tree Protocols	127
Chapter 12: Spanning Tree and Rapid Spanning Tree Protocols	129
STP and RSTP Overview	130
Bridge Priority and the Root Bridge.....	130
Forwarding Delay and Topology Changes.....	133
Mixed STP and RSTP Networks	135
Spanning Tree and VLANs	136
Configuring STP and RSTP Global Settings	138
Configuring STP and RSTP Port Settings	142
Chapter 13: Multiple Spanning Tree Protocol Overview	147
Overview.....	148
Multiple Spanning Tree Instance (MSTI).....	149
VLAN and MSTI Associations.....	149
Ports in Multiple MSTIs	149
Multiple Spanning Tree Regions	151
Region Guidelines.....	153
Common and Internal Spanning Tree (CIST).....	154
MSTP with STP and RSTP.....	155
Summary of Guidelines	156
Associating VLANs to MSTIs.....	158
Connecting VLANs Across Different Regions	160
Chapter 14: Multiple Spanning Tree Protocol	163
Configuring MSTP Global Settings.....	164
Configuring MSTP Port Settings.....	168
Configuring MST and MSTI Settings	172
Configuring MST Port Settings.....	175
Displaying MST Instance Information.....	178
Section III: Virtual LANs	181
Chapter 15: Port-Based Virtual LANs	183
VLAN Overview	184
Port-Based VLAN Overview	186
Port-Based Example 1	187
Port-Based Example 2	188
Adding or Viewing Port-Based VLANs	190
Assigning the Port-Based VLAN Mode to Ports	193
Modifying Port-Based VLANs	196
Deleting Port-Based VLANs	198
Step 1.....	198
Step 2.....	198
Selecting the MAC Address Forwarding Table Mode.....	200
Viewing the MAC Addresses of Network Devices in the Forwarding Table	202
Chapter 16: 802.1Q Tagged Virtual LANs	205
802.1Q Tagged VLAN Overview	206
Tagged Ports.....	206
Untagged Ports	207
802.1Q Tagged VLAN Components	207
Guidelines to Adding Tagged VLANs.....	209
Tagged VLAN Example.....	209
Adding or Viewing 802.1Q Tagged VLANs	212

Assigning the 802.1Q Tagged VLAN Mode to Ports.....	216
Configuring PVIDs and Filters for Tagged and Untagged Ports	219
Modifying 802.1Q Tagged VLANs	221
Deleting 802.1Q Tagged VLANs.....	222
Viewing Port-based and 802.1Q Tagged VLANs.....	223
Chapter 17: GARP VLAN Registration Protocol	225
GVRP Overview	226
Guidelines.....	228
GVRP and Network Security	228
GVRP-inactive Intermediate Switches.....	229
Enabling or Disabling GVRP.....	230
Configuring GVRP Port Settings.....	231
Configuring GVRP Time Settings	233
Chapter 18: Private Virtual LAN	235
Private VLAN Overview	236
Source Port.....	236
Forwarding Ports	236
Configuring the Private VLAN	237
Disabling the Private VLAN.....	240
Chapter 19: Voice VLAN	241
Voice VLAN Overview.....	242
802.1Q VLAN	242
Organization Unique Identifier (OUI)	242
CoS with Voice VLAN.....	242
Dynamic Port Auto-Detection	243
General Guidelines.....	244
Configuring the Voice VLAN	245
Managing the OUI Table	248
Adding OUI Entries.....	248
Modifying OUI Entries.....	249
Deleting OUI Entries.....	249
Chapter 20: MAC Address Forwarding Table	251
Viewing the MAC Addresses of Network Devices in the Forwarding Table.....	252
Selecting the MAC Address Forwarding Table Mode	254
Section IV: Network Traffic Management	257
Chapter 21: Traffic Statistics and Charts	259
Network Statistics Overview	260
Displaying Port Traffic Comparison Statistics	261
Displaying Port Error Statistics	265
Displaying Historical Statistics	266
Chapter 22: Static Unicast MAC Addresses	269
Static MAC Addresses Overview	270
Adding Static Unicast MAC Addresses.....	272
Modifying Static Unicast MAC Addresses.....	275
Deleting Static Unicast MAC Addresses.....	278

Chapter 23: Static Multicast MAC Addresses	279
Adding Static Multicast MAC Addresses	280
Modifying Static Multicast MAC Addresses	283
Deleting Static Multicast Addresses	286
Chapter 24: IGMP Snooping	287
IGMP Snooping Overview	288
Configuring IGMP Snooping	290
Adding or Deleting Static Multicast Router Ports	293
Viewing Multicast Addresses and Host Ports	297
Chapter 25: Quality of Service and Class of Service	299
QOS and COS Overview	300
Packet Priority	300
Egress Queue vs Packet Priority Mapping	301
Prioritizing Untagged Packets	302
Scheduling Algorithm	302
Mapping CoS Priorities to Egress Queues	304
Mapping CoS Priorities to Ports	306
Mapping DSCP Classes to Egress Queues	307
Setting the Queue Scheduling Algorithm	309
Mapping IPv6 Traffic Classes to Port Egress Queues	311
Chapter 26: Access Control Policies	313
Access Control Policies Overview	314
Policy Filters	314
Actions	314
How Ingress Packets are Compared Against Policies	315
Guidelines	315
Adding Layer 2 IPv4 Policies	317
Adding IPv6 Policies	325
Enabling or Disabling Policies	329
Modifying or Deleting Policies	330
Displaying Policies	331
Displaying Policies by Port Numbers	332
Managing Rate Control Settings	333
Adding Rate Control Entries	333
Modifying or Deleting Rate Control Entries	334
Chapter 27: Bandwidth Control	337
Setting Threshold Limits for Ingress Unknown Unicast, Broadcast, and Multicast Packets	338
Setting Ingress Traffic Rate Limits	340
Setting Egress Traffic Rate Limits	342
Chapter 28: RMON	345
RMON Overview	346
Enabling or Disabling RMON	347
RMON Port Statistics Groups	348
Adding RMON Port Statistics Groups	348
Modifying or Deleting RMON Port Statistics Groups	349
RMON Histories	350
Adding RMON History Groups	350
Modifying or Deleting RMON History Groups	352
RMON Events	353
Adding Events	353
Modifying or Deleting Events	355

RMON Alarms	356
Adding RMON Alarms	357
Modifying or Deleting Alarms	359
Section V: Port Trunks	361
Chapter 29: Static Port Trunks	363
Static Port Trunk Overview	364
Guidelines	364
Adding Static Port Trunks	366
Modifying Static Port Trunks	368
Deleting Static Port Trunks	370
Chapter 30: LACP Trunks	371
LACP Port Trunk Overview	372
System Priority and ID Numbers	372
Port Priority Value	373
General Guidelines	374
Displaying LACP Group Status	375
Adding LACP Trunks	377
Setting LACP Port Priorities	379
Modifying LACP Trunks	380
Deleting LACP Trunks	382
Section VI: Simple Network Management Protocol	383
Chapter 31: SNMPv1 and v2c	385
Enabling or Disabling SNMP Management	386
SNMPv1 and SNMPv2c User and Group Names	387
Adding User and Group Names	387
Modifying User and Group Names	388
Deleting User and Group Names	389
SNMP Community Strings	390
Adding SNMP Community Strings	390
Modifying SNMP Community Strings	391
Deleting SNMP Community Strings	391
SNMP Traps	392
Adding Trap Host Table Entries	392
Modifying Trap Host Table Entries	394
Deleting Trap Host Table Entries	395
Chapter 32: SNMPv3	397
SNMPv3 Overview	398
SNMPv3 Authentication Protocols	398
SNMPv3 Privacy Protocol	399
SNMPv3 MIB Views	399
SNMPv3 Configuration Process	400
SNMPv3 User and Group Names	402
Adding SNMPv3 User and Group Names	402
Modifying SNMPv3 User and Group Names	403
Deleting SNMPv3 User and Group Names	404
SNMPv3 View Names	405
Adding SNMPv3 View Names	405
Modifying SNMPv3 View Names	408

Deleting SNMPv3 View Names.....	408
SNMPv3 View Table.....	409
Adding SNMPv3 View Table Entries.....	409
Modifying SNMPv3 View Table Entries.....	410
Deleting SNMPv3 View Table Entries.....	411
SNMPv3 Traps	412
SNMP Engine ID	413
Modifying SNMP Engine ID.....	413
Resetting SNMP Engine ID.....	413
Section VII: Network Security	415
Chapter 33: Port Authentication	417
Overview.....	418
Authentication Devices	419
Authentication Methods	420
802.1x Port-based Network Access Control	420
MAC Address-based Authentication	420
Authenticator Port Operational Settings	421
Authenticator Port Operating Modes	422
Single Host Mode.....	422
Single Host Mode with Piggy Backing.....	422
Multiple Host Mode	424
Supplicant and VLAN Associations	426
Single Host Mode.....	426
Multiple Host Mode	427
Multiple Supplicant Mode	427
Supplicant VLAN Attributes on RADIUS Servers.....	427
Guest VLAN.....	428
RADIUS Accounting	429
General Steps.....	430
Guidelines.....	431
Configuring Port Access Control	433
Chapter 34: Local Dial-in User Accounts	439
Local Dial-in User Authentication Overview	440
Adding Local Dial-in User Accounts	441
Modifying or Deleting Local Dial-in User Accounts.....	443
Chapter 35: RADIUS Client	445
RADIUS Client Overview.....	446
Adding Server IP Addresses to the RADIUS Client.....	447
Modifying or Deleting Server IP Addresses from the RADIUS Client.....	449
Configuring RADIUS Accounting.....	450
Chapter 36: TACACS+ Client	451
TACACS+ Client Overview.....	452
Adding Server IP Addresses to the TACACS+ Client.....	453
Modifying or Deleting Server IP Addresses from the TACACS+ Client.....	455
Chapter 37: DHCP Snooping	457
DHCP Snooping Overview	458
Trusted Ports.....	458
Untrusted Ports	458
Unauthorized DHCP Servers	458

DHCP with Option 82	459
Option 82 Pass Through.....	460
General Guidelines	463
Configuring DHCP Snooping	464
Adding DHCP Snooping to VLANs	467
Designating Trusted and Untrusted Ports	469
Managing the Binding Database	471
Chapter 38: Destination MAC Address Filters	475
Destination MAC Address Filters Overview	476
Adding Destination MAC Address Filters	477
Deleting Destination MAC Address Filters	478
Chapter 39: Port Mirroring	479
Port Mirroring Overview	480
Enabling Port Mirroring	481
Disabling Port Mirroring	483
Chapter 40: Denial of Service	485
Configuring Denial of Service Protection	486
 Section VIII: Troubleshooting Tools and Maintenance	 489
Chapter 41: System Log and Syslog Client	491
Viewing the Event System Log	492
Sending System Log Events to a Syslog Server	495
Chapter 42: Troubleshooting Tools	499
Rebooting the Switch	500
Restoring the Factory Default Values	502
Locking the Factory Reset Options	503
Disabling the eco-Friendly Button and Default Settings	503
Enabling the eco-Friendly Button and Default Settings	504
Pinging Network Devices	505
Chapter 43: Loopback Detection	507
Configuring Loopback Detection	508
Disabling Loopback Detection	511
Chapter 44: Switch Firmware	513
Management Software Overview	514
Upgrading the Management Software with a Web Browser	515
Upgrading the Management Software with TFTP	518
Designating the Boot Management Software	520
Appendix A: Open Source Software	523

Figures

Figure 1: Front Panel of the GS950/10PS V2 Switch	26
Figure 2: Front Panel of the GS950/18PS V2 Switch	26
Figure 3: Front Panel of the GS950/28PS V2 Switch	27
Figure 4: Front Panel of the GS950/52PS V2 Switch	27
Figure 5: Login Window	33
Figure 6: IPv4 Setup	34
Figure 7: Main Menu	37
Figure 8: Example of a Sub-menu	37
Figure 9: Switch Info Menu Selection	38
Figure 10: Switch Information Window	38
Figure 11: Save Settings to Flash Menu Selection	41
Figure 12: Save Settings to Flash Window	42
Figure 13: Front Panel Menu Selection	46
Figure 14: Front Panel Window	46
Figure 15: PoE MAX and SYS LEDs	46
Figure 16: POWER LED	47
Figure 17: Copper Port LEDs	48
Figure 18: PoE Copper Ports 1 to 24 LEDs on the GS950/52PS V2 Switch	48
Figure 19: Non-PoE Copper Ports 25 to 48 LEDs on the GS950/52PS V2 Switch	49
Figure 20: Physical Interface Menu Selection	51
Figure 21: Physical Interface Window	51
Figure 22: Management Menu Selection	56
Figure 23: Management Window	56
Figure 24: IPv4 Setup Menu Selection	61
Figure 25: IPv4 Setup Window	62
Figure 26: IPv6 System Settings Menu Selection	67
Figure 27: IPv6 System Settings Window	68
Figure 28: IPv6 Neighbor Settings Menu Selection	71
Figure 29: IPv6 Neighbor Settings Window	71
Figure 30: DHCP Auto Config Menu Selection	74
Figure 31: DHCP Auto Configuration Window	75
Figure 32: System Time Menu Selection	78
Figure 33: System Time Window	79
Figure 34: Administration Menu Selection	84
Figure 35: Administration Window	85
Figure 36: Modify Administration Window	86
Figure 37: User Interface Menu Selection	87
Figure 38: User Interface Window	87
Figure 39: SSL Settings Menu Selection	89
Figure 40: SSH Settings	90
Figure 41: SSH Menu Selection	91
Figure 42: SSH Settings Window	92
Figure 43: IP Access List Menu Selection	93
Figure 44: IP Access List Window	94
Figure 45: Config File Backup/Restore via HTTP Menu Selection	99
Figure 46: Config File Backup/Restore via HTTP Window	99
Figure 47: Config File Backup/Restore via TFTP	102
Figure 48: Config File Backup/Restore via TFTP Window	102
Figure 49: LED ECO Mode Menu Selection	108

Figure 50: LED ECO Mode Page.....	108
Figure 51: IEEE 802.3az EEE Menu Selection.....	109
Figure 52: IEEE 802.3az EEE Window.....	109
Figure 53: Power Over Ethernet Menu Selection.....	115
Figure 54: Power Over Ethernet Window.....	115
Figure 55: LLDP Global Settings.....	121
Figure 56: LLDP Global Settings Window.....	121
Figure 57: LLDP Neighbors Information Menu Selection.....	124
Figure 58: LLDP Neighbors Information Window.....	124
Figure 59: Point-to-Point Ports.....	135
Figure 60: Edge Port.....	135
Figure 61: STP and VLAN Fragmentation with Untagged Ports.....	136
Figure 62: STP and VLAN Fragmentation with Untagged Ports.....	137
Figure 63: STP and VLAN Compatibility with Tagged Ports.....	137
Figure 64: Protocol Settings Menu Selection.....	138
Figure 65: Spanning Tree Protocol Settings Window.....	139
Figure 66: Port Settings Menu Selection.....	142
Figure 67: Port Settings Window.....	142
Figure 68: Multiple Spanning Tree Region.....	152
Figure 69: CIST and VLAN Guideline - Example 1.....	158
Figure 70: CIST and VLAN Guideline - Example 2.....	159
Figure 71: Spanning Regions.....	160
Figure 72: Protocol Settings Menu Selection.....	164
Figure 73: Spanning Tree Port Settings Menu Selection.....	168
Figure 74: MST Settings Menu Selection.....	172
Figure 75: MST Settings Window.....	172
Figure 76: MST Port Settings.....	175
Figure 77: MST Port Settings Page.....	175
Figure 78: Instance Information Menu Selection.....	178
Figure 79: Instance Information Window.....	178
Figure 80: Port-Based VLAN - Example 1.....	187
Figure 81: Port-Based VLAN - Example 2.....	188
Figure 82: Port-Based VLAN Menu Selection.....	190
Figure 83: Port-Based VLAN Window.....	191
Figure 84: VLAN Mode Menu Selection.....	193
Figure 85: VLAN Mode Window.....	194
Figure 86: Example of the VLAN Mode Window.....	194
Figure 87: Example of the Modify Port-Based VLAN Window.....	196
Figure 88: Forwarding Table Mode Menu Selection.....	200
Figure 89: Forwarding Table Mode Window.....	200
Figure 90: Dynamic Forwarding Table Menu Selection.....	202
Figure 91: Dynamic Forwarding Table Window.....	202
Figure 92: Example of a Tagged VLAN.....	210
Figure 93: Tagged VLAN Menu Selection.....	212
Figure 94: Tagged VLAN Window.....	213
Figure 95: VLAN Mode Menu Selection.....	216
Figure 96: VLAN Mode Window.....	217
Figure 97: Example of the VLAN Mode Window.....	217
Figure 98: Port Settings Menu Selection.....	219
Figure 99: Port Settings Window.....	219
Figure 100: VLAN Current Database Menu Selection.....	223
Figure 101: VLAN Current Database Window.....	223
Figure 102: GVRP Example.....	227
Figure 103: GVRP Global Settings Menu Selection.....	230
Figure 104: GVRP Global Settings Window.....	230
Figure 105: GVRP Port Settings Menu Selection.....	231
Figure 106: GVRP Port Settings Window.....	231
Figure 107: GVRP Time Settings Menu Selection.....	233
Figure 108: GVRP Time Settings Window.....	233
Figure 109: Private VLAN Menu Selection.....	237

Figure 110: Private VLAN Window	238
Figure 111: Voice VLAN Settings Menu Selection	245
Figure 112: Voice VLAN Settings Window	245
Figure 113: Voice VLAN OUI Settings Menu Selection	248
Figure 114: Voice VLAN OUI Settings Window	248
Figure 115: Dynamic Forwarding Table Menu Selection	252
Figure 116: Dynamic Forwarding Table Window	252
Figure 117: Forwarding Table Mode Menu Selection	254
Figure 118: Forwarding Table Mode Window	254
Figure 119: Traffic Comparison Menu Selection	261
Figure 120: Traffic Comparison Chart Window	261
Figure 121: Traffic Comparison Window Example	264
Figure 122: Error Group Menu Selection	265
Figure 123: Error Group Chart Window	265
Figure 124: Historical Status Menu Statistics	266
Figure 125: Historical Status Chart Window	266
Figure 126: Static Unicast Menu Selection	272
Figure 127: Static Unicast Address Table Window	273
Figure 128: Modifying a Static Unicast Address in an 802.1Q Tagged VLAN	276
Figure 129: Modifying a Static Unicast Address in a Port-Based VLAN	276
Figure 130: Static Multicast Menu Selection	280
Figure 131: Static Multicast Address Window	281
Figure 132: Static Multicast Address Table Window for an 802.1Q Tagged VLAN	284
Figure 133: Static Multicast Address Table Window for a Port-Based VLAN	284
Figure 134: IGMP Snooping Settings Menu Selection	290
Figure 135: IGMP Snooping Settings Window	290
Figure 136: IGMP Snooping Router Port Menu Selection	293
Figure 137: IGMP Snooping Router Port Window	294
Figure 138: Modify IGS Static Router Port Window for an 802.1Q Tagged VLAN	295
Figure 139: Modify IGS Static Router Port Window for the Port-Based VLANs	295
Figure 140: Multicast MAC Address Tables in the IGMP Snooping Window	297
Figure 141: CoS Menu Selection	304
Figure 142: CoS Window	305
Figure 143: Port Priority Menu Selection	306
Figure 144: Port Priority Window	306
Figure 145: DSCP Menu Selection	307
Figure 146: DSCP Class Mapping Window	308
Figure 147: Scheduling Algorithm Menu Selection	309
Figure 148: Scheduling Algorithm Window	309
Figure 149: IPv6 Traffic Class Priority Settings Menu Selection	311
Figure 150: IPv6 Traffic Class Priority Settings Window	312
Figure 151: Policy Settings Menu Selection	317
Figure 152: Policy Settings Window	317
Figure 153: IPv4 Policy Settings Window	318
Figure 154: IPv6 Policy Settings Window	325
Figure 155: Policy Database Menu Selection	332
Figure 156: Policy Database Window	332
Figure 157: Rate Control Settings Menu Selection	333
Figure 158: Rate Control Settings Window	334
Figure 159: Storm Control Menu Selection	338
Figure 160: Storm Control Window	339
Figure 161: Ingress Rate Limiting Menu Selection	340
Figure 162: Ingress Rate Limiting Window	340
Figure 163: Egress Rate Limiting Menu Selection	342
Figure 164: Egress Rate Limiting Window	342
Figure 165: RMON Global Settings Menu Selection	347
Figure 166: RMON Basic Settings window	347
Figure 167: RMON Statistics Menu Selection	348
Figure 168: Ethernet Statistics Settings Window	348
Figure 169: RMON History Menu Selection	350

Figure 170: History Control Settings Window	350
Figure 171: RMON Event Menu Selection.....	353
Figure 172: RMON Event Settings Window.....	353
Figure 173: RMON Alarms Menu Selection.....	357
Figure 174: RMON Alarm Settings Window.....	357
Figure 175: Static Port Trunk Example	364
Figure 176: Trunking Menu Selection.....	366
Figure 177: Trunking Window	366
Figure 178: LACP Group Status Menu Selection	375
Figure 179: LACP Group Status Window	375
Figure 180: Trunking Menu Selection.....	377
Figure 181: Trunking Window	377
Figure 182: Port Priority Menu Selection	379
Figure 183: Port Priority Window	379
Figure 184: SNMP User/Group Menu Selection.....	387
Figure 185: SNMP User/Group Window	388
Figure 186: Community Table Menu Selection.....	390
Figure 187: SNMP Community Table Window	391
Figure 188: Trap Management Menu Selection.....	392
Figure 189: Trap Management Window.....	393
Figure 190: MIB Tree	399
Figure 191: SNMPv3 Table Relationships.....	401
Figure 192: SNMP User Group, SNMPv3 Window Example.....	403
Figure 193: Group Access Table Menu Selection	405
Figure 194: SNMP Group Access Table Window	406
Figure 195: SNMP Group Access Table Example for SNMPv3	407
Figure 196: View Table Menu Selection	409
Figure 197: SNMP View Table Window.....	409
Figure 198: SNMP View Table Window Example.....	410
Figure 199: SNMP Engine ID Settings Menu Selection.....	413
Figure 200: SNMP Engine ID Settings Window.....	413
Figure 201: Single Host Mode	422
Figure 202: Multiple Host Operating Mode.....	423
Figure 203: Multiple Supplicant Mode.....	425
Figure 204: Port Access Control Menu Selection	433
Figure 205: Port Access Control Settings Window	433
Figure 206: Port Access Control Settings Window - Port Settings	434
Figure 207: Dial-in User Menu Selection	441
Figure 208: Dial-In User Window	441
Figure 209: RADIUS Menu Selection	447
Figure 210: RADIUS Window	447
Figure 211: RADIUS Accounting Settings Menu Selection.....	450
Figure 212: RADIUS Accounting Global Settings Window	450
Figure 213: TACACS+ Menu Selection	453
Figure 214: TACACS+ Window	453
Figure 215: DHCP with Option 82.....	460
Figure 216: DHCP with Option 82 Pass Through Disabled	461
Figure 217: DHCP with Option 82 Pass Through Enabled.....	462
Figure 218: DHCP Snooping General Settings.....	464
Figure 219: DHCP Snooping General Settings Window.....	464
Figure 220: DHCP Snooping VLAN Settings Menu Selection	467
Figure 221: DHCP Snooping VLAN Settings Window	467
Figure 222: DHCP Snooping Trusted Interfaces Menu Selection.....	469
Figure 223: DHCP Snooping Trusted Interfaces Window.....	469
Figure 224: DHCP Snooping Binding Database Menu Selection	471
Figure 225: DHCP Snooping Binding Database Window	472
Figure 226: Destination MAC Filter Menu Selection.....	477
Figure 227: Destination MAC Filter Window	477
Figure 228: Mirroring Menu Selection.....	481
Figure 229: Mirroring Window.....	481

Figure 230: Denial of Service Menu Selection.....	486
Figure 231: Front Panel Window	486
Figure 232: System Log Settings Menu Selection.....	492
Figure 233: System Log Settings Window.....	493
Figure 234: Reboot Menu Selection	500
Figure 235: Factory Default Reset/Reboot Window	500
Figure 236: Ping Menu Selection	505
Figure 237: Ping Test Settings Window.....	505
Figure 238: Loopback Detection Menu Selection.....	508
Figure 239: Loopback Detection Window.....	508
Figure 240: Firmware Upgrade via HTTP Menu Selection.....	516
Figure 241: Firmware Upgrade via HTTP Window.....	516
Figure 242: Firmware Upgrade via TFTP Menu Selection	518
Figure 243: Firmware Upgrade via TFTP Window	519
Figure 244: Firmware Upgrade via HTTP Menu Selection.....	520
Figure 245: Firmware Upgrade via HTTP Window.....	520
Figure 246: OSS Information Menu Selection	523
Figure 247: OSS Information Window	523

Tables

Table 1: Switch Information Window	39
Table 2: SYS LED	46
Table 3: POWER LED	47
Table 4: PoE MAX LED	47
Table 5: Link and Activity (L/A) LEDs for Copper Ports	49
Table 6: PoE LEDs for the Copper Ports	50
Table 7: Link LEDs for the SFP Ports	50
Table 8: Spanning Tree Protocol Port LEDs	50
Table 9: Physical Interface (Port) Settings	52
Table 10: Management Window	57
Table 11: IPv4 Setup Window	62
Table 12: IPv6 System Settings Window	68
Table 13: Manually Setting the Clock and Calender	79
Table 14: Setting the Calender and Clock from an NTP Server	81
Table 15: User Interface Window	88
Table 16: SSH Settings Window	92
Table 17: PoE Maximum Power Budgets	112
Table 18: IEEE Powered Device Classes	113
Table 19: Power Over Ethernet Window - Power Budget Fields	115
Table 20: Power Over Ethernet Window - Port Table	116
Table 21: LLDP Global Settings	122
Table 22: LLDP Port Settings	122
Table 23: LLDP System Information	123
Table 24: LLDP Neighbors Information Window	124
Table 25: Valid Port Priority Values	132
Table 26: Spanning Tree Protocol Settings Window for STP and RSTP	139
Table 27: STP and RSTP Port Settings	143
Table 28: Spanning Tree Protocol Settings Window for MSTP	164
Table 29: MSTP Port Settings	168
Table 30: MST Settings Window - MST Configuration Identification Settings	173
Table 31: MST Settings Window - MST Instance Settings	173
Table 32: MST Port Settings Window	175
Table 33: Instance Information Window	178
Table 34: Components of Port-Based VLANs	186
Table 35: Example 1 of Port-Based VLANs	187
Table 36: Example 2 of Port-Based VLANs	189
Table 37: Port-Based VLAN Window	191
Table 38: Dynamic Forwarding Table Window	203
Table 39: Components of 802.1Q Tagged VLANs	207
Table 40: Example of Tagged VLANs	211
Table 41: Tagged VLAN Window	214
Table 42: Port Settings Window	220
Table 43: VLAN Current Database Window	224
Table 44: GVRP Port Settings Window	231
Table 45: GVRP Time Settings Window	234
Table 46: Private VLAN Window	238
Table 47: Voice VLAN Settings Window	246
Table 48: Voice VLAN OUI Settings Window	248
Table 49: Dynamic Forwarding Table Window	253

Table 50: Traffic Comparison Chart Statistics	262
Table 51: Graph Colors	263
Table 52: Statistics in the Historical Status Chart Window	267
Table 53: Static Unicast MAC Address Tables	273
Table 54: Static Multicast Address Table	281
Table 55: IGMP Snooping Settings Window	291
Table 56: IGMP Snooping Router Port Window	294
Table 57: Multicast MAC Address Tables in the IGMP Snooping Settings Window	297
Table 58: Default Mappings Priority Levels to Priority Queues	301
Table 59: Example of Customized Mappings Priority Levels to Priority Queues	301
Table 60: Weighted Round Robin Priority	303
Table 61: IPv4 Policy Settings Window	318
Table 62: IP Protocol Numbers	322
Table 63: IPv6 Policy Settings Window	326
Table 64: Ethernet Statistics Settings Window	349
Table 65: History Control Settings Window	351
Table 66: RMON Event Settings Window	354
Table 67: RMON Alarm Settings Window	357
Table 68: LACP Group Status Window	375
Table 69: SNMP User/Group Window for v1/v2c	388
Table 70: SNMP Community Table Window	391
Table 71: Trap Management Window	394
Table 72: SNMP User/Group Window for v3	402
Table 73: SNMP Group Access Table Window for v3	406
Table 74: SNMP View Table Window for v3	410
Table 75: Port Access Control Settings Window	433
Table 76: Port Access Control Settings Window - Advanced Settings	435
Table 77: Dial-In User Window	442
Table 78: RADIUS Window	447
Table 79: TACACS+ Window	454
Table 80: DHCP Snooping General Settings Window	464
Table 81: DHCP Snooping Binding Database Window	472
Table 82: DHCP Snooping Binding Database Window	473
Table 83: Denial of Service Protection	487
Table 84: System Log Settings Window - Syslog Client	496
Table 85: Ping Test Settings Window	506
Table 86: Ping Test Results Window	506
Table 87: Loopback Detection Global Settings	509

Preface

This guide contains management instructions for the GS950 PS V2 Series of Gigabit Ethernet switches, with Power over Ethernet Plus (PoE+). The preface contains the following sections:

- “Document Conventions” on page 20
- “Contacting Allied Telesis” on page 21

Document Conventions

This document uses the following conventions:

Note

Notes provide additional information.



Caution

Cautions inform you that performing or omitting a specific action may result in equipment damage or loss of data.



Warning

Warnings inform you that performing or omitting a specific action may result in bodily injury.

Contacting Allied Telesis

If you need assistance with this product, you may contact Allied Telesis technical support by going to the Support & Services section of the Allied Telesis web site at **www.alliedtelesis.com/support**. You can find links for the following services on this page:

- ❑ 24/7 Online Support — Enter our interactive support center to search for answers to your product questions in our knowledge database, to check support tickets, to learn about Return Merchandise Authorizations (RMAs), and to contact Allied Telesis technical experts.
- ❑ USA and EMEA phone support — Select the phone number that best fits your location and customer type.
- ❑ Hardware warranty information — Learn about Allied Telesis warranties and register your product online.
- ❑ Replacement Services — Submit a RMA request via our interactive support center.
- ❑ Documentation — View the most recent installation and user guides, software release notes, white papers, and data sheets for your products.
- ❑ Software Downloads — Download the latest software releases for your managed products.

For sales or corporate information, go to **www.alliedtelesis.com/contact** and select your region.

Section I

Basic Switch Management

This section contains the following chapters:

- ❑ Chapter 1, “Introduction” on page 25
- ❑ Chapter 2, “Port LEDs and Basic Port and Switch Settings” on page 45
- ❑ Chapter 3, “IPv4 Address” on page 59
- ❑ Chapter 4, “IPv6 Address” on page 65
- ❑ Chapter 5, “DHCP Auto Configuration” on page 73
- ❑ Chapter 6, “System Time” on page 77
- ❑ Chapter 7, “Management Security” on page 83
- ❑ Chapter 8, “Configuration Files” on page 97
- ❑ Chapter 9, “LED ECO Mode and Energy-Efficient Ethernet” on page 107
- ❑ Chapter 10, “Power over Ethernet (PoE)” on page 111
- ❑ Chapter 11, “Link Layer Discovery Protocol” on page 119

Chapter 1

Introduction

This chapter contains introductory information about the web browser management interface and instructions on how to start a management session. The chapter contains the following sections:

- ❑ “Introduction” on page 26
- ❑ “Hardware Features” on page 28
- ❑ “Command Line Interface” on page 31
- ❑ “Starting the First Management Session” on page 32
- ❑ “Starting a Web Browser Management Session” on page 36
- ❑ “Main Menu” on page 37
- ❑ “Displaying General Switch Information” on page 38
- ❑ “Saving Your Changes to a Configuration File” on page 41
- ❑ “Unsupported Special Characters” on page 43

Introduction

This manual describes the web browser management interface for the GS950 PS V2 Series of Gigabit Ethernet Switches. The instructions explain how to configure the parameter settings and features of the devices, as well as view status information and traffic statistics.

Switch Models This manual applies to the following GS950 PS V2 Switches:

- ❑ GS950/10PS V2
- ❑ GS950/18PS V2
- ❑ GS950/28PS V2
- ❑ GS950/52PS V2

Front Panels Figure 1 illustrates the front panel of the GS950/10PS V2 Switch.

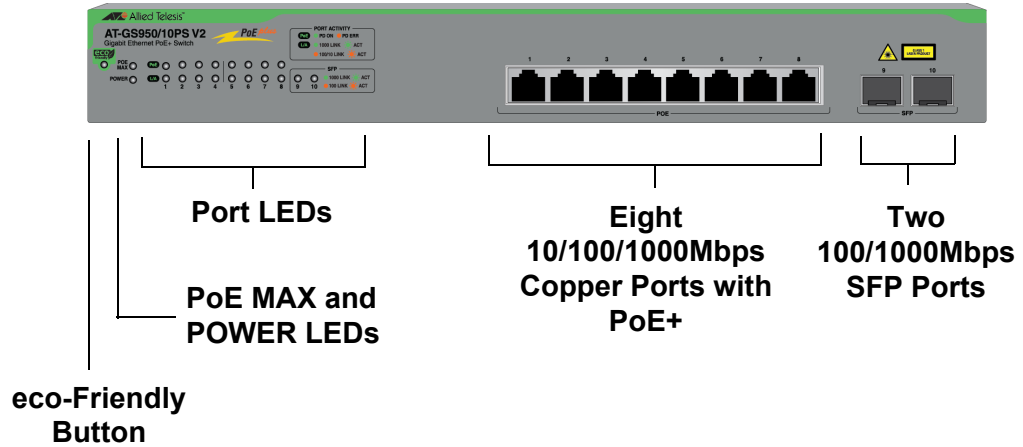


Figure 1. Front Panel of the GS950/10PS V2 Switch

Figure 2 illustrates the front panel of the GS950/18PS V2 Switch.

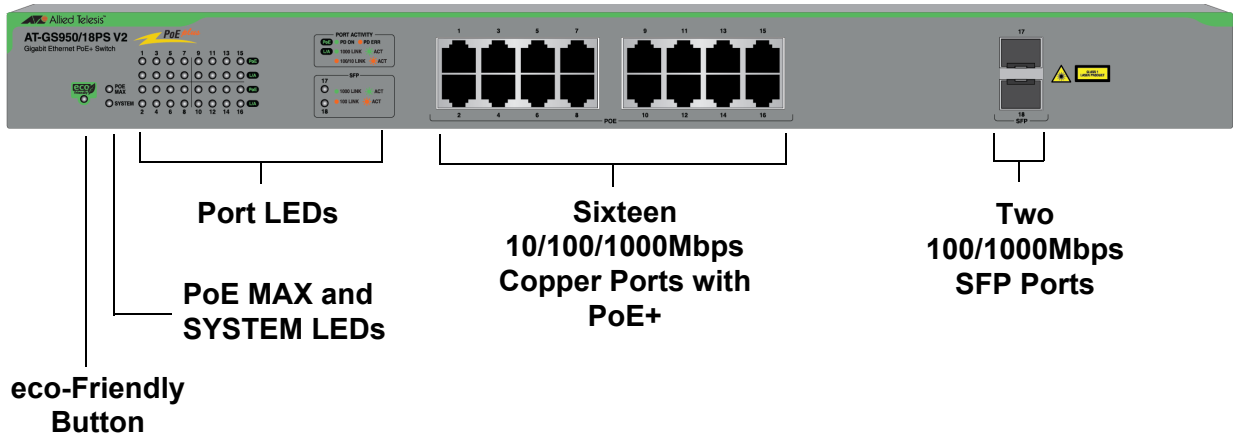


Figure 2. Front Panel of the GS950/18PS V2 Switch

Figure 3 illustrates the front panel of the GS950/28PS V2 Switch.




Figure 3. Front Panel of the GS950/28PS V2 Switch

Figure 4 illustrates the front panel of the GS950/52PS V2 Switch.




Figure 4. Front Panel of the GS950/52PS V2 Switch

Hardware Features

Here are the hardware features of the switches.

Copper Ports

Here are the basic features of the copper ports:

- ❑ 10/100/1000Mbps
- ❑ 10Base-T, 100Base-TX and 1000Base-T compliant
- ❑ IEEE 802.3u Auto-Negotiation compliant
- ❑ Auto-MDI/MDIX
- ❑ 100 meters (328 feet) maximum operating distance
- ❑ IEEE 802.3x Flow Control in 10/100Mbps full-duplex mode
- ❑ IEEE 802.3x Back Pressure in 10/100Mbps half-duplex mode
- ❑ IEEE803.3z 1000Base-T Flow Control
- ❑ Support for Jumbo frames up to 10KB
- ❑ RJ-45 connectors

Power Over Ethernet Plus

Here are the basic PoE+ features:

- ❑ Supported on all copper ports on the GS950/10PS V2, GS950/18PS V2, and GS950/28PS V2 Switches
- ❑ Supported on copper ports 1 to 24 on the GS950/52PS V2 Switch
- ❑ Maximum power budgets:
 - GS950/10PS V2: 75W
 - GS950/18PS V2: 185W
 - GS950/28PS V2: 185W
 - GS950/52PS V2: 370W
- ❑ PoE (15.4W af port) and PoE+ (30W at port)
- ❑ Powered device classes 0 to 4
- ❑ Per port power control
- ❑ Port priorities to control power distribution
- ❑ Supports hardware 2-event detect
- ❑ Alternative A mode wiring (pins 1, 2, 3, and 6)

SFP Slots

The SFP ports support the following types of transceivers:

- ❑ 100Mbps SFP 100Base-FX fiber optic transceivers
- ❑ 1000Mbps SFP 1000Base-SX/LX fiber optic transceivers

- Single-port bidirectional 1000Mbps SFP 1000Base-SX/LX fiber optic transceivers
- 1000Mbps 1000Base-T copper transceiver

Note

Transceivers are purchased separately. For a list of supported transceivers, refer to the product's data sheet on the Allied Telesis web site.

System LEDs

The system LEDs are listed here:

- Power SYS LED. Refer to "SYSTEM LED" on page 46 and "POWER LED" on page 47.
- Maximum PoE+ LED. Refer to "PoE MAX LED" on page 47.

The system LEDs are described in "Viewing the Front Panel of the Switch" on page 46.

Port LEDs

The port LEDs are briefly described here:

- Link/activity and PoE+ LEDs for the copper ports. Refer to "Copper Port LEDs" on page 47.
- Link/activity LEDs for the SFP ports. Refer to "SFP Port LEDs" on page 50.
- Spanning tree protocol LEDs for the copper and SFP ports. Refer to "Spanning Tree LEDs" on page 50
- eco-Friendly button for turning off the LEDs to conserve electricity. Refer to "Enabling or Disabling the LED ECO Mode" on page 108.

The port LEDs are described in "Viewing the Front Panel of the Switch" on page 46.

Installation Options

The switch supports the following installation options:

- Standard 19-inch equipment rack
- Desk or table
- Wall

Power Conservation

The switch has the following power conservation features:

- eco-Friendly button to turn off the port LEDs when the system is not being monitored
- IEEE 802.3az Energy-Efficient Ethernet (EEE) that reduces power during periods of no data activity
- High-efficiency power supply

MAC Address Table

Here are the basic features of the MAC address table:

- Storage capacity up to 16K MAC address entries
- Automatic learning and aging

Management Interfaces

The switches have these three management interfaces:

- Web browser
- Command line
- SNMPv1, v2c, and v3

Note

You can use the web browser interface to configure all of the features of the switches. The command line and SNMP interface support only a subset of features.

Command Line Interface

You can perform the following management functions from the command line interface:

- Display basic switch information, such as the management software version number and IPv4 and IPv6 addresses.
- Return all parameter settings to their default values.
- Reboot the switch.
- Add local manager accounts or change the passwords of existing accounts.
- Enable or disable the HTTP or HTTPS secure server.
- Configure IPv4 and IPv6 management addresses.
- Download new management software to the switch.
- Upload or download configuration files.
- Configure SNMPv1, v2c, and v3.
- Ping other network devices.

You have to use the web browser management interface to perform all other management functions.

Starting the First Management Session

When you power on the GS950 PS V2 Switch at its default setting (Static IP address 192.168.1.1) you can change the IP setting to DHCP client mode. When you switch to DHCP client mode, the switch begins transmitting DHCP queries for a DHCP server on all its ports. If a DHCP server responds, the unit uses the IP address the server assigns to it. If there is no DHCP server, the switch uses the default IP address 192.168.1.1.

If your network has a DHCP server, use the IP address the server assigns it to start the management session. Refer to “Starting the First Management Session with a DHCP Server” on page 33.

If your network does not have a DHCP server, or if you want to configure the unit before connecting it to your network, you can start the first management session by creating a direct connection between your management workstation and the switch. This involves connecting an Ethernet cable to the Ethernet port on the computer and any port on the switch. This procedure requires changing the IP address on your computer to make it a member of the same subnet as the default IP address on the switch. Refer to “Starting the First Management Session with a Direct Connection” on page 33.

You can also start the first management session while the GS950 PS V2 Switch is connected to your network. However, if your network does not have a DHCP server, you still have to change the IP address of your computer to match the subnet of the default address of the GS950 PS V2 Switch. Furthermore, if your network contains virtual LANs (VLANs), you have to be sure to connect the switch and your computer to ports on an Ethernet switch that are members of the same VLAN. Refer to “Starting the First Management Session without a DHCP Server” on page 34.

The instructions for starting the first management session are found in the following sections:

- ❑ “Starting the First Management Session with a DHCP Server” on page 33
- ❑ “Starting the First Management Session with a Direct Connection” on page 33
- ❑ “Starting the First Management Session without a DHCP Server” on page 34

Starting the First Management Session with a DHCP Server

This procedure explains how to start the first management session on the GS950/28PS V2 Switch when the LAN port is connected to a network that has a DHCP server. To start the management session, perform the following procedure:

1. Add the GS950/28PS V2 Switch to the DHCP server. The MAC address for the switch can be found on a label on the back panel. The switch supports both IPv4 and IPv6 addresses.
2. Connect at least one port on the GS950/28PS V2 Switch to a network device, such as another Ethernet switch.
3. Power on the GS950/28PS V2 Switch.
4. Start the web browser on your computer.
5. Enter the IP address (192.168.1.1) of the switch in the URL field of the browser, and change the IP setting to DHCP client mode. This is the IP address assigned to the unit by the DHCP server. If you do not know the address, refer to the DHCP server.

You should see the login window in Figure 5.

The screenshot shows a web browser window titled "Login". It contains two input fields: "User Name" and "Password". Below the input fields are two buttons: "Sign in" and "Clear".

Figure 5. Login Window

6. Enter "manager" for the username and "friend" for the password. The username and password are case-sensitive.
7. Click the Sign In button.

Starting the First Management Session with a Direct Connection

To start the management session with a direct Ethernet connection between your computer and the GS950/28PS V2 Switch, perform the following procedure:

1. Connect one end of a network cable to any port on the switch and the other end to the Ethernet network port on your computer.
2. Change the IPv4 address on your computer to 192.168.1.*n*, where *n* is any number from 2 to 254. Refer to the documentation that accompanies your computer for instructions on how to set the IP address. See Figure 6 on page 34.

IPv4 Setup

System MAC Address: 00:34:5A:7B:69:30

System IP Address: 192 . 168 . 1 . 1

System Subnet Mask: 255 . 255 . 255 . 0

System Default Gateway: 0 . 0 . 0 . 0

System IP Mode: Static

Apply

Figure 6. IPv4 Setup

Note

The switch does not have a default IPv6 address.

3. Set the subnet mask on your computer to 255.255.255.0.
4. Power on the switch.
5. Start the web browser on your computer.
6. Enter the IP address 192.168.1.1 in the URL field of the browser.

You should see the logon window. Refer to Figure 5 on page 33.

7. Enter “manager” for the username and “friend” for the password. The username and password are case-sensitive.
8. Click the Sign In button.

Starting the First Management Session without a DHCP Server

This procedure explains how to start the first management session when the GS950 PS V2 Switch is connected to a network that does not have a DHCP server. To start the management session, perform the following procedure:

1. If your network has VLANs, verify that your computer and the GS950 PS V2 Switch are connected to ports on an Ethernet switch that are members of the same VLAN. This might require accessing the switch’s management software and listing the VLANs and their port assignments. For example, if the GS950 PS V2 Switch is connected to a port that is a member of the Sales VLAN, your computer must be connected to a port that is also a member of the same VLAN. If your network is small and does not have VLANs or routers, you may connect your computer to any port on the Ethernet switch.

2. Change the IPv4 address on your computer to 192.168.1.*n*, where *n* is any number from 2 to 254. Refer to the documentation that accompanies your computer for instructions on how to set the IP address. The switch does not have a default IPv6 address. See Figure 6 on page 34.
3. Set the subnet mask on your computer to 255.255.255.0.
4. Power on the GS950 PS V2 Switch.
5. Start the web browser on your computer.
6. Enter the IP address 192.168.1.1 in the URL field of the browser.
You should see the logon window, shown in Figure 5 on page 33.
7. Enter “manager” for the username and “friend” for the password. The username and password are case-sensitive.
8. Click the Sign In button.

Starting a Web Browser Management Session

This section contains the procedures for starting or ending a web browser management session on the switch.

Starting a Management Session

To start a web browser management session with the switch, perform the following procedure:

Note

If you are using the switch's default IPv4 address 192.168.1.1, start with step 1. If you have already assigned the switch a new address, start with step 3.

1. Change the IPv4 address of your computer to 192.168.1.*n*, where *n* is a number from 2 to 254. See Figure 6 on page 34.
2. Connect the Ethernet network port on your computer to any of the Ethernet ports on the switch.

Note

The Console port is not supported.

3. Start the web browser on your computer and enter the IP address of the switch in the URL field.

The default address is 192.168.1.1 with the subnet mask 255.255.255.0.

The switch displays the logon window, shown in Figure 5 on page 33.

4. Enter the username and password for the switch. The default settings are "manager" and "friend", respectively. The username and password are case sensitive. (The password appears in the Password field as a series of asterisks.)

The switch displays the Device Monitoring - System Information window, shown in Figure 7 on page 37.

Note

To end a web browser management session, close your web browser.

Main Menu

Figure 7 shows the main menu of the web browser management interface. The menu is displayed on the left side of the windows.

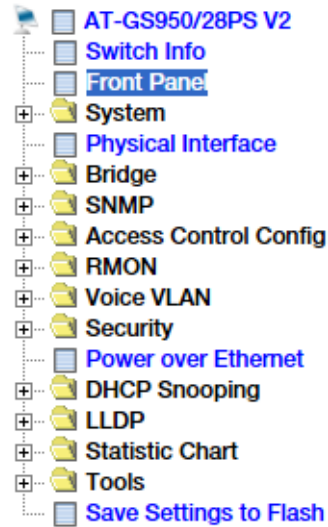


Figure 7. Main Menu

Some menu selections have sub-menus that you can display by clicking the selections. Figure 8 shows the System sub-menu.

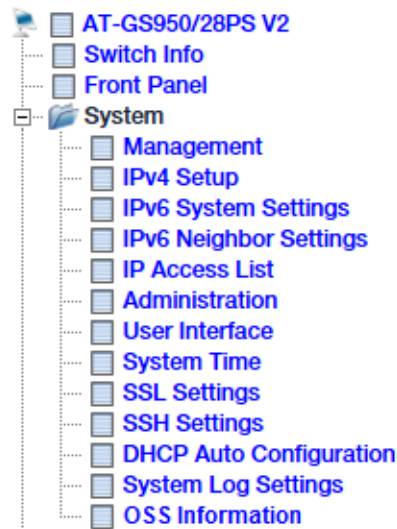


Figure 8. Example of a Sub-menu

Displaying General Switch Information

The first window you will see after starting a management session is the Switch Information window shown in Figure 10. You can also display the window by selecting **Switch Info** from the main menu, as shown in Figure 9.




Figure 9. Switch Info Menu Selection

Switch Information	
System Up For:	9 day(s),5 hr(s),11 min(s),39 sec(s)
Runtime Image:	1.00.005
Boot Loader:	1.00.000
Hardware Information	
DRAM Size:	256 MB
Flash Size:	32 MB
Administration Information	
System Name:	
System Location:	
System Contact:	
System MAC Address, IPv4 Information	
MAC Address:	E0:1A:EA:56:3F:54
IP Address:	192.168.1.1
Subnet Mask:	255.255.255.0
Default Gateway:	0.0.0.0
IPv6 Information	
IPv6 Unicast Address / Prefix Length:	
IPv6 Default Gateway:	
Link Local Address / Prefix length:	
Automatic Network Features	
IPv4 DHCP Client Mode:	Disabled
IPv4 BOOTP Client Mode:	Disabled
IPv6 DHCP Client Mode:	Disabled

Figure 10. Switch Information Window

The fields are described in Table 1.

Table 1. Switch Information Window

Field	Description
Switch Information	
System Up For	Displays the number of days, hours, minutes, and seconds the switch has been running since it was last rebooted.
Runtime Image	Displays the version number of the runtime firmware.
Boot Loader	Displays the version number of the bootloader firmware.
Hardware Information	
DRAM Size	Displays the size of the DRAM, in megabytes.
Flash Size	Displays the size of the flash memory, in megabytes.
Administration Information	
System Name	Displays the name assigned to the switch. Refer to “Setting the Name, Location, and Contact” on page 56.
System Location	Displays the location of the switch. Refer to “Setting the Name, Location, and Contact” on page 56.
System Contact	Displays the contact person responsible for managing the switch. Refer to “Setting the Name, Location, and Contact” on page 56.
System MAC Address, IPv4 Information	
MAC Address	Displays the MAC address of the switch.
IP Address	Displays the switch’s IPv4 address. Refer to Chapter 3, “IPv4 Address” on page 59.
Subnet Mask	Displays the subnet mask of the switch.
Default Gateway	Displays the default gateway IP address.

Table 1. Switch Information Window (Continued)

Field	Description
IPv6 Information	
IPv6 Unicast Address/Prefix Length	Displays the switch's IPv6 address and prefix length. Refer to Chapter 4, "IPv6 Address" on page 65.
IPv6 Default Gateway	Displays the IPv6 default gateway address.
Link Local Address/Prefix Length	Displays the IPv6 link local address.
Automatic Network Features	
IPv4 DHCP Client Mode	Displays the status of the IPv4 DHCP client on the switch. Refer to "Setting the IPv4 Address from a DHCP or BOOTP Server" on page 64.
IPv4 BOOTP Client Mode	Displays the status of the BOOTP client on the switch. Refer to "Setting the IPv4 Address from a DHCP or BOOTP Server" on page 64.
IPv6 DHCP Client Mode	Displays the status of the IPv6 DHCP client. Refer to "Setting the IPv6 Address" on page 67.

Saving Your Changes to a Configuration File

Your changes to the settings of the features are immediately implemented by the switch as soon you enter them in the web browser interface. However, the changes are not permanently saved and will be lost if you reboot or power cycle the switch without first saving them in one of its two configuration files in flash memory.

The two configuration files are labeled Config1 and Config2. Config1 is the primary configuration file and Config2 is the secondary file. Unless you specify otherwise, the switch always uses Config1 as the active primary configuration file for saving your changes and reapplying its settings when rebooted or powered on.

You can use the secondary Config2 file for several functions. For instance, you can use it as a backup to the Config1 file so that if the latter becomes damaged or corrupted, you can quickly restore the switch settings. You can also use it to return the switch to an earlier configuration. For example, if you make changes to the switch's configuration and save them to Config1 but not Config2, you could restore the previous configuration by instructing the switch boot up with the Config2 file. For instructions, refer to "Booting the Switch with the Config2 Configuration File" on page 106.

Perform the following procedure to save the switch's configuration to the Config1 or Config2 configuration file:

1. Select **Save Setting to Flash** at the bottom of the main menu. Refer to Figure 11.




Figure 11. Save Settings to Flash Menu Selection

The switch displays the Save Settings to Flash window, shown in Figure 12.



Note: The switch will stop responding while saving the current configuration to flash.

Figure 12. Save Settings to Flash Window

2. To save your changes to the Config1 file, click the **Save Settings to Flash** button. You do not need to change the window to save to the Config1 file.
3. To save your changes to the Config2 file, do the following:
 - a. Click the **Startup-Config** option to remove the check mark from its dialog box.
 - b. Select **Config2** from the menu.
 - c. Click the **Save Settings to Flash** button. The switch's configuration is save in the Config2 file.

Unsupported Special Characters

Text fields in the management interface do not support the following special characters:

- Vertical bar (|)
- Comma (,)
- Forward slash and backslash (\ /)
- Angle brackets (< >)
- Question mark (?)
- Colon (:)
- Semicolon (;)
- Single quote (')
- Double quotes (“ ”)

Chapter 2

Port LEDs and Basic Port and Switch Settings

This chapter contains the following sections:

- “Viewing the Front Panel of the Switch” on page 46
- “Configuring Basic Port Settings” on page 51
- “Setting the Name, Location, and Contact” on page 56

Viewing the Front Panel of the Switch

To view the front panel and port LEDs of the switch, select **Front Panel** from the main menu. Refer to Figure 13.

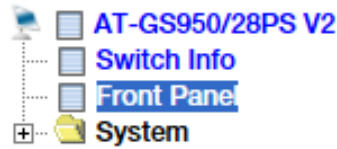


Figure 13. Front Panel Menu Selection

The example in Figure 14 is of the GS950/28PS V2 Switch.



Figure 14. Front Panel Window

SYSTEM LED

The GS950/18PS V2, GS950/28PS V2, and GS950/52PS V2 Switches have a SYS LED on the left side of the faceplate. Refer to Figure 15.




Figure 15. PoE MAX and SYS LEDs

Table 2 defines the SYS LED states.

Table 2. SYS LED

State	Description
Off	The switch is not receiving AC power.
Green	The switch is operating normally.

Table 2. SYS LED (Continued)

State	Description
Red	A cooling fan has failed.

POWER LED

The GS950/10PS V2 Switch has a POWER LED on the left side of the faceplate. Refer to Figure 16.



Figure 16. POWER LED

Table 3 defines the LED states.

Table 3. POWER LED

State	Description
Off	The switch is not receiving AC power or has experienced a system failure.
Green	The switch is operating normally.

PoE MAX LED


The switches have a PoE MAX LED on the left side of the faceplate. Refer to Figure 15 on page 46. Table 4 defines the states of the PoE MAX LED.

Table 4. PoE MAX LED

State	Description
Off	The total power requirements of the powered devices connected to the ports are below the switch's maximum power budget. The switch is providing power to all powered devices on its ports.
Red	The total power requirements of the powered devices exceed Guard Band range (switch's maximum power budget - Guard Band).

Copper Port LEDs

The port LEDs for the GS950/10PS V2, GS950/18PS V2, and GS950/28PS V2 Switches are identified in Figure 17.



1. PoE LEDs
2. Link/Activity LEDs

Figure 17. Copper Port LEDs

The LEDs for PoE copper ports 1 to 24 on the GS950/52PS V2 Switch are identified in Figure 18.




Figure 18. PoE Copper Ports 1 to 24 LEDs on the GS950/52PS V2 Switch

The LEDs for the non-PoE copper ports 25 to 48 on the GS950/52PS V2 Switch are identified in Figure 19.




Figure 19. Non-PoE Copper Ports 25 to 48 LEDs on the GS950/52PS V2 Switch

Table 5 defines the states of the link/ activity (L/A) LEDs.

Table 5. Link and Activity (L/A) LEDs for Copper Ports

State	Description
Off	The port has not established a link to a network device.
Steady Green	The port has established a 1000Mbps link to a network device.
Steady Amber	The port has established a 10 or 100Mbps link to a network device

Table 6 defines the states of the PoE LEDs.

Table 6. PoE LEDs for the Copper Ports

State	Description
Off	This state has the following possible causes: - The port is not connected to a powered device.
Green	The port is supplying power to a powered device.
Amber	This state has the following possible causes: - A powered device is requiring more power than its device class. - There is a terminal short in the network cable or connector. - The switch is denying power to a powered device because it has reached its maximum power budget.

SFP Port LEDs

Each SFP port has one LED that indicates link status. Refer to Table 7.

Table 7. Link LEDs for the SFP Ports

State	Description
Off	The port has not established a link to a network device.
Steady Green	The port has established a 100Mbps or 1000Mbps link to a network device.

Spanning Tree LEDs

The LEDs inside the ports display spanning tree status. Refer to Table 8.

Table 8. Spanning Tree Protocol Port LEDs

State	Description
Off	The port has not established a link to a network device or STP is disabled on the port.
Steady Red	The port is in the spanning tree discarding state.
Steady Green	The port is in the spanning tree forwarding state or spanning tree is disabled on the switch.
Steady Yellow	The port is in the spanning tree learning state.

Configuring Basic Port Settings

This procedure explains how to configure the following port settings:

- Enable or disable ports (referred to as Admin status)
- Set the speed and duplex mode
- Forward or block jumbo frames
- Enable or disable flow control
- Forward or block Extensible Authentication Protocol (EAP) frames
- Forward or block spanning tree protocol BPDU frames

The procedure also displays the following parameters:

- Trunk group number
- Port type
- Link status

Perform the following procedure to view or configure the basic parameter settings of the switch ports:

1. Select **Physical Interface** from the main menu. Refer to Figure 20.




Figure 20. Physical Interface Menu Selection

The Physical Interface window is shown in Figure 21.

Physical Interface										
Port	Trunk	Type	Link Status	Admin Status	Mode	Jumbo	Flow Ctrl	EAP	BPDU	Action
All	-	-	-	Ignore ▾	Ignore ▾	Ignore ▾	Ignore ▾	Ignore ▾	Ignore ▾	Apply
1	---	1000TX	Up	Enabled ▾	Auto (1000F) ▾	Enabled ▾	Disabled ▾	Disabled ▾	Enabled ▾	Apply
2	---	1000TX	Up	Enabled ▾	Auto (1000F) ▾	Enabled ▾	Disabled ▾	Disabled ▾	Enabled ▾	Apply
3	---	1000TX	Up	Enabled ▾	Auto (1000F) ▾	Enabled ▾	Disabled ▾	Disabled ▾	Enabled ▾	Apply
4	---	1000TX	Up	Enabled ▾	Auto (1000F) ▾	Enabled ▾	Disabled ▾	Disabled ▾	Enabled ▾	Apply
5	---	1000TX	Up	Enabled ▾	Auto (1000F) ▾	Enabled ▾	Disabled ▾	Disabled ▾	Enabled ▾	Apply
6	---	1000TX	Up	Enabled ▾	Auto (1000F) ▾	Enabled ▾	Disabled ▾	Disabled ▾	Enabled ▾	Apply
7	---	1000TX	Up	Enabled ▾	Auto (1000F) ▾	Enabled ▾	Disabled ▾	Disabled ▾	Enabled ▾	Apply

Figure 21. Physical Interface Window

For a quick way to set all ports to the same setting, use the All row at the top of the table. Parameters set in the All row are applied to all

ports. The Ignore setting means the All row is ignored for that parameter.

2. Configure the port settings in Table 9.

Table 9. Physical Interface (Port) Settings

Parameter	Definition
Port	Displays the port number.
Trunk	Displays the static or LACP trunk number. This value is empty for ports that are not members of trunks. Refer to Chapter 29, “Static Port Trunks” on page 363 or Chapter 30, “LACP Trunks” on page 371.
Type	Displays the port type: <ul style="list-style-type: none"> - 100TX: 10/100/1000Mbps for copper ports - 100FX or 1000TX: 100/1000Mbps for SFP ports
Link Status	Displays the link state: <ul style="list-style-type: none"> - Up -The port has a valid link to a network device. - Down -The port does not have a link to a network device. Examples of the Down state are listed here: <ul style="list-style-type: none"> - The port is not connected to a network device. - The port is disabled. - The network device is not powered on. - There is a problem with the copper or fiber optic cable.

Table 9. Physical Interface (Port) Settings (Continued)

Parameter	Definition
Admin Status	<p>Enables or disables ports. The menu has these options:</p> <ul style="list-style-type: none"> - Enabled: Enables the port. The switch forwards network traffic on the port. This is the default setting. - Disabled: Disables the port. The switch blocks all ingress and egress traffic on the port. <p>Reasons for disabling ports include:</p> <ul style="list-style-type: none"> - Secure unused ports to protect the network from unauthorized connections. - Ports are under network attack. - The copper or fiber optic cable has a problem. - The network device is not operating correctly.
Mode	<p>Sets the speed and duplex mode of the port. The menu has these options:</p> <ul style="list-style-type: none"> - Auto -Sets the port to Auto-Negotiation. Once a port establishes a link to a network device, the operating speed and duplex mode are displayed in parentheses (for example, "1000F" for 1000Mbps full duplex mode). This is the default setting. - 1000/Full: Sets the port to 1000Mbps, full-duplex mode. - 100/Full: Sets the port to 100Mbps, full-duplex mode. - 10/Full: Sets the port to 10Mbps, full-duplex mode. - 100/Half: Sets the port to 100Mbps, half-duplex mode. - 10/Half: Sets the port to 10Mbps, half-duplex mode.

Table 9. Physical Interface (Port) Settings (Continued)

Parameter	Definition
Mode (Continued)	<p>Here are guidelines to setting speed and duplex mode:</p> <ul style="list-style-type: none"> - Fiber port support force mode setting, default is Auto mode. - Do not use the Auto setting for copper ports that are connected to remote network devices that do not support Auto-Negotiation. Instead, manually set the speed and duplex mode on copper port to match the remote devices.
Jumbo	<p>Enables or disables support for jumbo frames of up to 10KB on the port. The menu has the following selections:</p> <ul style="list-style-type: none"> - Enabled: Enables jumble frame support. This is the default setting. - Disabled: Disables jumble frame support. The port discards ingress and egress jumbo frames. <p>You cannot enable jumbo frame support on ports where CoS is also enabled. Refer to “Mapping CoS Priorities to Egress Queues” on page 304.</p>
Flow Control	<p>Enables or disables flow control. Ports use flow control to temporarily stop their remote counterparts from sending packets, to allow them time to process packets already stored in their buffers. Ports initiate flow control by sending pause packets. The switch can also respond to pause packets from other devices by temporarily delaying the transmission of network packets on ports. The pull-down menu has these options:</p> <ul style="list-style-type: none"> - Enabled: Enables flow control. The port transmits pause packets and responds to pause packets from its network counterpart. - Disabled: Disables flow control. This is the default setting.

Table 9. Physical Interface (Port) Settings (Continued)

Parameter	Definition
EAP	<p>Enables or disables forwarding of Extensible Authentication Protocol (EAP) packets. The menu has these options:</p> <ul style="list-style-type: none"> - Enabled: The port forwards ingress and egress EAP packets. - Disabled: The port discards ingress and egress EAP packets. This is the default setting,
BPDU Pass-through	<p>Enables or disables forwarding of spanning tree BPDU packets by the switch from other network devices. The menu has these options:</p> <ul style="list-style-type: none"> - Enabled: The port forwards ingress spanning tree BPDU packets from other network devices. This is the default setting. - Disabled: The port does not pass-through ingress spanning tree BPDU packets. You have to disable BPDU pass-through on all ports to enable spanning tree on the switch.

3. Select **Save Settings to Flash** from the main menu to save your changes.

Setting the Name, Location, and Contact

This procedure explains how to configure the following values:

- System name
- Location
- Contact

Assigning a name, location and contact administrator will make identifying the switch easier and so reduce the chances of a network administrator configuring the wrong device. The procedure also lets you view these values:

- System description (model name)
- System object ID for SNMP

Perform the following procedure to view or set the values:

1. Select **System** -> **Management** from the main menu. Refer to Figure 13.




Figure 22. Management Menu Selection

The Management window is shown in Figure 14.

Management

System Description: AT-GS950/28PS V2 28-Port Gigabit Web Smart Switch with 4 SFP Slots

System Object ID: 1.3.6.1.4.1.207.1.4.327

System Name:

System Location:

System Contact:

Figure 23. Management Window

2. Configure the fields in Table 6.

Table 10. Management Window

Field	Description
System Description	Displays the Allied Telesis switch model. You cannot change this value.
System Object ID	Displays the SNMP MIB object identifier of the switch model. You cannot change this value.
System Name	Enter a name of up to 15 characters for the switch, for example Sales. The name is optional. Spaces and most special characters are allowed. Refer to “Unsupported Special Characters” on page 43:
System Location	Enter the physical location of up to 15 characters for the switch. The location is optional. Spaces and most special characters are allowed.
System Contact	Enter the name of the network administrator responsible for maintaining the switch. The system contact can have up to 15 characters. The system contact is optional. Spaces and most special characters are allowed.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 3

IPv4 Address

This chapter contains the following sections:

- “IPv4 Address Overview” on page 60
- “Manually Setting the IPv4 Address” on page 61
- “Setting the IPv4 Address from a DHCP or BOOTP Server” on page 64

IPv4 Address Overview

The sections in this chapter explain how to assign the switch an IP4 address. You can assign the address manually or have the switch automatically obtain the address from a DHCP or BOOTP server on your network. The switch has to have an IP address to support the following features:

- Web browser management
- Command line management with a Telnet or Secure Shell client
- Syslog server
- RADIUS authentication
- TACACS+ authentication
- RMON
- SNMPv1, v2c, or v3
- HTTP software updates
- TFTP software updates
- Ping tests

Manually Setting the IPv4 Address

This procedure explains how to manually assign an IPv4 address configuration to the switch. The configuration consists of the following values:

- IPv4 address
- Subnet mask
- Default gateway

The switch comes with the default static IP address 192.168.1.1 and subnet mask 255.255.255.0. The switch can have only one IPv4 address.

Note

Changing the switch's IPv4 address will interrupt your management session. To resume managing the device, you will have to start a new management session using its new IPv4 address.

This procedure also displays the MAC address of the switch.

Perform the following procedure to manually assign an IPv4 address to the switch:

1. Select **System** -> **IPv4 Setup** from the main menu. Refer to Figure 24.




Figure 24. IPv4 Setup Menu Selection

The IPv4 Setup window is shown in Figure 25 on page 62.

IPv4 Setup

System MAC Address: E0:1A:EA:56:3F:54

System IP Address: . . .

System Subnet Mask: . . .

System Default Gateway: . . .

System IP Mode: ▼

Figure 25. IPv4 Setup Window

Note

System IP Mode has to be set to Static, the default setting, for you to assign a static IPv4 address to the switch. If it is set to DHCP or BOOTP, change it to Static.

2. Configure the parameters in Table 11.

Table 11. IPv4 Setup Window

Field	Description
System MAC Address	Displays the MAC address of the switch. You cannot change this parameter.
System IP Address	Enter a new IPv4 address for the switch. The switch can have only one IPv4 address. The default address is 192.168.1.1.
System Subnet Mask	Enter new subnet mask. The default mask is 255.255.255.0.
System Default Gateway	Enter the IPv4 address of the default gateway. This is the address of an interface on a router or other Layer 3 routing device that represents the first hop to reaching the subnets or networks where management devices, such as management workstations or RADIUS servers, are located.

Table 11. IPv4 Setup Window (Continued)

Field	Description
System IP Mode	Specify whether the IPv4 address is assigned manually "Static" or by a DHCP or BOOTP server. The options are listed here: <ul style="list-style-type: none"> - Static: The IPv4 address is added manually. This is the default setting. - DHCP: The IPv4 address is provided by a DHCP server on the network. - BOOTP: The IPv4 address is provided by a BOOTP server on the network

3. Click the **Apply** button.

The system displays a confirmation prompt.

4. Click **OK** to implement the change or **Cancel** to keep the previous values.

Note

If you click OK, your management session is interrupted. To resume managing the switch, start a new management session on the device using its new IPv4 address.

5. After starting the new session, select **Save Settings to Flash** from the main menu to save your changes.

Setting the IPv4 Address from a DHCP or BOOTP Server

The switch has DHCP and BOOTP clients that allow it to automatically obtain its IPv4 address configuration from a DHCP or BOOTP server on your network. Here are the guidelines:

- The switch can have only one IPv4 address.
- Only one client can be active at a time.
- The switch transmits the client queries from the ports of the management VLAN only. The default VLAN includes all ports and has the VID 1.
- The switch operates without an IP address if it does not receive a response from a DHCP or BOOTP server.
- If the switch's current IP address is from a DHCP or BOOTP server, it retains the address if you disable the client, until the lease time on the address expires.

Perform this procedure to enable or disable the DHCP or BOOTP client:

1. Select **System** -> **IPv4 Setup** from the main menu. Refer to Figure 24 on page 61. The IPv4 Setup window is shown in Figure 25 on page 62.
2. In the System IP Mode field, select one of the following:
 - DHCP**: Activates the DHCP client.
 - BOOTP**: Activates the BOOTP client.
 - Static**: Deactivates the DHCP and BOOTP clients so that you can enter a static address. This is the default setting. Refer to "Manually Setting the IPv4 Address" on page 61.
3. Click the **Apply** button.

The system displays a confirmation prompt.

4. Click **OK** to implement the change or **Cancel** to keep the previous values.

Note

If you click OK, your management session is interrupted. To resume managing the switch, start a new management session using its new IPv4 address.

5. After starting the new session, select **Save Settings to Flash** from the main menu to save your changes.

Chapter 4

IPv6 Address

This chapter contains the following sections:

- “IPv6 Address Overview” on page 66
- “Setting the IPv6 Address” on page 67
- “Setting the IPv6 Neighbor Settings” on page 71

IPv6 Address Overview

The sections in this chapter explain how to assign the switch an IP6 address. You can assign the address manually or have the switch automatically obtain it from a DHCP IPv6 server on your network. The switch needs an IP address to support the following features:

- Web browser management
- Command line management with a Telnet or Secure Shell client
- Syslog server
- RADIUS authentication
- TACACS+ authentication
- RMON
- SNMPv1, v2c, or v3
- HTTP software updates
- TFTP software updates
- Ping tests

Setting the IPv6 Address

This procedure explains how to assign an IPv6 address configuration to the switch or to activate the DHCP IPv6 client to assign the switch an address from a DHCP IPv6 server.

Note

Changing the switch's IPv6 address will interrupt your management session if you are using the address to manage the switch. To resume managing the unit, start a new management session using the new IPv6 address.

Perform the following procedure to set the IPv6 address:

1. Select **System > IPv6 System Settings** from the main menu. Refer to Figure 26.

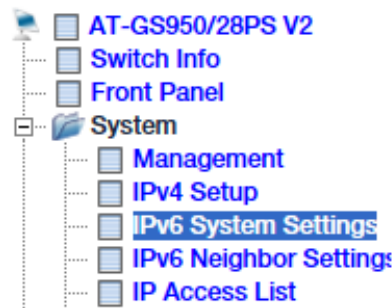


Figure 26. IPv6 System Settings Menu Selection

The IPv6 System Settings window is shown in Figure 27 on page 68.

IPv6 System Settings

IPv6 State

DHCPv6 Client

IPv6 Unicast Address/Prefix Length (e.g.:3710::1/64)

IPv6 Static Gateway (e.g.:3710::9)

IPv6 Dynamic Gateway

NS Retransmit Time Settings

NS Retransmit Time (1-3600) sec

Link Local Address Settings

Automatic Link Local Address

Link Local Address/Prefix length (e.g.:FE80::6/10)

Figure 27. IPv6 System Settings Window

2. Configure the settings in Table 12.

Table 12. IPv6 System Settings Window

Field	Description
IPv6 System Settings	
IPv6 State	Select one of the following: <ul style="list-style-type: none"> - Enabled: Enables the IPv6 address configuration on the switch. The default is disabled. You have to enable the IPv6 State to configure the window settings. - Disabled: Disables the IPv6 address configuration and window settings.

Table 12. IPv6 System Settings Window (Continued)

Field	Description
DHCPv6 Client	<p>Select one of the following:</p> <ul style="list-style-type: none"> - Enabled: Enables the DHCP IPv6 client. The switch obtains its IPv6 address configuration from a DHCP IPv6 server on your network. Enabling the client disables the settings in the window. - Disabled: Disables the DHCP IPv6 client. This is the default setting. This setting lets you manually assign an IP address.
IPv6 Unicast Address/Prefix Length	<p>Enter the IPv6 unicast address and prefix length for the switch. The default is no address. The address is entered in this format:</p> <p>xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/prefix</p> <p>Each x is a hexadecimal digit representing 4 bits. The eight groups of digits have to be separated by colons. Groups where all four digits are '0' can be omitted. Leading '0's in groups can also be omitted.</p> <p>As an example, the following IPv6 addresses are equivalent:</p> <p>3710:421e:09a8:0000:0000:0000:00a4:1c50/56 3710:421e:9a8::a4:1c50/56</p> <p>If the DHCP IPv6 client is enabled on the switch, this field displays the IPv6 address assigned by the DHCP server.</p>
IPv6 Static Gateway	Enter the IPv6 static gateway of the switch.
IPv6 Dynamic Gateway	Displays the IPv6 address of the dynamic gateway on the switch from the DHCP IPv6 server. This parameter is empty if the DHCP IPv6 client is disabled.
NS Retransmit Time Settings	
NS Retransmit Time	Enter a Neighbor Solicitation (NS) retransmit time. The range is from 1 to 3600 seconds. The default is 1 second.

Table 12. IPv6 System Settings Window (Continued)

Field	Description
Link Local Address Settings	
Automatic Link Local Address	Select one of the following: Enabled: Automatically assigns a link local address (follow EUI-64 format). Disabled: Manually assign the link local address. This is the default setting.
Link Local Address/Prefix Length	Enter the link local address and prefix length. Automatic Link Local Address has to be disabled for you to set this value.

3. Click the **Apply** buttons.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Setting the IPv6 Neighbor Settings

Perform the following procedure to manage the list of IPv6 neighbors:

1. Select **System** -> **IPv6 Neighbor Settings** from the main menu. Refer to Figure 28.

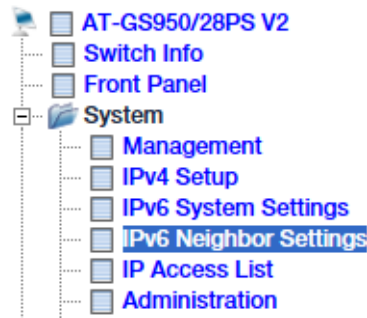


Figure 28. IPv6 Neighbor Settings Menu Selection

The IPv6 Neighbor Settings window is shown in Figure 29.

IPv6 Neighbor Settings

Neighbor IPv6 Address *

Link Layer MAC Address : : : : : * (XX:XX:XX:XX:XX:XX)

Neighbor IPv6 Address	Link Layer MAC Address	Status	Action
*	*	All <input type="button" value="v"/>	<input type="button" value="Find"/> <input type="button" value="Delete"/>

Page 0 / 0 Page

Figure 29. IPv6 Neighbor Settings Window

2. To add a new address, do the following:
 - a. Enter an IPv6 address in the **Neighbor IPv6 Address** field in this format:

XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX

Each x is a hexadecimal digit representing 4 bits. The eight groups of digits have to be separated by colons. Groups where all four digits are '0' can be omitted. Leading '0's in groups can also be omitted. For example, the following IPv6 addresses are equivalent:

3710:421e:09a8:0000:0000:0000:00a4:1c50

3710:421e:9a8::a4:1c50

- b. Enter a link layer MAC address in the **Link Layer MAC Address** field.
 - c. Click the **Add** button.
 3. To delete addresses, do one of the following:
 - To delete a single entry, click its **Delete** button in the Action column.
 - To delete all the entries, select **All** from the drop-down menu in the Status column and then click the **Delete** button in the Action column.
 - To delete all static entries, select **Static** from the drop-down menu in the Status column and then click the **Delete** button in the Action column.
 - To delete all dynamic entries, select **Dynamic** from the drop-down menu in the Status column and then click the **Delete** button in the Action column.
 4. To find an address in the table, do one of the following:
 - To view all static or dynamic IPv6 neighbors, type asterisks in the **Neighbor IPv6 Address** and **Link Layer MAC Address** fields, then select **Static** or **Dynamic** from the drop-down menu in the Status column.
 - To find a neighbor by its IPv6 address, enter the neighbor address in the **Neighbor IPv6 Address** field and an asterisk in the **Link Layer MAC Address** field. The asterisk serves as a wildcard character.
 - To find a specific IPv6 neighbor by its MAC address, enter an asterisk in the **Neighbor IPv6 Address** field and the MAC address in the **Link Layer MAC Address** field.
 - To find an address by both its IPv6 address and MAC address, enter the values in the **Neighbor IPv6 Address** field and **Link Layer MAC Address** fields.
 5. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 5

DHCP Auto Configuration

This chapter contains the following section:

- “Enabling or Disabling DHCP Auto Configuration” on page 74

Enabling or Disabling DHCP Auto Configuration

The switch stores its parameter settings in a configuration file in flash memory. It updates the file with the latest settings whenever you select the Save Settings to Flash main menu selection. It is good network practice to maintain backup copies of the configuration files on a server on your network, so that you can restore the files to the devices or to configure replacement devices, if needed. Refer to “Backing Up Configuration Files from the Switch with HTTP” on page 99 or “Backing Up Configuration Files from the Switch with TFTP” on page 102.

The switch has a DHCP auto configuration feature that can make adding or replacing switches easier. With DHCP auto configuration, the switch can automatically obtain its configuration file from a DHCP server along with its IP address when booted or powered on. Here are the guidelines:

- ❑ You have to enable the DHCP client on the switch. Refer to “Setting the IPv4 Address from a DHCP or BOOTP Server” on page 64 or “Setting the IPv6 Address” on page 67.
- ❑ The DHCP server has to support option 54 (server address). Use the option to specify the IP address of a TFTP server.
- ❑ The DHCP server also has to support option 67 (filename). Use this option to specify the filename of the configuration file for the switch.
- ❑ The TFTP and DHCP servers have to reside on the same device.
- ❑ The switch configuration files have to reside on the DHCP server.

Perform this procedure to enable or disable DHCP auto configuration:

1. Select **System -> DHCP Auto Config** from the main menu. Refer to Figure 30.

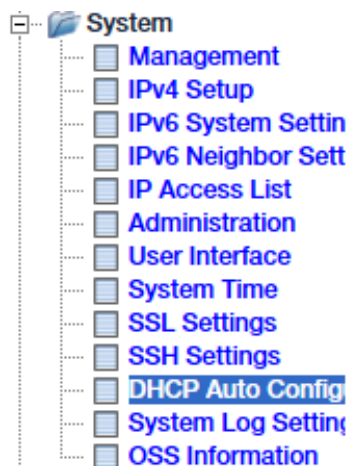


Figure 30. DHCP Auto Config Menu Selection

The DHCP Auto Config window is shown in Figure 31.




Figure 31. DHCP Auto Configuration Window

2. Do one of the following:
 - To enable the feature, select **Enabled**. The switch will query the DHCP server for both its IP address assignment and configuration file the next time it is booted or powered on.
 - To disable the feature, select **Disabled**. This is the default setting.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.
5. If you enabled the feature and want the switch to obtain its configuration file from the DHCP server now, reboot the switch. Refer to “Rebooting the Switch” on page 500.
6. After the switch obtains its configuration file from the DHCP server, you should disable the DHCP Auto Configuration feature. That way, any additional configuration changes you make to the switch will not be overwritten by the configuration file on the DHCP server the next time you reboot or power cycle the device.

Chapter 6

System Time

This chapter contains the following section:

- “Setting the Date and Time” on page 78

Setting the Date and Time

The switch has an internal clock and calendar. The switch uses the clock and calendar to add the time and date to the events it stores in the event log and sends to syslog servers, You can manually set the clock and calendar yourself or have the switch automatically obtain the time and date from a Network Time Protocol (NTP) server on your network or the Internet.

Note

The switch maintains the time and date with its internal battery when powered off.

Here are the procedures:

- ❑ “Manually Setting the Date and Time” on page 78
- ❑ “Setting the Date and Time from an SNTP Server” on page 80

Manually Setting the Date and Time

Perform the following procedure to manually set the date and time:

1. Select **System** -> **System Time** from the main menu. Refer to Figure 32.

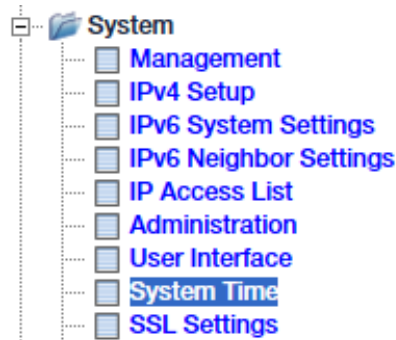


Figure 32. System Time Menu Selection

The System Time window is shown in Figure 33 on page 79.

System Time

Clock Mode: Local Time
 Current Time: 1 Jan 2020 00:37:50
 Time Zone:

Date/Time Settings

Clock Mode:

Local Time Settings

Date Settings: / / (YYYY:MM:DD)
 Time Settings: : : (HH:MM:SS)

Simple Network Time Protocol (SNTP) Settings

SNTP Primary Server: . . . IPv4
 SNTP Secondary Server: . . . IPv4
 SNTP Poll Interval: Min(1-60)
 Time Zone:

Additional Time Parameters

Daylight Saving Time Status:
 From: (Month:Day:HH:MM)
 To: (Month:Day:HH:MM)
 DST Offset:

Figure 33. System Time Window

- Configure the fields in Table 13.

Table 13. Manually Setting the Clock and Calendar

Field	Description
Date/Time Settings	
Clock Mode	Set to Local . This is the default setting.
Local Time Settings	

Table 13. Manually Setting the Clock and Calendar (Continued)

Field	Description
Date Settings	Set the current year, month, and date from the menus.
Time Settings	Set the current hour, minutes, and seconds from the menus. The hour is set in 24-hour format.
Addition Time Parameters Set these fields if the switch’s location observes Daylight Savings Time.	
Daylight Savings Time Status	Choose one of the following: - Enabled: Select this if the switch’s location observes Daylight Savings Time. - Disabled: Select this if the switch’s location does not observe Daylight Savings Time. This is the default setting.
From	Set the month, day, hour, and minute of the start of Daylight Savings Time.
To	Set the month, day, hour, and minute of the end of Daylight Savings Time.
DST Offset	Set the Daylight Savings Time offset from the menu. The choices are 1 hour, the default, and 30 minutes.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Setting the Date and Time from an SNTP Server

Perform the following procedure to configure the switch to obtain the date and time from an SNTP server:

1. Select **System** -> **System Time** from the main menu. Refer to Figure 32 on page 78. The System Time window is shown in Figure 33 on page 79.

2. Configure the fields in Table 14.

Table 14. Setting the Calender and Clock from an NTP Server

Field	Description
Date/Time Settings	
Clock Mode	Set to SNTP .
Simple Network Time Protocol (SNTP) Settings	
SNTP Primary Server	<p>Select IPv4 or IPv6 from the menu and enter the IP address of the primary SNTP server. The default is IPv4. This field is required.</p> <p>The format for an IPv4 address is shown here: nnn nnn nnn nnn Each N is a decimal number from 0 to 255.</p> <p>The format for an IPv6 address is shown here: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx Each x is a hexadecimal digit representing four bits. The eight groups of digits have to be separated by colons. Groups where all four digits are '0' can be omitted. Leading '0's in groups can also be omitted.</p> <p>As an example, the following IPv6 addresses are equivalent: 3710:421e:09a8:0000:0000:0000:00a4:1c50 3710:421e:9a8::a4:1c50</p>
SNTP Secondary Server	Select IPv4 or IPv6 from the menu and enter the IP address of a secondary SNTP server. The default is IPv4. This field is optional.
SNTP Poll Interval	Enter the number of minutes the switch waits to poll the NTP server. The range is 1 to 60 minutes. The default is 1 minute.
Time Zone	Select the time zone of the location of the switch from the pull-down menu.
Addition Time Parameters Set these fields if the location of the switch observes Daylight Savings Time.	

Table 14. Setting the Calendar and Clock from an NTP Server (Continued)

Field	Description
Daylight Savings Time Status	Select one of the following: <ul style="list-style-type: none"> - Enabled: Select this if the switch's location observes Daylight Savings Time. - Disabled: Select this if the switch's location does not observe Daylight Savings Time. This is the default setting.
From	Set the month, day, hour, and minute of the start of Daylight Savings Time.
To	Set the month, day, hour, and minute of the end of Daylight Savings Time.
DST Offset	Set the Daylight Savings Time offset from the menu. The choices are 1 hour, the default, and 30 minutes.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 7

Management Security

This chapter contains the following sections:

- ❑ “Adding or Deleting Manager Accounts” on page 84
- ❑ “Configuring the Management Interfaces” on page 87
- ❑ “Enabling or Disabling HTTP or HTTPS Web Browser Management” on page 89
- ❑ “Enabling or Disabling the Secure Shell Server for the Command Line Interface” on page 91
- ❑ “Identifying Authorized Management Workstations” on page 93

Adding or Deleting Manager Accounts

You need a user name and password of a manager account to log on and manage the switch. The switch has one default manager account, with the user name “manager” and default password “friend”. If more than one manager will be managing the switch, you can add more accounts so that the administrators do not have to share the same account. There are three options for manager accounts:

- ❑ Local: Local manager accounts are stored on the switch itself. The switch authenticates the user names and passwords when managers log on. The switch can store up to 64 local accounts.
- ❑ Local and RADIUS: If your network has a RADIUS server, you can use both the local manager accounts in the switch as well as a RADIUS server to store and authenticate user names and passwords of the manager accounts. When a manager logs on, the switch first checks its own manager accounts to see if there is a match for the user name and password. If there is no match, it sends the user name and password to the RADIUS server for authentication. For instructions on configuring the RADIUS client on the switch, refer to Chapter 35, “RADIUS Client” on page 445.
- ❑ Local and TACACS+: If your network has a TACACS+ server, you can use it in the same way as a RADIUS server to store and authenticate the user names and passwords of manager accounts. When a manager logs on, the switch sends the user name and password to a TACACS+ server if there is no match in its own local manager accounts. For instructions on configuring the TACACS+ client on the switch, refer to Chapter 36, “TACACS+ Client” on page 451.

Perform the following procedure to set the authentication method and add or delete local manager accounts:

1. Select **System** -> **Administration** from the main menu. Refer to Figure 34.



Figure 34. Administration Menu Selection

The Administration window is shown in Figure 35.

Administration

User Authentication Method:

User Name: (Maximum length is 12)

Password: (Maximum length is 12)

Confirm Password:

Index	User Name	Password	Action
1	manager	*****	<input type="button" value="Modify"/>
2	JaneSmith	*****	<input type="button" value="Modify"/> <input type="button" value="Delete"/>
3	BillWilson	*****	<input type="button" value="Modify"/> <input type="button" value="Delete"/>

Figure 35. Administration Window

2. To select the authentication method of the manager accounts, do the following:
 - a. From the User Authentication Method, select **Local**, **Local & RADIUS**, or **Local & TACACS+**. You can select only one method. The default is Local.
 - b. Click the **Apply** button.
3. To add a new local manager account, do the following:
 - a. Enter a user name in the **User Name** field. Here are the guidelines.
 - The name is case sensitive.
 - It can be up to 12 characters.
 - Spaces and special characters are not recommended.
 - b. Enter a password for the manager account in the **Password** field. The password has the same guidelines as the user name.
 - c. Confirm the password by entering it again in the **Confirm Password** field.
 - d. Click the **Add** button.
 - e. Go to step 6.

4. To modify the password of a local manager account, do the following:
 - a. Click the **Modify** button of the account to be changed. You can modify only one account at a time.

The switch displays the Modify Administration window. Refer to Figure 36.

Modify Administration

Modify Administration

Entry number:

User Name:

Password:

Confirm Password:

Figure 36. Modify Administration Window

- b. Change the password of the account by referring to step 3.

You cannot change the entry numbers or user names of accounts.
 - c. Go to step 6.
5. To delete a local account, do one of the following:
 - To delete a specific account, click its **Delete** button in the Action column.
 - To delete all the accounts, click the **Delete All** button.

You cannot delete the default manager account. Your management session is not interrupted if you delete the account you are currently using to manage the switch.
6. Select **Save Settings to Flash** from the main menu to save your changes.

Configuring the Management Interfaces

The switch has three management interfaces:

- Web browser
- Command line (Telnet)
- SNMPv1, v2c, v3

This procedure explains how to enable or disable the individual management interfaces. It also explains how to set the idle timeout values, which control how long the switch waits before ending inactive management sessions. You can improve the security of the switch by disabling management interfaces you do not plan to use.

Perform the following procedure to configure the management interfaces:

1. Select **System** -> **User Interface** from the main menu. Refer to Figure 37.




Figure 37. User Interface Menu Selection

The User Interface window is shown in Figure 38.

User Interface

SNMP Agent:	<input type="text" value="Enabled"/>	
Web Server Status:	<input type="text" value="Enabled"/>	
CLI(Telnet) Status:	<input type="text" value="Enabled"/>	
<input type="button" value="Apply"/>		
<hr/>		
Web Idle Timeout:	<input type="text" value="10"/>	(3-60) min
Group Interval:	<input type="text" value="120"/>	(0 or 120-1225, 0 is Disabled) sec
CLI Idle Timeout:	<input type="text" value="10"/>	(3-60) min
<input type="button" value="Apply"/>		

Figure 38. User Interface Window

2. Configure the settings in Table 15.

Table 15. User Interface Window

State	Description
SNMP Agent	Select one of the following: <ul style="list-style-type: none"> - Enabled: Enables the SNMP agent to allow switch management with SNMPv1, v2c, or v3. This is the default setting - Disabled: Disables SNMP management. Refer to Chapter 31, “SNMPv1 and v2c” on page 385 or Chapter 32, “SNMPv3” on page 397.
Web Server Status	Displays the status of the web browser interface. The interface is always enabled. You cannot disable it.
CLI (Telnet) Status	Select one of the following: <ul style="list-style-type: none"> - Enabled: Enables the command line interface so that you can manage the switch with a Telnet or SHH client. This is the default setting. - Disabled: Disables the command line interface.
Web Idle Timeout	Enter the maximum amount of time the switch waits before ending inactive web browser management sessions. The range is 3 to 60 minutes. The default is 10 minutes.
Group Interval	Enter the maximum amount of time the switch waits before ending inactive SNMPv3 management sessions. The range is 120 to 1225 seconds. The default is 120 seconds. Entering 0 disables the timer so that SNMPv3 sessions are never timed out.
CLI Idle Timeout	Enter the maximum amount of time the switch waits before ending inactive command line management sessions. The range is 3 to 60 minutes. The default is 10 minutes.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Enabling or Disabling HTTP or HTTPS Web Browser Management

The switch has a web browser server so that you can remotely manage the device with a web browser from your management workstation. The server has two modes:

- HTTP
- HTTPS (SSL)

Although HTTP is the default active mode for the web server, it is not recommended because it is not secure. The packets exchanged by the switch and your workstation during HTTP management sessions are transmitted in plain text. This makes them vulnerable to snooping from network intruders and may compromise the security of the switch and your network. In contrast, HTTPS management sessions are secure because the transmitted packets are encrypted.

Here are the guidelines:

- The default active mode is HTTP.
- Both HTTP and HTTPS cannot be active on the switch at the same time.
- Changing the web server mode interrupts your management session. You will have to log on again.
- If you enable the HTTPS (SSL) mode, remember to include the prefix "HTTPS://" in the URL field of your web browser when specifying the IP address of the switch at the start of your management sessions.

Perform the following procedure to enable or disable HTTP or HTTPS on the switch:

1. Select **System** -> **SSL Settings** from the main menu. Refer to Figure 39.




Figure 39. SSL Settings Menu Selection

The SSL Settings window is shown in Figure 40.



Figure 40. SSH Settings

2. Select one of the following options:
 - Enabled:** Enables the SSL mode for encrypted HTTPS management sessions, and disables non-secure HTTP management.
 - Disabled:** Disables the SSL mode and HTTPS management, and activates non-secure HTTP management. This is the default setting.
3. Click the **Apply** button.
4. The switch displays a confirmation prompt.
5. Click **OK** to implement the change or **Cancel** to cancel the change.

Note

If you click OK, your management session is interrupted.

6. To resume managing the switch, start a new management session. When specifying the IP address of the switch in the URL field of the web browser, include the prefix “HTTPS://” if you enabled SSL or “HTTP://” if you disabled it.
7. Select **Save Settings to Flash** from the main menu to save your changes.

Enabling or Disabling the Secure Shell Server for the Command Line Interface

The switch has two management interfaces. The primary interface is the web browser interface, which is explained in this guide. You can configure all the features and functions of the switch from the web browser interface.

The switch also has a command line management interface that has a series of commands you enter at a command line prompt. You can configure only a limited number of features of the switch with the command line interface.

You can access the command line interface with either Telnet or Secure Shell (SSH). You need to have a Telnet or SSH client on your management workstation. SSH is recommended because it protects your management sessions from snooping by encrypting the packets. In contrast, management sessions conducted with Telnet are conducted in clear text, meaning packets are not encrypted, leaving your switch and network vulnerable to network intruders. The default settings are the Telnet server is enabled and the SSH server is disabled.

Note

You cannot disable the Telnet server on the switch.

Perform this procedure to enable or disable the SSH server:

1. Select **System** -> **SSH Settings** from the main menu. Refer to Figure 41.




Figure 41. SSH Menu Selection

The SSH Settings window is shown in Figure 42.




Figure 42. SSH Settings Window

2. Configure the settings in Table 16.

Table 16. SSH Settings Window

State	Description
SSH Status	Select one of the following: <ul style="list-style-type: none"> - Enabled: Enables both the SSH server. - Disabled: Disables the SSH servers. This is the default setting.
Port	Enter the SSH protocol port number. The range is 1 to 65535. The default is 22.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Identifying Authorized Management Workstations

This is a security feature of the web browser management interface and the Telnet and SSH command line interfaces. It lets you define the management workstations that are authorized to access the management interfaces. You identify the workstations by their IPv4 or IPv6 addresses. After activating this feature, you can manage the switch only from those workstations whose IP addresses you enter in the table.

Note

Your management session will be interrupted if you activate this feature without adding the IP address of your workstation to the list. To resume managing the switch, you have to start a management workstation from another workstation whose IP address you entered in the table or change the IP address of your workstation to an approved address.

Perform the following procedure to enter the IPv4 or IPv6 addresses of management workstations that are approved to manage the switch:

1. Select **System** -> **IP Access List** from the main menu. Refer to Figure 43.




Figure 43. IP Access List Menu Selection

The IP Access List window is shown in Figure 44 on page 94.

IP Access List

IP Restriction Status:

IP Address: . . . IPv4
 IPv6

IP Access Table

Index	Accessible IP	Action
1	192.168.1.87	<input type="button" value="Delete"/>
2	192.168.1.94	<input type="button" value="Delete"/>

Figure 44. IP Access List Window

2. To enable or disable the feature, select **Enabled** or **Disabled** from the IP Restriction Status pull-down menu and click the **Apply** button.

Note

Do not enable the feature without first entering the IP address of your workstation. Otherwise, your management session will be interrupted and you will not be able to restart it from your workstation.

3. To add an address, do the following:
 - a. To add an IPv4 address, enter the address in the **IPv4** fields.
 - b. To add an IPv6 address, click the IPv6 radio circle and enter the address in this format:

xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx

Each x is a hexadecimal digit representing 4 bits. The eight groups of digits have to be separated by colons. Groups where all four digits are '0' can be omitted. Leading '0's in groups can also be omitted. For example, the following IPv6 addresses are equivalent:

3710:421e:09a8:0000:0000:0000:00a4:1c50

3710:421e:9a8::a4:1c50

- c. Click the **Add** button.

4. To delete an address, do one of the following:
 - a. To delete a specific address, click its **Delete** button in the Action column.
 - b. To delete all the addresses, click the **Delete All** button.
The switch displays a confirmation prompt.
 - c. Click **OK** to delete the address or **Cancel** to cancel the change.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 8

Configuration Files

This chapter explains how to upload or download switch configuration files to your management workstation or TFTP server in the following sections:

- ❑ “Overview to Switch Configuration Files” on page 98
- ❑ “Backing Up Configuration Files from the Switch with HTTP” on page 99
- ❑ “Restoring Configuration Files to the Switch with HTTP” on page 101
- ❑ “Backing Up Configuration Files from the Switch with TFTP” on page 102
- ❑ “Restoring Configuration Files to the Switch with TFTP” on page 104
- ❑ “Booting the Switch with the Config2 Configuration File” on page 106

Overview to Switch Configuration Files

The switch immediately implements your changes to its parameter settings as soon you enter them in the web browser interface. However, the changes are not permanently saved and will be lost if you reboot or power cycle the switch without first saving them in its Config1 configuration file in flash memory with the **Save Settings to Flash** from the main menu. Once your changes are saved in the Config1 file, the switch retains them even when booted or powered off.

The switch also has a configuration file called Config2, which you might use for several functions, including these:

- ❑ You might use it as a backup to the Config1 file so that if the latter becomes damaged or corrupted, you can quickly restore the switch settings.
- ❑ You can also use it to return the switch to an earlier configuration. For example, if you make changes to the switch's configuration and save them to Config1 but not Config2, you can restore the previous configuration by instructing the switch to boot up with the Config2 file, as explained in "Booting the Switch with the Config2 Configuration File" on page 106.
- ❑ Finally, you can use the Config2 to download a new configuration for switch, but not immediately implement it on the device. For instance, you might download a new configuration file to the Config2 file early in the day, but not configure the switch with it until non-business hours, to avoid disrupting the work of your network users. This is explained in "Booting the Switch with the Config2 Configuration File" on page 106.

You can upload the Config1 and Config2 configuration files from the switch and store them on your computer or TFTP server, and later download them back to the switches, if needed. This is useful in maintaining configuration histories of switches, configuring replacement units, or quickly configuring switches that are to have similar configurations.

Backing Up Configuration Files from the Switch with HTTP

This section contains the procedure for uploading the Config1 or Config2 configuration file from the switch to your computer. You might perform this procedure to maintain configuration histories of the switches so that you can quickly configure replacement units or switches that are to have similar settings. Refer to “Overview to Switch Configuration Files” on page 98.

Perform the following procedure to upload a configuration file from the switch to your computer:

1. Select **Tools > Config File Backup/Restore > via HTTP** from the main menu. Refer to Figure 45.




Figure 45. Config File Backup/Restore via HTTP Menu Selection

The Config File Backup/Restore via HTTP window is shown in Figure 46.

Config File Backup/Restore via HTTP

Next Bootup Configure File: Config 1 Config 2

Current Configure File: Config 1

Select File:

Config File: Startup-Config

Figure 46. Config File Backup/Restore via HTTP Window

2. Use the Config File option to select the configuration file you want to upload. The default is the Config1 file. To upload the currently designated file, click the **Backup** button. A prompt is displayed for saving the file. Go to step 4.
3. To upload the other configuration file, do the following:

- a. Click the **Startup-Config** option to remove the check mark from its dialog box.
 - b. Select the other configuration file (for example, Config2) from the menu.
 - c. Click the **Backup** button. A prompt is displayed for saving the file.
4. At the prompt, select the location for saving the file on your computer and enter a filename. The default is “config.bin”. Do not change the “bin” extension. The file is save on your computer.

Restoring Configuration Files to the Switch with HTTP

Perform the following procedure to restore a configuration file on your computer to replace the Config1 or Config2 file on the switch:

Note

Replacing the Config1 file causes the switch to immediately reboot, disrupting network operations.

1. Select **Tools > Config File Backup/Restore> via HTTP** from the main menu. Refer to Figure 45 on page 99. The Config File Backup/Restore via HTTP window is shown in Figure 46 on page 99.
2. Click the **Browse** button in the Select File field.
3. Locate the configuration file on your computer.
4. Use the Config File option to select the configuration file you want to replace on the switch with the downloaded file. The default is the Config1 file. To replace the currently designated file, click the **Restore** button.
5. To replace the other configuration file, do the following:
 - a. Click the **Startup-Config** option to remove the check mark from its dialog box.
 - b. Select the other configuration file (for example, Config2) from the menu.
 - c. Click the **Restore** button.

The file is downloaded from your computer to the switch and replaces the designated Config1 or Config2 file. Note the following:

- If you replace the Config1 file, the switch automatically reboots to load the new configuration.
- Restoring to Config2 behavior is the same as Config1, the switch will automatically reboot after the restore configuration is finished.
- If the IP address for the switch in the new configuration file is different than its previous address, be sure to use the new address when starting new management sessions.

Backing Up Configuration Files from the Switch with TFTP

Perform the following procedure to backup the Config1 or Config2 configuration file on the switch to a TFTP server:

1. Start the TFTP server on your network.
2. Select **Tools > Config File Backup/Restore > via TFTP** from the main menu. Refer to Figure 47.

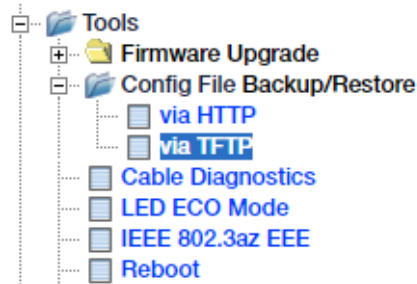


Figure 47. Config File Backup/Restore via TFTP

The Config File Backup/Restore via TFTP Window is shown in Figure 48.

Config File Backup/Restore via TFTP

TFTP Server IP: . . . IPv4
 IPv6

Config File: Startup-Config

Config File Name: (Max. length: 64 characters)

Figure 48. Config File Backup/Restore via TFTP Window

3. Enter the IP address of the TFTP server in the TFTP Server IP field. It can be either an IPv4 or IPv6 address. The default is IPv4.
4. Use the Config File option to select the configuration file you want to upload. The default is the Config1 file. To upload the currently designated file, skip step 5 and go to step 6.
5. To upload the other configuration file, do the following:
 - a. Click the **Startup-Config** option to remove the check mark from its dialog box.

- b. Select the configuration file (for example, Config2) from the Config File menu.
6. Enter a filename for the configuration file in the **Config File Name** field. The file will be given this filename on the TFTP server. The maximum length is 64 alphanumeric characters. Spaces and special characters are not allowed. The extension has to be "bin".
7. Select the **Backup** button. The switch transmits the file to the TFTP server.

Restoring Configuration Files to the Switch with TFTP

Perform the following procedure to restore the Config1 or Config2 configuration file to the switch from a TFTP server:

Note

Replacing the Config1 file causes the switch to immediately reboot, disrupting network operations.

1. Start the TFTP server on your network and store the configuration file on the server.
2. Select **Tools > Config File Backup/Restore > via TFTP** from the main menu. Refer to Figure 47 on page 102. The Config File Backup/Restore via TFTP window is shown in Figure 48.
3. Enter the IP address of the TFTP server in the TFTP Server IP field. It can be either an IPv4 or IPv6 address. The default is IPv4.
4. Use the Config File option to select either the Config1 or Config2 file to replace with the downloaded file. Do one of the following:
 - To replace the currently designated file, skip to step 5.
 - To replace the other configuration file, do the following:
 - a. Click the **Startup-Config** option to remove the check mark from its dialog box.
 - b. Select the other configuration file (for example, Config2) from the menu.
5. In the **Config File Name** field, enter the filename of the configuration file on the TFTP server. If necessary, include the directory path of the file on the server. The extension has to be “bin”.
6. Click the **Restore** button.

The file is downloaded from the TFTP server to the switch. It replaces the Config1 or Config2 file, according to your selection in step 4. Note the following:

- If you replace the Config1 file, the switch automatically reboots to load its new configuration.
- Restoring to Config2 behavior is the same as Config1, the switch will automatically reboot after the restore configuration is finished.

- ❑ If the IP address for the switch in the new configuration file is different than its previous address, be sure to use the new address when starting new management sessions.

Booting the Switch with the Config2 Configuration File

Under normal operating conditions, the switch uses the Config1 configuration file to restore its settings when booted or power cycled. However, you can, if needed, instruct the switch to temporarily boot with the Config2 file instead. You might do this if the Config1 file is damaged or corrupted, or if the Config2 file has a newer configuration for the switch.

Here is an example scenario. Assume you are about to download a new configuration file to the switch but do not want the switch to implement it yet. So you download it as the Config2 file on the switch, which stores it in flash memory but does not use it. At a later time, when you are ready to configure the switch with the settings in Config2, you perform this procedure to boot the switch with the file.

It is important to know that booting the switch with the Config2 file does not change the fact that the Config1 file is still the active configuration file and that its contents are not changed. If you want to overwrite the configuration in the Config1 file with the settings from the Config2 file, select **Save Settings to Flash** from the main menu.

Perform the following procedure to instruct the switch to use the Config2 file to configure its settings at the next reboot:

1. Select **Tools > Config File Backup > via HTTP** from the main menu. Refer to Figure 45 on page 99. The Config File Backup/Restore via HTTP window is shown in Figure 46 on page 99.
2. Click the **Config2** radio button in Next Bootup Configure File option.
3. Click the **Apply** button.
4. Boot the switch. Refer to “Rebooting the Switch” on page 500.
5. Wait two minutes for the switch to initialize its management software, and then start a new management session to resume managing the unit. Note the following:
 - ❑ If the IP address for the switch in the Config2 file is different than its previous address, be sure to use the new address to start new management sessions.
 - ❑ At this point the switch is operating with the configuration from the Config2 file. However, the Config1 file, which the switch normally uses to store and configure its settings after a reboot, is unchanged. To instruct the switch to retain and use the new configuration after future reboots or power cycles, save it in the Config1 file with **Save Setting to Flash** from the main menu.

Chapter 9

LED ECO Mode and Energy-Efficient Ethernet

This chapter has the following sections:

- “Enabling or Disabling the LED ECO Mode” on page 108
- “Enabling or Disabling IEEE 802.3az Energy-Efficient Ethernet” on page 109

Enabling or Disabling the LED ECO Mode

You can turn off the port LEDs on the front panel of the switch to conserve electricity when you are not using them to monitor the device. This is referred to as the LED ECO mode. Turning off the port LEDs does not interfere with the network operations of the switch.

Perform the following procedure to enable or disable the LED ECO mode:

1. Select **Tools > LED ECO Mode** from the main menu. Refer to Figure 49.




Figure 49. LED ECO Mode Menu Selection

The LED ECO Mode is shown in Figure 50.




Figure 50. LED ECO Mode Page

2. Select one of the following from the LED ECO Mode menu:
 - Enabled:** Turns off the port LEDs.
 - Disabled:** Turns on the port LEDs. This is the default setting.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Enabling or Disabling IEEE 802.3az Energy-Efficient Ethernet

The switch supports IEEE 802.3az Energy-Efficient Ethernet (EEE). EEE is an energy saving feature that reduces power consumption during periods of no data activity. When the feature is enabled, the switch saves electricity by placing the Ethernet circuitry in a special sleep mode when all the ports are inactive. When data activity resumes, the circuitry automatically resumes normal operations. The default setting for the feature is disabled.

Perform the following procedure to enable or disable IEEE 802.3az EEE:

1. Select **Tools > IEEE 802.3az EEE** from the main menu. Refer to Figure 51.




Figure 51. IEEE 802.3az EEE Menu Selection

The IEEE 802.3az EEE window is shown in Figure 52.

IEEE 802.3az EEE

IEEE 802.3az Energy Efficient Ethernet Status:

Figure 52. IEEE 802.3az EEE Window

2. Select **Enabled** or **Disabled** from the menu to enable or disable the feature. The default is disabled.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 10

Power over Ethernet (PoE)

This chapter describes Power over Ethernet in the following sections:

- “Overview” on page 112
- “Configuring PoE on the Ports” on page 115

Overview

The GS950 PS V2 Switches feature PoE on the copper ports. This feature allows the switch to supply power to network devices over the same cables that carry the network traffic. The value of PoE is that it can make installing a network easier. Selecting locations for network devices are often limited by whether there are power sources nearby. This often limits equipment placement or requires the added time and cost of having additional electrical sources installed. But with PoE, you can install PoE-compatible devices wherever they are needed without having to worry about whether there are adjacent power sources.

A device that provides PoE to other network devices is referred to as *power sourcing equipment* (PSE). The GS950 PS V2 Switches act as PSE units by adding DC power on the network cables connected to its ports, thus functioning as a central power source for other network devices.

Devices that receive their power from a PSE are called *powered devices* (PD). Examples include wireless access points, IP telephones, webcams, and even other Ethernet switches.

The switches automatically determine whether or not a device connected to a port is a powered device. Ports that are connected to network nodes that are not powered devices (that is, devices that receive their power from another power source) function as regular Ethernet ports, without PoE. The PoE feature remains activated on the ports but no power is delivered to the devices.

Maximum PoE Budgets

The maximum PoE budgets are the total amounts of power the switches can supply to powered devices on their ports. The maximum PoE budgets of the switches are listed in Table 17.

Table 17. PoE Maximum Power Budgets

Switch	PoE Budget	PoE Ports
GS950/10PS V2	75W	1 to 8
GS950/18PS V2	185W	1 to 16
GS950/28PS V2	185W	1 to 24
GS950/52PS V2	370W	1 to 24 ^a

a. Copper ports 25 to 48 on the GS950/52PS V2 Switch do not support PoE.

PoE Standards

The GS950 PS V2 Switches support these PoE standards:

- ❑ PoE (IEEE 802.3af): This standard provides up to 15.4 watts at the switch port to support powered devices that require up to 12.95 watts.
- ❑ PoE+ (IEEE 802.3at): This standard provides up to 30.0 watts at the switch port to support powered devices that require up to 25.5 watts.

Powered Device Classes

Powered devices are grouped into classes, based on their power requirements. The GS950 PS V2 Switches support the five classes listed in Table 18.

Table 18. IEEE Powered Device Classes

Class	Maximum Power Output at the Switch Port	Powered Device Power Range
0	15.4W	0.44W to 12.95W
1	4.0W	0.44W to 3.84W
2	7.0W	3.84W to 6.49W
3	15.4W	6.49W to 12.95W
4	30.0W	12.95W to 25.5W

Note

The switches can support any combination of PoE class devices up to their maximum PoE power budgets.

PoE Port Priorities

If the power requirements of the powered devices exceed the switch's power budget, the switch will deny power to some ports based on a system called PoE port priorities. You can use this feature to ensure that powered devices critical to the operations of your network are given preferential treatment by the switch in the allocation of power should the demands of the devices exceed the available power budget.

There are three priority levels:

- ❑ Critical
- ❑ High
- ❑ Low

Ports set to the Critical level, the highest priority level, are guaranteed power before any of the ports assigned to the other two priority levels. Ports assigned to the other priority levels receive power only if all the

Critical ports are receiving power. Ports that are connected to your most critical powered devices should be assigned to this level. If there is not enough power to support all the ports set to the Critical priority level, power is allocated to ports based on port number, in ascending order.

The High level is the second highest level. Ports set to this level receive power only if all the ports set to the Critical level are already receiving power. If there is not enough power to support all of the ports set to the High priority level, power is provided to the ports based on port number, in ascending order.

The lowest priority level is Low. This is the default setting. Ports set to this level receive power only if all of the ports assigned to the other two levels are already receiving power. As with the other levels, if there is not enough power to support all of the ports set to the Low priority level, power is provided to the ports based on port number, in ascending order.

Power allocation is dynamic. Ports supplying power to powered devices can cease power transmission if the switch's power budget is at maximum usage and new powered devices connected to ports with higher priorities become active.

Configuring PoE on the Ports

Perform the following procedure to configure PoE on the ports:

1. Select **Power over Ethernet** from the main menu. Refer to Figure 53

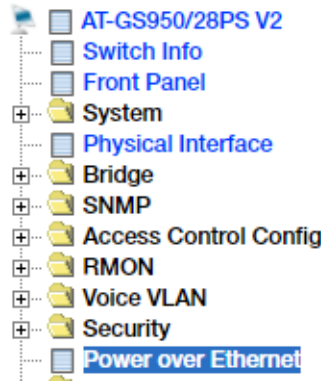


Figure 53. Power Over Ethernet Menu Selection

The Power over Ethernet window is shown in Figure 54.

Power Over Ethernet						
Power Budget:		185 W				
Power Consumption:		105 W				
Power Remaining:		80 W				
Port	Admin	Status	Class	Priority	Power(mW)	Action
All	Ignore ▼	-	-	Ignore ▼	-	Apply
1	Enabled ▼	POWER ON	Class4	Low ▼	5400	Apply
2	Enabled ▼	POWER ON	Class4	Low ▼	5400	Apply
3	Enabled ▼	POWER ON	Class4	Low ▼	7500	Apply
4	Enabled ▼	POWER ON	Class4	Low ▼	7500	Apply

Figure 54. Power Over Ethernet Window

The three fields at the top of the window are defined in Table 20.

Table 19. Power Over Ethernet Window - Power Budget Fields

Field	Description
Power Budget	Displays the switch's total power budget for powered devices. Refer to Table 17 on page 112.
Power Consumption	Displays the total consumed power by powered devices.

Table 19. Power Over Ethernet Window - Power Budget Fields

Field	Description
Power Remaining	Displays the remaining unused power budget.

2. Configure the parameters in Table 20.

Table 20. Power Over Ethernet Window - Port Table

Column	Description
Port	Displays the port number.
Admin	Enable or disable PoE on the port by selecting one of the following: <ul style="list-style-type: none"> - Enabled: Enables PoE on the port. This is the default setting. - Disabled: Disables PoE.
Status	Displays PoE status on the port: The PoE port status is given as follows: <ul style="list-style-type: none"> - POWER ON: The port is supplying PoE power to a powered device. - POWER OFF: The port is not supplying PoE power. The port is not connected to a powered device, PoE is disabled on the port, or the switch does not have sufficient free power because it has reached its power budget.
Class	Displays the PoE Class of the device. Refer to Table 18 on page 113. The column displays N/A if the port is not providing power to a powered device.
Priority	Select one of the following to set the port's PoE priority: <ul style="list-style-type: none"> - Low: This is the default setting. - High - Critical Refer to "PoE Port Priorities" on page 113.
Power(nW)	Displays the power consumption (milliwatts) of the powered device.
Action	Contains the Apply button.

3. If you changed a port setting, click its **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 11

Link Layer Discovery Protocol

This chapter describes Link Layer Discovery Protocol (LLDP) in the following sections:

- ❑ “LLDP Overview” on page 120
- ❑ “Configuring LLDP” on page 121
- ❑ “Displaying Neighbor Information” on page 124

LLDP Overview

Link Layer Discovery Protocol (LLDP) allows Ethernet network devices, such as switches and routers, to share and store device-related information with each other. Neighboring devices that use LLDP can advertise parts of their Layer 2 configuration to each other, enabling you to identify several types of incorrect configurations more easily.

LLDP is a “one-hop” protocol. LLDP information can only be sent to and received by devices that are directly connected to each other, or connected via a hub or repeater. Devices that are directly connected to each other are called neighbors. Advertised information is not forwarded on to other devices on the network. Also, LLDP is a one-way protocol. That is, the information transmitted in LLDP advertisements flows in one direction only, from one device to its neighbors, and the communication ends there. Transmitted advertisements do not solicit responses, and received advertisements do not solicit acknowledgments. LLDP cannot solicit any information from other devices. LLDP operates over physical ports only. For example, you can configure it on switch ports that belong to static port trunks or LACP trunks, but not on the trunks themselves, and on switch ports that belong to VLANs, but not on the VLANs themselves.

Each port can be configured to transmit local information, receive neighbor information, or both. LLDP transmits information in packets called LLDP Data Units (LLDPDUs). An LLDPDU consists of a set of Type-Length-Value elements (TLV), each of which contains a particular type of information about the device or port transmitting it.

Note

GS950 PS V2 Switches do not support Link Layer Discovery Protocol for Media Endpoint Devices (LLDP-MED).

Configuring LLDP

Perform the following procedure to configure LLDP:

1. Select **LLDP > LLDP Global Settings** from the main menu. Refer to Figure 55.



Figure 55. LLDP Global Settings

The LLDP Global Settings window is shown in Figure 56.

LLDP Global Settings

LLDP

Message TX Hold Multiplier (2-10)
Message TX Interval (5-32768) sec
LLDP Reinit Delay (1-10) sec
LLDP TX Delay (1-8192) sec

Note : (LLDP TX Delay)<= (0.25* (Message TX Interval)) and (Message TX Interval)*(Message TX Hold Multiplier)<65535.

LLDP System Information	
Chassis ID Subtype	
Chassis ID	
System Name	
System Description	

Figure 56. LLDP Global Settings Window

2. To enable or disable LLDP on the switch, perform the following:
 - a. Select **Enabled** or **Disabled** from the LLDP menu to enable or disable the feature. The default is disabled. You have to enable LLDP to configure the settings.
 - b. Click the **Apply** button.

3. To configure the global settings that apply to all switch ports, perform the following steps:
 - a. Configure the fields in Table 21.

Table 21. LLDP Global Settings

Field	Description
Message TX Hold Multiplier	Enter the multiplier value for the Message TX Interval. This is used to calculate the Time To Live (TTL) that the switch advertises to neighbors. TTL controls the length of time in seconds that advertisements are valid. The range is 2 to 10. The default is 4.
Message TX Interval	Enter the interval between regular transmissions of LLDP advertisements. The interval must be a least four times the LLDP TX Delay. The range is from 5 to 32768 seconds. The default is 30 seconds.
LLDP Reinit Delay	Enter the number of seconds that must elapse after LLDP is disabled on a port before it can be reinitialized. The range is from 1 to 10 seconds. The default is 2 seconds.
LLDP TX Delay	Enter the minimum time interval between transmissions of LLDP advertisements from changes to LLDP local information. The transmission delay timer cannot be greater than a quarter of the Message TX Interval. The range is from 1 to 8192 seconds. The default is 2 seconds.

- b. Click the **Apply** button.
4. To configure port settings, perform the following:
 - a. Choose a setting for a port from the State menu. Refer to Table 22.

Table 22. LLDP Port Settings

Field	Description
Disabled	Configures the port to block ingress and egress LLDPDUs.
RxTx	Configures the port to transmit and receive LLDPDUs. This is the default setting.

Table 22. LLDP Port Settings (Continued)

Field	Description
RxOnly	Configures the port to receive but not transmit LLDPDUs.
TxOnly	Configures the port to transmit but not receive LLDPDUs.

b. Click the **Apply** button in the Action column.

5. Select **Save Settings to Flash** from the main menu to save your changes.

The fields in the LLDP System Information section of the window are described in Table 23.

Table 23. LLDP System Information

Field	Description
Chassis ID Subtype	Displays “MAC Address” as the Chassis ID subtype. This value is not adjustable.
Chassis ID	Displays the switch’s MAC address. This value is not adjustable.
System Name	Displays the system name of the switch. Refer to “Setting the Name, Location, and Contact” on page 56.
System Description	Displays the product name. This value is not adjustable.

Displaying Neighbor Information

To view information from neighboring LLDP devices, select **LLDP > LLDP Neighbors Information** from the main menu. Refer to Figure 57.



Figure 57. LLDP Neighbors Information Menu Selection

The LLDP Neighbors Information window is shown in Figure 58.




Figure 58. LLDP Neighbors Information Window

The table columns are described in Table 24.

Table 24. LLDP Neighbors Information Window

Column	Description
Entity	Displays the number the switch assigned to the reporting neighbor in the order it received the LLDP information.
Port	Displays the port number that received the LLDP information.
Chassis ID Subtype	Displays Chassis ID subtype of the neighboring network device.
Chassis ID	Displays Chassis ID of the neighboring network device. For Allied Telesis products, Chassis ID is the device’s MAC address.
Port ID Subtype	Displays the Port ID subtype of the port on the neighboring network device port.
Port ID	Displays the port number on the neighboring network device port
Port Description	Displays a description of the neighboring port.

Table 24. LLDP Neighbors Information Window (Continued)

Column	Description
Show Detail	Click the button to display a detailed report on the neighboring port.

Section II

Spanning Tree Protocols

This section contains the following chapters:

- ❑ Chapter 12, “Spanning Tree and Rapid Spanning Tree Protocols” on page 129
- ❑ Chapter 13, “Multiple Spanning Tree Protocol Overview” on page 147
- ❑ Chapter 14, “Multiple Spanning Tree Protocol” on page 163

Chapter 12

Spanning Tree and Rapid Spanning Tree Protocols

This chapter explains the Spanning Tree Protocol (STP) and Rapid Spanning Tree Protocol (RSTP) in the following sections:

- ❑ “STP and RSTP Overview” on page 130
- ❑ “Configuring STP and RSTP Global Settings” on page 138
- ❑ “Configuring STP and RSTP Port Settings” on page 142

STP and RSTP Overview

The performance of an Ethernet network can be adversely affected by the existence of physical loops in the network topology. Loops exist when two or more nodes on a network can transmit packets to each other over more than one physical path. Packets can become caught in repeating cycles, referred to as broadcast storms, that needlessly consume network bandwidth and can significantly reduce network performance.

STP and RSTP prevent data loops from forming by ensuring that only one path exists between the end nodes in your network. Where multiple paths exist, these protocols place the extra paths in a standby or blocking mode, leaving only one main active path. STP and RSTP can also activate a redundant path if a main path goes down, thereby maintaining network connectivity.

Where the two protocols differ is in the time each takes to complete the process referred to as *convergence*. With the convergence process, when a change is made to the network topology, such as the addition of a new bridge, the spanning tree protocol has to determine whether there are redundant paths that must be blocked to prevent data loops, or activated to maintain connections between the various network segments.

With STP, convergence can take up to a minute or more to complete in a large network. This can result in the loss of communications between various parts of the network during the convergence process and the subsequent loss of data packets.

RSTP is much faster. It can complete convergence in seconds, and as such, greatly diminishes the possible impact the process can have on your network. The STP implementation complies with the IEEE 802.1d standard.

Only one spanning tree at a time can be active on the switch.

The following subsections provide a basic overview on how STP and RSTP operate and define the different adjustable parameters. For further information about STP and RSTP, refer to IEEE Std 802.1D and IEEE Std 802.1w, respectively.

Bridge Priority and the Root Bridge

When a spanning tree protocol is activated on a network, the bridges select a *root bridge*. A root bridge distributes network topology information to the other network bridges and is used by the other bridges to determine if there are redundant paths in the network.

A root bridge is selected by the *bridge priority* number, also referred to as the bridge identifier, and sometimes the bridge's MAC address. The bridge with the lowest bridge priority number in the network is selected as the root bridge. If two or more bridges have the same lowest bridge priority

number, the one with the lowest MAC address is designated as the root bridge.

You can change the bridge priority number on the switch. You can designate which switch on your network is the root bridge by giving it the lowest bridge priority number. You may also consider which bridge should function as the backup root bridge in the event you need to take the primary root bridge off line and assign that bridge the second lowest bridge identifier number.

The bridge priority has a range of 0 to 61440 in increments of 4096. You specify the increment that represents the desired bridge priority value. The range is divided into the following sixteen increments:

- 0
- 4096
- 8192
- 12288
- 16384
- 20480
- 24576
- 28672
- 32768
- 36864
- 40960
- 45056
- 49152
- 53248
- 57344
- 61440

Path Costs and Port Costs

After selecting the root bridge, the bridges determine if the network contains redundant paths and, if one is found, select a preferred path while placing the redundant paths in a backup or blocking state.

Where there is only one path between a bridge and the root bridge, the bridge is referred to as the *designated bridge*, and the port through which the bridge is communicating with the root bridge is referred to as the *root port*.

If redundant paths exist, the bridges that are a part of the paths must determine which path is the primary, active path, and which path(s) are

placed in the standby, blocking mode. This is accomplished by calculating *path costs*. The path offering the lowest cost to the root bridge becomes the primary path, and all other redundant paths are placed into a blocking state.

Path cost is determined by evaluating *port costs*. Every port on a bridge participating in STP has a cost associated with it. The cost of a port on a bridge is typically based on port speed. The faster the port, the lower the port cost. The exception to this is the ports on the root bridge, where all ports have a port cost of 0.

Path cost is the sum of the port costs between a bridge and the root bridge.

The port cost of a port on the switch is adjustable. For STP and RSTP, the range is from 0 to 200,000,000.

Port Priority

If two paths have the same port cost, the bridges must select a preferred path. In some instances, this can involve the use of the *port priority* parameter which is used as a tie breaker when two paths have the same cost.

The range for port priority is 0 to 240. As with bridge priority, this range is broken into increments, in this case, multiples of 16. To select a port priority for a port, you enter the desired value. Table 25 on page 132 lists the values that are valid.

Table 25. Valid Port Priority Values

Step	Port Priority
1	0
2	16
3	32
4	48
5	64
6	80
7	96
8	112
9	128
10	144

Table 25. Valid Port Priority Values (Continued)

Step	Port Priority
11	160
12	176
13	192
14	208
15	224
16	240

Forwarding Delay and Topology Changes

If there is a change in the network topology due to a failure, removal, or addition of active components, the active topology might also change. This might trigger a change in the state of some blocked ports. However, a change in a port state is not activated immediately.

It may take time for the root bridge to notify all bridges that a topology change has occurred, especially if it is a large network. A temporary data loop could occur if a topology change is made before all bridges have been notified and that could adversely impact network performance.

To forestall the formation of temporary data loops during topology changes, a port designated to change from blocking to forwarding passes through two additional states - listening and learning - before it begins to forward frames. The amount of time a port spends in these states is set by the forwarding *delay* value. This value states the amount of time that a port spends in the listening and learning states prior to changing to the forwarding state.

The forwarding delay value is adjustable on the switch. The appropriate value for this parameter depends on a number of variables; the size of your network is a primary factor. For large networks, you should specify a value large enough to allow the root bridge sufficient time to propagate a topology change throughout the entire network. For small networks, you should specify a smaller value so that the time for a topology change is optimized for minimum data loss.

Note

The forwarding delay parameter applies only to ports that are operating in the STP mode.

Hello Time and Bridge Protocol Data Units (BPDU)

The bridges that are part of a spanning tree domain communicate with each other using a bridge broadcast frame that contains a special section devoted to carrying STP or RSTP information. This portion of the frame is referred to as the bridge protocol data unit (BPDU). When a bridge is brought on-line, it issues a BPDU in order to determine whether a root bridge has already been selected on the network, and if not, whether it has the lowest bridge priority number of all the bridges and should therefore become the root bridge.

The root bridge periodically transmits a BPDU to determine whether there have been any changes to the network topology and to inform other bridges of topology changes. The frequency with which the root bridge sends out a BPDU is called the *hello time*. This is a value that you can set on the switch. The interval is measured in seconds. If the switch is selected as the root bridge of a spanning tree domain, and the hello time is set to the default of two seconds, it transmits a BPDU every two seconds.

Point-to-Point and Edge Ports

This section applies only to RSTP. Part of the task of configuring RSTP is defining the port types on the bridge, which is directly related to the device(s) connected to the port. With the port types defined, RSTP can reconfigure a network much quicker than STP when a change in network topology is detected.

There are two possible selections:

- Point-to-point port
- Edge port

If a bridge port is connected to another bridge or router port, it normally operates in full-duplex mode and is functioning as a point-to-point port. Figure 59 on page 135 illustrates two switches that are connected with one data link. This link is operating between two point-to-point ports.




Figure 59. Point-to-Point Ports

A port operates as an edge port when it is connected to a network terminal device such as a workstation or a server. An edge port on a bridge should not have any STP or RSTP devices connected to it either directly or through another device connected to that port. In this configuration, since the port has no STP or RSTP devices connected to it, it will always forward network traffic. Figure 60 illustrates a port functioning as an edge port.




Figure 60. Edge Port

Mixed STP and RSTP Networks


RSTP IEEE 802.1w is fully compliant with STP IEEE 802.1d. As such, RSTP is compatible with STP. A network can have bridges running STP and RSTP.

If you decide to activate spanning tree on the switch, Allied Telesis recommends RSTP instead of STP, even if all of other switches in the network are running STP. The switch can combine RSTP with the STP of the other switches. The switches monitor the traffic on each port for BPDU packets. Ports that receive RSTP BPDU packets operate in RSTP mode, while ports receiving STP BPDU packets operate in STP mode.

Spanning Tree and VLANs

The spanning tree implementation on the switch supports a single-instance spanning tree that encompasses all switch ports. If the ports are grouped into VLANs, spanning tree crosses VLAN boundaries. This can pose a problem in networks containing multiple VLANs that span two bridges and are connected with untagged ports. In this situation, spanning tree might block data links because it detects suspected data loops, which can cause VLAN fragmentation.

This is illustrated in Figure 61. VLANs 1 to 3 span two switches. The separate VLAN parts on the two switches are connected by dedicated links of untagged ports.



U = Untagged VLAN port

Figure 61. STP and VLAN Fragmentation with Untagged Ports

When STP or RSTP is activated on the switches, two links are disabled because the three links represent loops, even though they reside in different VLANs. Refer to Figure 62 on page 137. As a result, two VLANs are disconnected between the bridges. In this example, the ports on the top switch linking the two parts of the VLANs 2 and 3 are changed to the blocking state, disrupting their VLAN connections.





Figure 62. STP and VLAN Fragmentation with Untagged Ports

You can avoid this problem by either not activating the spanning tree protocol or connecting the switches using tagged instead of untagged ports. As explained in Chapter 16, “802.1Q Tagged Virtual LANs” on page 205, a link of tagged ports can carry traffic from multiple VLANs simultaneously. Refer to Figure 63.



T = Tagged VLAN port

Figure 63. STP and VLAN Compatibility with Tagged Ports

Configuring STP and RSTP Global Settings

Perform the following procedure to enable or disable spanning tree protocol on the switch, choose STP or RSTP as the active protocol. and configure global (non port-specific) settings:

Note

You have to disable BPDU pass-through on all the ports before you can enable the spanning tree protocol. Refer to “Configuring Basic Port Settings” on page 51.

Note

The switch briefly stops forwarding Ethernet traffic when spanning tree is enabled or disabled, need to wait for forward delay time.

1. Select **Bridge > Spanning Tree > Protocol Settings** from the main menu. Refer to Figure 64.

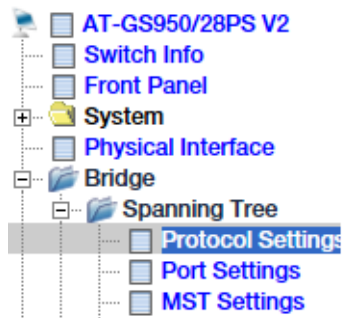


Figure 64. Protocol Settings Menu Selection

The Spanning Tree Protocol Settings page is shown in Figure 65 on page 139.

Spanning Tree Protocol Settings

Global STP Status:

Protocol Version:

Bridge Priority:

Maximum Age: (6-40) sec

Hello Time: (1-10) sec

Forward Delay: (4-30) sec

Transmit Hold Count: (1-10)

Max Hop Count: (6-40)

Note: Enabling Spanning-Tree will temporarily cause the system to stop responding.

Root Bridge: 00:00:00:00:00:00:00:00

Root Cost: 0

Root Maximum Age: 20 sec

Root Forward Delay: 15 sec

Root Port: 0

Figure 65. Spanning Tree Protocol Settings Window

2. Configure the settings in Table 26:

Table 26. Spanning Tree Protocol Settings Window for STP and RSTP

Field	Definition
Global STP Status	Select one of the following from the menu: <ul style="list-style-type: none"> - Enabled: Activates the spanning tree protocol. You have to enable the protocol to configure the settings in the window. - Disable: Deactivates the spanning tree protocol. This is the default setting.
Protocol Version	Select the spanning tree version from the menu: <ul style="list-style-type: none"> - STP - RSTP - MSTP

Table 26. Spanning Tree Protocol Settings Window for STP and RSTP

Field	Definition
Bridge Priority	<p>Select the priority number for the bridge. This number is used to determine the root bridge of the spanning tree domain. The bridge with the lowest priority number is selected as the root bridge. If two or more bridges have the same priority value, the bridge with the numerically lowest MAC address becomes the root bridge. If a root bridge goes off-line, the bridge with the next lowest priority number automatically becomes the root bridge. The range is 0 (zero) to 61,440 in increments of 4096. The highest priority is 0.</p>
Maximum Age	<p>Enter the maximum length of time the bridge stores bridge protocol data units (BPDUs). Bridges in a bridged LAN use the aging time to test the age of stored configuration messages called bridge protocol data units (BPDUs) and to delete expired data units. For example, when set to the default value 20 seconds, bridges delete configuration messages after 20 seconds. The range is 6 to 40 seconds.</p> <p>Observe the following rules when setting the maximum age:</p> <ul style="list-style-type: none"> - Maximum Age must be greater than $(2 \times (\text{Hello Time} + 1))$ - Maximum Age must be less than $(2 \times (\text{Forward Delay} - 1))$ - The aging time for BPDUs is different from the aging time for the MAC address table.
Hello Time	<p>Enter the time interval at which the root bridge transmits BPDUs to determine whether there have been any changes to the network topology and to inform other bridges of topology changes. The range is 1 to 10 seconds. The default is 2 seconds.</p>

Table 26. Spanning Tree Protocol Settings Window for STP and RSTP

Field	Definition
Forward Delay	Enter the waiting period in seconds before a bridge changes to a new state, for example, becomes the new root bridge after a topology change. If the bridge transitions too soon, not all links may have yet adapted to the change, resulting in network loops. The range is 4 to 30 seconds. The default is 15 seconds.
Transmit Hold Count	Applies to MSTP. Refer to “Configuring MSTP Global Settings” on page 164.
Max Hop Count	Applies to MSTP. Refer to “Configuring MSTP Global Settings” on page 164.
Root Bridge	Displays the MAC address of the root bridge of the spanning tree domain. This value is zero when the spanning tree protocol is disabled.
Root Cost	Displays the sum of the root port costs of the bridges between the switch’s root port and the root bridge.
Root Maximum Age	Displays the maximum age setting on the root bridge.
Root Forward Delay	Displays the forward delay setting on the root bridge.
Root Port	Displays the port on the switch leading to the root bridge. This value will be zero when the switch is the root bridge or spanning tree is disabled.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Configuring STP and RSTP Port Settings

Perform the following procedure to configure the ports for STP and RSTP:

Note

You have to enable spanning tree on the switch before you can configure the ports. Refer to “Configuring STP and RSTP Global Settings” on page 138.

1. Select **Bridge > Spanning Tree > Port Settings** from the main menu. Refer to Figure 66.

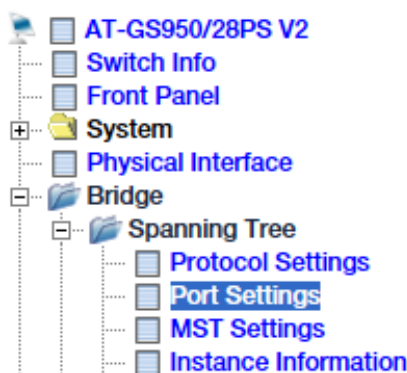


Figure 66. Port Settings Menu Selection

The Port Settings window is shown in Figure 67.

Port Settings											
Port	STP Status	Priority	Admin Cost (0 = Auto)	External Cost	Status	Edge	P2P	Restricted Role	Restricted TCN	Migrate	Apply
All	Ignore	Ignore		-	-	Ignore	Ignore	Ignore	Ignore	Restart	Apply
1	Enabled	128	0	20000	Disabled	Auto	Auto	False	False	Restart	Apply
2	Enabled	128	0	20000	Disabled	Auto	Auto	False	False	Restart	Apply
3	Enabled	128	0	20000	Disabled	Auto	Auto	False	False	Restart	Apply
4	Enabled	128	0	20000	Disabled	Auto	Auto	False	False	Restart	Apply
5	Enabled	128	0	20000	Disabled	Auto	Auto	False	False	Restart	Apply

Figure 67. Port Settings Window

2. Configure the settings in Table 27 on page 143.

Table 27. STP and RSTP Port Settings

Field	Definition
Port	<p>Displays the port numbers on the switch.</p> <p>You can use the All row at the top of the table as a shortcut to applying the same parameter settings to all the ports. The default Ignore parameter setting means to ignore the All row for that parameter.</p>
STP Status	<p>Enable or disable the spanning tree protocol on the port by selecting one of the following from the menu:</p> <ul style="list-style-type: none"> - Enabled: Activates the protocol on the port. This is the default setting. - Disabled: Deactivates the protocol.
Priority	<p>Select the port priority from the menu. This parameter is used as a tie breaker when two or more ports have equal costs to the root bridge. The range is 0 to 240 in increments of 16. The default value is 128.</p>
Admin Cost	<p>Enter the port cost. The spanning tree algorithm uses the cost to select the port that provides the lowest cost path to the root bridge when there are multiple physical paths. The range is 0 to 65,535. The default setting is 0 (Auto-detect), which sets port cost depending on port speed. Auto-Detect assigns a value of 100 to a port operating at 10Mbps, 10 for 100Mbps, and 4 for 1Gbps.</p>
External Cost	<p>Reserved for MSTP. Refer to “Configuring MSTP Port Settings” on page 168.</p>

Table 27. STP and RSTP Port Settings (Continued)

Field	Definition
State	<p>Displays the current port state. STP and RSTP have different port states. The STP states are listed here:</p> <ul style="list-style-type: none"> - Blocking - The port is blocking all traffic, except for BPDUs, because the protocol detected a network loop. The path is blocked because its cost to the root bridge is higher than another path. - Listening - The port is in the convergence state, prior to transitioning to the forwarding or blocking state. - Learning - The port is in the convergence state, learning source addresses from ingress frames, prior to transitioning to the forwarding or blocking state. - Forwarding - The port is forwarding network traffic. This indicates normal operation. STP continues monitoring the port for BPDUs that indicate whether the port should return to the blocking state to prevent a loop. - Disabled - The port does not have a link to a network device or the spanning tree protocol is disabled on the port. <p>The possible RSTP states are listed here:</p> <ul style="list-style-type: none"> - Discarding - The port is discarding all traffic, except for BPDUs, because a loop is detected. The path is blocked because its cost to the root bridge is higher than another path. - Forwarding - The port is forwarding network traffic. This indicates normal operation. RSTP continues to monitor the port for BPDUs that indicate whether the port should return to the blocking state to prevent a loop.

Table 27. STP and RSTP Port Settings (Continued)

Field	Definition
Edge	<p>Select a setting for the edge port parameter. Edge ports are connected to edge devices, such as computers or printers. Applies to RSTP only. Here are the settings:</p> <ul style="list-style-type: none"> - Auto: The switch automatically determines whether the port is an edge port. This is the default setting. The port is automatically designated an edge port and placed in the forwarding state if it does not receive any BPDUs for three seconds: - ForceTrue: Designates the port as an edge port. The port will always be in a forwarding state. - ForceFalse: Designates the port as not an edge port.
P2P	<p>Select a setting for the point-to-point port parameter. The settings are listed here:</p> <ul style="list-style-type: none"> - Auto: The switch automatically determines whether the port is a point-to-point port. This is the default setting. - ForceTrue: Designates the port as a point-to-point port. - ForceFalse: Designates the port as not a point-to-point port.
Restricted Role	Reserved for MSTP. Refer to “Configuring MSTP Port Settings” on page 168.
Restricted TNC	Reserved for MSTP. Refer to “Configuring MSTP Port Settings” on page 168.
Migrate	Click the button to reset the port so that it resumes sending RSTP BPDUs. An RSTP port that receives STP BPDUs changes to the STP mode and transmits STP BPDUs. You can use this button to return the port to the RSTP mode so that it transmits RSTP BPDUs again.

3. Click the **Apply** button to implement your changes.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 13

Multiple Spanning Tree Protocol Overview

This chapter provides background information on the Multiple Spanning Tree Protocol (MSTP). The sections are listed here:

- ❑ “Overview” on page 148
- ❑ “Multiple Spanning Tree Instance (MSTI)” on page 149
- ❑ “Multiple Spanning Tree Regions” on page 151
- ❑ “Common and Internal Spanning Tree (CIST)” on page 154
- ❑ “MSTP with STP and RSTP” on page 155
- ❑ “Summary of Guidelines” on page 156
- ❑ “Associating VLANs to MSTIs” on page 158
- ❑ “Connecting VLANs Across Different Regions” on page 160

Overview

MSTP searches for loops in the wiring topology of a network and, where loops exist, blocks bridge ports to prevent broadcast storms. This is the same function as that performed by STP and RSTP, as explained in Chapter 12, “Spanning Tree and Rapid Spanning Tree Protocols” on page 129. The main difference between MSTP and the other spanning tree protocols is that it lets you group the bridges of a network into multiple spanning tree domains. This can be useful in networks with large number of bridges because it enables the spanning tree protocol to react to and resolve loops more quickly than if all of the bridges are one domain.

The following sections describe some of the terms and concepts related to MSTP.

Note

Do not activate MSTP on the switch without first familiarizing yourself with the following concepts and guidelines. Unlike STP and RSTP, you cannot activate this spanning tree protocol without configuring the protocol parameters.

Note

The MSTP implementation on the switch complies with the IEEE 802.1s standard and is compatible with other vendors' compliant 802.1s implementations.

Multiple Spanning Tree Instance (MSTI)

The individual spanning trees domains in MSTP are referred to as Multiple Spanning Tree Instances (MSTIs). An MSTI can span any number of switches.

To create an MSTI, you assign it a number, referred to as the MSTI ID. The range is 1 to 31. (The switch is shipped with a default MSTI with an ID of 0. This default spanning tree instance is discussed later in “Common and Internal Spanning Tree (CIST)” on page 154.)

After selecting an MSTI ID, you need to define the scope of the MSTI by assigning one or more VLANs to it. An instance can contain any number of VLANs, but a VLAN can belong to only one MSTI at a time.

Here are the MSTI guidelines:

- ❑ The switch supports up to 16 spanning tree instances, including the CIST.
- ❑ An MSTI can contain any number of VLANs.
- ❑ A VLAN can belong to only one MSTI at a time.
- ❑ A switch port can belong to more than one spanning tree instance at a time by being an untagged and tagged member of VLANs belonging to different MSTIs. This is possible because a port can be in different MSTP states for different MSTIs simultaneously. For example, a port can be in the MSTP blocking state for one MSTI and the forwarding state for another spanning tree instance. For further information, refer to “Ports in Multiple MSTIs” on page 149.

VLAN and MSTI Associations

Part of the task to configuring MSTP involves assigning VLANs to spanning tree instances. The mapping of VLANs to MSTIs is called associations. A VLAN, either port-based or tagged, can belong to only one instance at a time, but an instance can contain any number of VLANs.

Ports in Multiple MSTIs

A port can be a member of more than one MSTI at a time if it is a tagged member of one or more VLANs assigned to different MSTIs. In this circumstance, a port might have to operate in different spanning tree states simultaneously, depending on the requirements of the MSTIs. For example, a port that belongs to two different VLANs in two different MSTIs might operate in the forwarding state in one MSTI and the blocking state in the other.

A port's MSTI parameter settings are divided into two groups. The first group is referred to as generic parameters. These are set only once on a port and apply to all the MSTIs where the port is a member. One of these parameters is the external path cost, which sets the operating cost of a port connected to a device outside its region. A port, even if it belongs to

multiple MSTIs, can have only one external path cost. Another generic parameter designates a port as an edge port or a point-to-point port.

The second group of port parameters can be set differently for each MSTI in which a port is a member. One parameter, the internal path cost, specifies the operating cost of a port when it is connected to a bridge in the same MSTP region. The other parameter in this group sets the port priority, which acts as a tie breaker when two or more ports have equal costs to a regional root bridge.

Multiple Spanning Tree Regions

Another important concept of MSTP is regions. An MSTP region is defined as a group of bridges that share exactly the same MSTI characteristics. The characteristics are listed here:

- Configuration name
- Revision number
- VLANs
- VLAN to MSTI ID associations

A configuration name is a name that identifies a region. You must assign each bridge in a region exactly the same name; even the same upper and lowercase lettering. Identifying the regions in your network is easier if you choose names that are characteristic of the functions of the nodes and bridges of the region. Examples are Sales Region and Engineering Region.

The revision number is an arbitrary number assigned to a region. You might use this number to keep track of the revision level of a region's configuration. For example, you might use this value to maintain the number of times you revise a particular MSTP region. It is not important that you maintain this number, only that all of the bridges in a region have the same number.

The bridges of a particular region must also have the same VLANs. The names of the VLANs and the VIDs must be same on all of the bridges in a region.

Finally, the VLANs in the bridges must be associated to the same MSTIs.

If any of the above information is different on two bridges, MSTP considers the bridges as residing in different regions.

Figure 68 on page 152 illustrates the concept of regions. It shows one MSTP region with two switches. The switches have the same configuration names and revision levels. They also have the same five VLANs and the VLANs are associated with the same MSTIs.

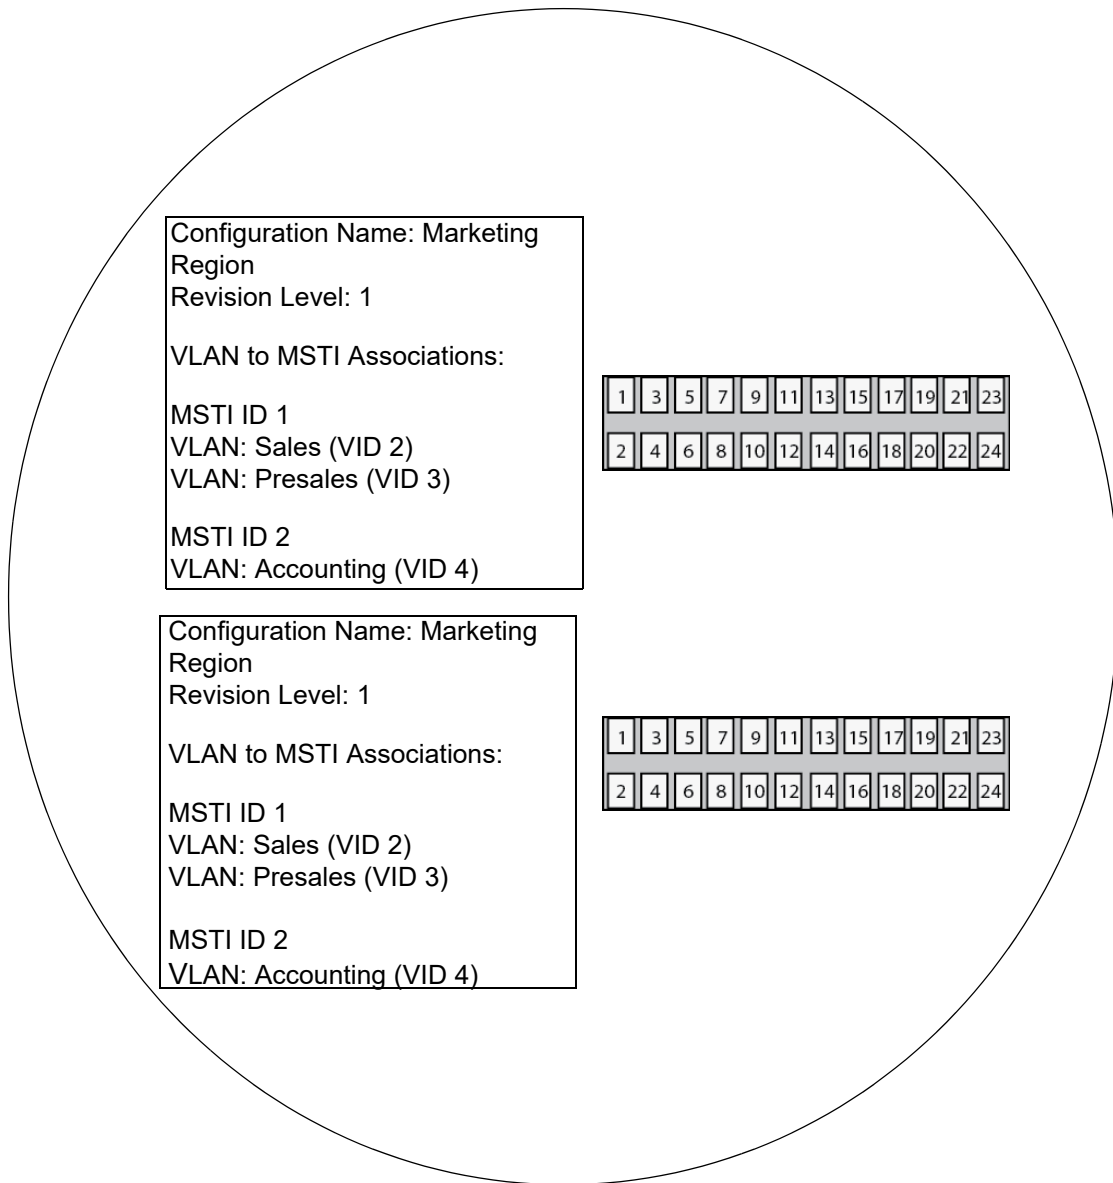


Figure 68. Multiple Spanning Tree Region

The switch determines regional boundaries by examining the MSTP BPDUs it receives on the ports. A port that receives an MSTP BPDU from another bridge with regional information different from its own is considered to be a boundary port and the bridge connected to the port as belonging to another region.

The same is true for ports connected to bridges running STP or RSTP. Those ports are also considered as part of another region.

Each MSTI functions as an independent spanning tree within a region. Consequently, each MSTI must have a root bridge to locate physical loops

within the spanning tree instance. An MSTI's root bridge is called a regional root. The MSTIs within a region may share the same regional root or they can have different regional roots.

A regional root of an MSTI must be within the region where the MSTI is located. An MSTI cannot have a regional root that is outside its region.

A regional root is selected by a combination of the MSTI priority value and the bridge's MAC address. The MSTI priority is analogous to the RSTP bridge priority value. Where they differ is that while the RSTP bridge priority is used to determine the root bridge for an entire bridged network, MSTI priority is used to determine the regional root of a particular MSTI.

The range for this parameter is the same as the RSTP bridge priority; from 0 to 61,440 in sixteen increments of 4,096.

Region Guidelines

Here are the guidelines for regions.

- A network can contain any number of regions and a region can contain any number of switches.
- A switch can belong to only one region at a time.
- A region can contain any number of VLANs.
- All of the bridges in a region must have the same configuration name, revision level, VLANs, and VLAN to MSTI associations.
- An MSTI cannot span multiple regions.
- Each MSTI must have a regional root for locating loops in the instance. MSTIs can share the same regional root or have different roots. A regional root is determined by the MSTI priority value and a bridge's MAC address.
- The regional root of an MSTI must be in the same region as the MSTI.

Common and Internal Spanning Tree (CIST)

MSTP has a default spanning tree instance called the Common and Internal Spanning Tree (CIST). This instance has an MSTI ID of 0.

This instance has unique features and functions that make it different from the MSTIs you create yourself. First, you cannot delete this instance or change its MSTI ID. Second, when you create a new port-based or tagged VLAN, it is by default associated with the CIST and is automatically given an MSTI ID of 0. The default VLAN is also associated by default with CIST.

Another important difference is that when you assign a VLAN to another MSTI, it still partially remains a member of CIST. This is because CIST is used by MSTP to communicate with other MSTP regions and with any RSTP and STP bridges in a network. MSTP uses CIST to participate in the creation of a spanning tree between different regions and between regions and STP and RSTP bridges, to form one spanning tree for the entire bridged network.

MSTP uses CIST to form the spanning tree of an entire bridged network because CIST can cross regional boundaries, while an MSTI cannot. If a port is a boundary port, that is, if it is connected to another region, that port automatically belongs solely to CIST, even if it was assigned to an MSTI, because only CIST is active outside of a region.

As mentioned earlier, every MSTI must have a root bridge, referred to as a regional root, in order to locate loops that might exist within the instance. CIST must also have a regional root. However, the CIST regional root communicates with the other MSTP regions and STP and RSTP bridges in the bridged network.

The CIST regional root is set with the CIST Priority parameter. This parameter, which functions similar to the RSTP bridge priority value, selects the root bridge for the entire bridged network. If the switch has the lowest CIST Priority value among all the spanning tree bridges, it functions as the root bridge for all the MSTP regions and STP and RSTP bridges in the network.

MSTP with STP and RSTP

MSTP is fully compatible with STP and RSTP. If a port on the switch running MSTP receives STP BPDUs, the port sends only STP BPDU packets. If a port receives RSTP BPDUs, the port sends MSTP BPDUs because RSTP can process MSTP BPDUs.

A port connected to a bridge running STP or RSTP is considered to be a boundary port of the MSTP region and the bridge as belonging to a different region.

An MSTP region can be considered as a virtual bridge. The implication is that other MSTP regions and STP and RSTP single-instance spanning trees cannot discern the topology or constitution of an MSTP region. The only bridge they are aware of is the regional root of the CIST instance.

Summary of Guidelines

Careful planning is essential for the successful implementation of MSTP. This section reviews all of the rules and guidelines mentioned in earlier sections, and provides a few new ones:

- ❑ The switch can support up to 16 spanning tree instances, including the CIST.
- ❑ An MSTI can contain any number of VLANs.
- ❑ A VLAN can belong to only one MSTI at a time.
- ❑ An MSTI ID can be from 1 to 31.
- ❑ The CIST ID is 0. You cannot change this value.
- ❑ A switch port can belong to more than one spanning tree instance at a time. This allows you to assign a port as an untagged and tagged member of VLANs that belong to different MSTIs. What makes this possible is a port's ability to be in different MSTP states for different MSTIs simultaneously. For example, a port can be in the MSTP blocking state for one MSTI and the forwarding state for another spanning tree instance.
- ❑ A router or Layer 3 network device is required to forward traffic between VLANs.
- ❑ A network can contain any number of regions and a region can contain any number of switches.
- ❑ The switch can belong to only one region at a time.
- ❑ A region can contain any number of VLANs.
- ❑ All of the bridges in a region must have the same configuration name, revision level, VLANs, and VLAN to MSTI associations.
- ❑ An MSTI cannot span multiple regions.
- ❑ Each MSTI must have a regional root for locating loops in the instance. MSTIs can share the same regional root or have different roots. A regional root is determined by the MSTI priority value and a bridge's MAC address.
- ❑ The regional root of an MSTI must be in the same region as the MSTI.
- ❑ The CIST must have a regional root for communicating with other regions and single-instance spanning trees.
- ❑ MSTP is compatible with STP and RSTP.
- ❑ A port transmits CIST information even when it is associated with another MSTI ID. However, in determining network loops, MSTI takes precedence over CIST. (This is explained more in "Associating VLANs to MSTIs" on page 158.)

Note

The MSTP implementation on the switch complies with the IEEE 802.1s standard and is compatible with similar products from other vendors, provided that their products are also compliant with the standard.

Associating VLANs to MSTIs

Allied Telesis recommends that you assign all of the VLANs on the switch, including the default VLAN, to an MSTI. You should not leave VLANs assigned to only the CIST. This is to prevent the switch from blocking ports that should be in the forwarding state. The reason for this guideline is explained here.

An MSTP BPDUs contains the instance to which the port transmitting the packet belongs. By default, all of the ports belong to the CIST instance. So CIST is included in the BPDUs. If a port is a member of a VLAN that has been assigned to another MSTI, that information is also included in the BPDUs.

This is illustrated in Figure 69. Port 8 in switch A is a member of a VLAN assigned to MSTI ID 7 while port 1 is a member of a VLAN assigned to MSTI ID 10. The BPDUs transmitted by port 8 to switch B indicate that the port is a member of both CIST and MSTI 7, while the BPDUs from port 1 indicates the port is a member of the CIST and MSTI 10.

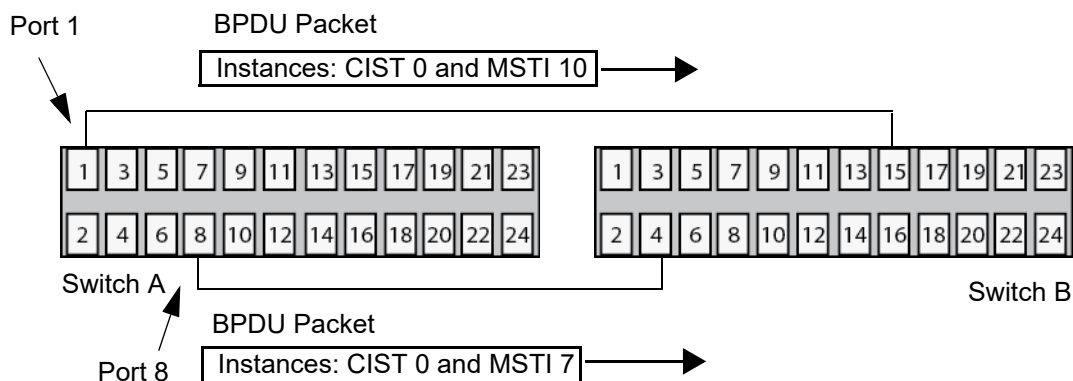


Figure 69. CIST and VLAN Guideline - Example 1

At first glance, it might appear that because both ports belong to CIST, a loop would exist between the switches and that MSTP would block a port to stop the loop. However, within a region, MSTI takes precedence over CIST. When switch B receives a packet from switch A, it uses MSTI, not CIST, to determine whether a loop exists. And because both ports on switch A belong to different MSTIs, switch B determines that no loop exists.

A problem can arise if you assign some VLANs to MSTIs while leaving others only in CIST. The problem is illustrated in Figure 70 on page 159. The network is the same as the previous example. The difference is that the VLAN containing port 8 on Switch A has not been assigned to an MSTI, and belongs only to CIST with its MSTI ID 0.

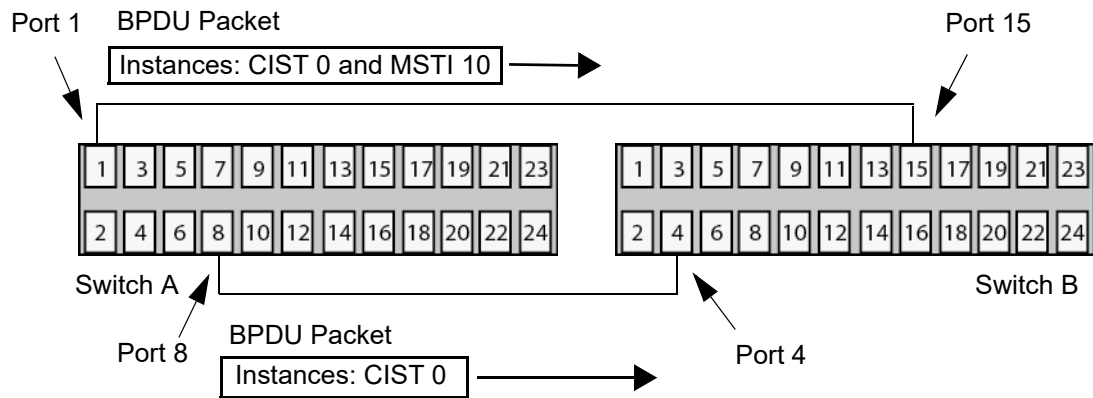


Figure 70. CIST and VLAN Guideline - Example 2

When port 4 on switch B receives a BPDU, the switch notes the port sending the packet belongs only to CIST. Therefore, switch B uses CIST to determine whether a loop exists. The result would be that the switch detects a loop because the other port is also receiving BPDU packets from CIST 0. Switch B would block a port to block the loop.

To avoid this issue, always assign all of the VLANs on the switch, including the Default VLAN, to MSTIs. This guarantees that all of the ports on the switch have an MSTI ID and ensures that loop detection is based on the MSTIs and not CIST.

Connecting VLANs Across Different Regions

Special consideration needs to be taken into account when connecting different MSTP regions or an MSTP region and an STP or RSTP region. Unless planned properly, VLAN fragmentation can occur between the VLANs of the network.

As mentioned previously, only the CIST can span regions. Consequently, you may run into problems if you use more than one physical data link to connect together various parts of VLANs that reside in bridges in different regions. The result can be a physical loop, which spanning tree disables by blocking ports.

This is illustrated in Figure 71. The example shows two switches that reside in different regions. Port 1 in switch A is a boundary port. It is an untagged member of the Accounting VLAN, which has been associated with MSTI 4. Port 16 is a tagged and untagged member of three different VLANs, all associated to MSTI 12.

If both switches were a part of the same region, there would be no problem because the ports reside in different spanning tree instances. However, the switches are part of different regions and MSTIs do not cross regions. Consequently, the result is that spanning tree would determine that a loop exists between the regions, and Switch B would block a port.

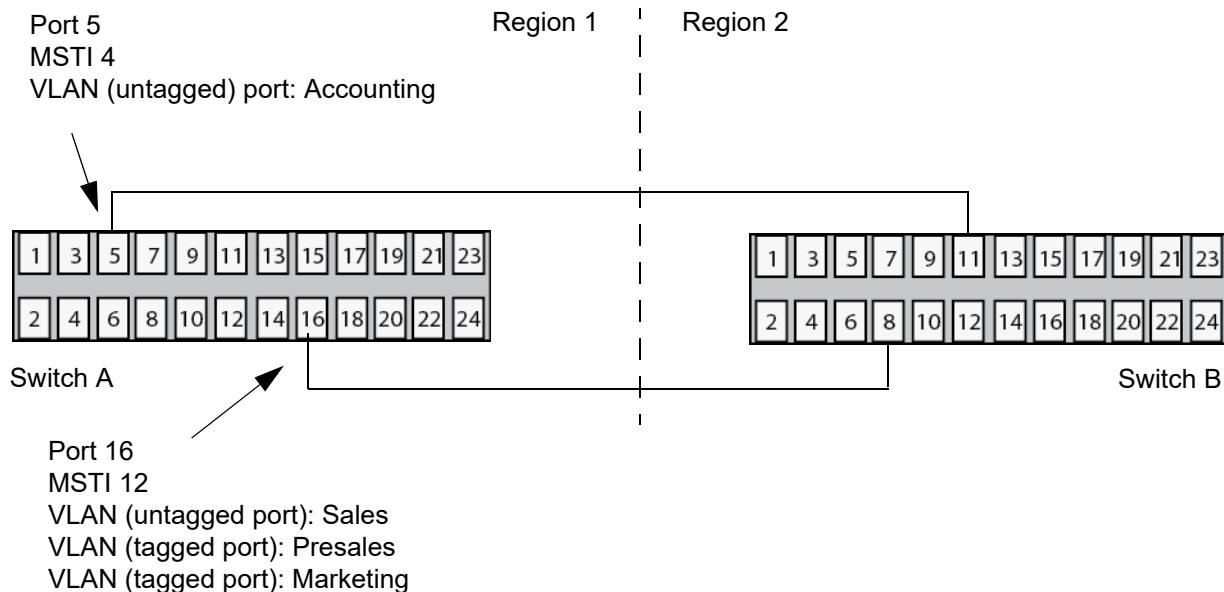


Figure 71. Spanning Regions

There are several ways to address this issue. One way is to have only one MSTP region for each subnet in your network.

Another approach is to group those VLANs that need to span regions into the same MSTI. Those VLANs that do not span regions can be assigned to other MSTIs.

Here is an example. Assume that you have two regions that contain the following VLANs:

Region 1 VLANs

Sales
Presales
Marketing
Advertising
Technical Support
Product Management
Project Management
Accounting

Region 2 VLANs

Hardware Engineering
Software Engineering
Technical Support
Product Management
CAD Development
Accounting

The two regions share three VLANs: Technical Support, Product Management, and Accounting. You could group those VLANs into the same MSTI in each region. For instance, for Region 1 you might group the three VLANs in MSTI 11 and in Region 2 you could group them into MSTI 6. After they are grouped, you can connect the VLANs across the regions using a link of tagged ports.

Chapter 14

Multiple Spanning Tree Protocol

This chapter contains the configuration procedures for the Multiple Spanning Tree Protocol (MSTP) in the following sections:

- ❑ “Configuring MSTP Global Settings” on page 164
- ❑ “Configuring MSTP Port Settings” on page 168
- ❑ “Configuring MST and MSTI Settings” on page 172
- ❑ “Configuring MST Port Settings” on page 175
- ❑ “Displaying MST Instance Information” on page 178

Configuring MSTP Global Settings

Perform the following procedure to enable or disable the spanning tree protocol, choose MSTP as the active protocol, and configure global switch (non-port) settings:

Note

You have to disable BPDU pass-through on all ports before you can enable the spanning tree protocol. Refer to “Configuring Basic Port Settings” on page 51.

Note

The switch briefly stops forwarding Ethernet traffic when you enable or disable the spanning tree protocol.

1. Select **Bridge > Spanning Tree > Protocol Settings** from the main menu. Refer to Figure 72.

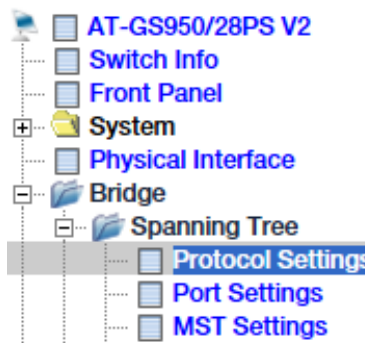


Figure 72. Protocol Settings Menu Selection

The Protocol Settings window is shown in Figure 65 on page 139.

2. Configure the settings in Table 28:

Table 28. Spanning Tree Protocol Settings Window for MSTP

Field	Definition
Global STP Status	Select one of the following from the menu: - Enabled: Activates the spanning tree protocol. You have to enable the protocol to configure the settings. - Disabled: Deactivates the spanning tree protocol. This is the default setting.

Table 28. Spanning Tree Protocol Settings Window for MSTP (Continued)

Field	Definition
Protocol Version	Select MSTP from the menu:
Bridge Priority	<p>Select the priority number for the bridge. This number is used to determine the root bridge the regional root for a particular MSTI. The bridge with the lowest priority number is selected as the root bridge. If two or more bridges have the same priority value, the bridge with the numerically lowest MAC address becomes the root bridge. If a root bridge goes off-line, the bridge with the next lowest priority number automatically takes over as the root bridge. The range is 0 (zero) to 61,440 in increments of 4096, with 0 as the highest priority.</p>
Maximum Age	<p>Enter the maximum length of time the bridge stores bridge protocol data units (BPDUs). Bridges in a bridged LAN use the aging time to test the age of stored configuration messages called bridge protocol data units (BPDUs) and to delete expired data units. For example, when set to the default value 20 seconds, bridges delete configuration messages after 20 seconds. The range is 6 to 40 seconds.</p> <p>Observe the following rules when setting the maximum age:</p> <ul style="list-style-type: none"> - Maximum Age must be greater than (2 x (Hello Time + 1)) - Maximum Age must be less than (2 x (Forward Delay - 1)) - The aging time for BPDUs is different from the aging time for the MAC address table. <p>MSTP uses this parameter when interacting with STP/RSTP domains on boundary ports.</p>

Table 28. Spanning Tree Protocol Settings Window for MSTP (Continued)

Field	Definition
Hello Time	Enter the time interval at which the root bridge transmits BPDUs to determine whether there have been any changes to the network topology and to inform other bridges of topology changes. This parameter is active only on the root bridge. The range is 1 to 10 seconds. The default is 2 seconds.
Forward Delay	Enter the waiting period in seconds before a bridge changes to a new state, for example, becomes the new root bridge after a topology change. If the bridge transitions too soon, not all links may have yet adapted to the change, resulting in network loops. The range is 4 to 30 seconds. The default is 15 seconds.
Transmit Hold Count	Enter the maximum number of BPDUs the bridge can send per second. The range is 1 to 10 BPDUs. The default is 6 BPDUs.
Max Hop Count	Enter the maximum number of hops before BPDUs are deleted. The Max Hop counter in a BPDU is decremented every time a BPDU crosses an MSTP region boundary. BPDUs are deleted after the counter reaches zero.
Root Bridge	Displays the MAC address of the root bridge of the spanning tree domain. This value is zero when STP is disabled on the switch.
Root Cost	Displays the sum of the root port costs of the bridges between the switch's root port and the root bridge.
Root Maximum Age	Displays the maximum age setting on the root bridge.
Root Forward Delay	Displays the forward delay setting on the root bridge.
Root Port	Displays the port on the switch that leads to the root bridge. This value will be zero if the switch is the root bridge or spanning tree is disabled.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Configuring MSTP Port Settings

Perform the following procedure to configure MSTP port settings:

1. Select **Bridge > Spanning Tree > Port Settings**. Refer to Figure 73.

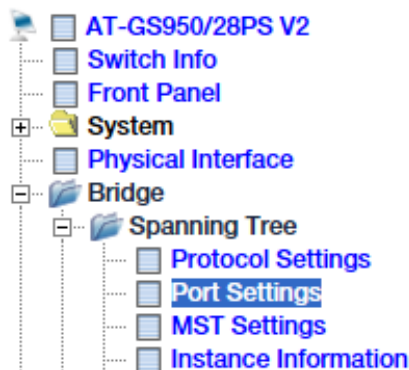


Figure 73. Spanning Tree Port Settings Menu Selection

The Port Settings window is shown in Figure 67 on page 142.

2. Configure the settings in Table 29.

Table 29. MSTP Port Settings

Field	Definition
Port	Displays the port numbers on the switch. You can use the All row at the top of the table as a shortcut to applying the same parameter settings to all the ports. The default Ignore parameter setting means to ignore the All row for that parameter.
STP Status	Enable or disable the spanning tree protocol on the port by selecting one of the following from the menu: - Enabled: Activates the protocol on the port. This is the default setting. - Disabled: Deactivates the protocol.
Priority	Select the port priority from the menu. This parameter is used as a tie breaker when two or more ports have equal costs to the root bridge. The range is 0 to 240 in increments of 16. The default value is 128.

Table 29. MSTP Port Settings (Continued)

Field	Definition
Admin Cost	Enter the port cost. The spanning tree algorithm uses the cost to select the port that provides the lowest cost path to the root bridge when there are multiple physical paths. The range is 0 to 65,535. The default setting is 0 (Auto-detect), which sets port cost depending on port speed. Auto-Detect assigns a value of 100 to a port operating at 10Mbps, 10 for 100Mbps, and 4 for 1Gbps.
External Cost	Displays the operating cost of the port when connected to a device outside its region.
State	<p>Displays the current MSTP port state, listed here:</p> <ul style="list-style-type: none"> - Blocking - The port is blocking all traffic, except for BPDUs, because the protocol detected a network loop. The path is blocked because its cost to the root bridge is higher than another path. - Listening: The port is in the convergence state, prior to transitioning to the forwarding or blocking state. - Learning - The port is in the convergence state, learning source addresses from ingress frames, prior to transitioning to the forwarding or blocking state. - Forwarding - The port is forwarding network traffic. This indicates normal operation. STP continues monitoring the port for BPDUs that indicate whether the port should return to the blocking state to prevent a loop. - Disabled - The port does not have a link to a network device or the spanning tree protocol is disabled on the port.

Table 29. MSTP Port Settings (Continued)

Field	Definition
Edge	<p>Select a setting for the edge port parameter. Edge ports are connected to edge devices, such as computers or printers. Applies to RSTP only. Here are the settings:</p> <ul style="list-style-type: none"> - Auto: The switch automatically determines whether the port is an edge port. This is the default setting. The port is automatically designated an edge port and placed in the forwarding state if it does not receive any BPDUs for three seconds: - ForceTrue: Designates the port as an edge port. The port will always be in a forwarding state. - ForceFalse: Designates the port as not an edge port.
P2P	<p>Select a setting for the point-to-point port parameter from the menu:</p> <ul style="list-style-type: none"> - Auto: The switch automatically determines whether the port is a point-to-point port. This is the default setting. - ForceTrue: Designates the port as a point-to-point port. - ForceFalse: Designates the port as not a point-to-point port.
Restricted Role	<p>Select the restricted role of the port from the menu:</p> <ul style="list-style-type: none"> - True: Blocks the port from being a root port or from communicating with the root bridge. - False: Permits the port to become a root port. This is the default setting <p>The net effect of setting all ports on the switch to True is that it forces the switch into the role of the root bridge regardless of other path costs in the network.</p>

Table 29. MSTP Port Settings (Continued)

Field	Definition
Restricted TNC	<p>Select the Restricted TCN setting from the menu:</p> <ul style="list-style-type: none"> - True: Blocks the port from transmitting Topology Change Notification (TCN) BPDUs when there is a topology change. - False: Allows the port to transmit TCN BPDUs when there is a topology change. This is the default setting.
Migrate	<p>Click the button to reset the port so that it resumes sending MSTP BPDUs. MSTP ports that receive STP BPDUs change to the STP mode and transmit STP BPDUs. You can use this button to return ports to the MSTP mode so that they transmit MSTP BPDUs again.</p>

3. Click the **Apply** button in the **Action** column.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Configuring MST and MSTI Settings

Perform the following procedure to configure MST settings and manage MST instances:

1. Select **Bridge > Spanning Tree > MST Settings** from the main menu. Refer to Figure 74.

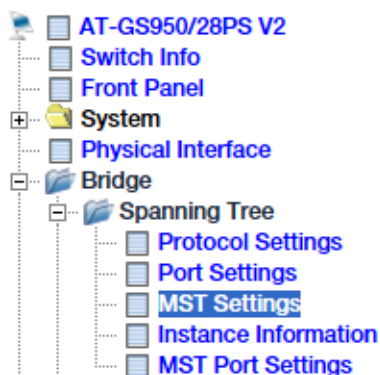


Figure 74. MST Settings Menu Selection

. The MST Settings window is shown in Figure 75.

MST Settings

MST Configuration Identification Settings

Configuration Name:

Revision Level: (0-65535)

MST Instance Settings

MSTI ID: *(1-31)

VID List: (1-4093)

Priority:

MSTI ID	VID List	Priority	Action
CIST	1-4093	<input type="text" value="32768"/>	<input type="button" value="Apply"/> <input type="button" value="Delete"/>

Figure 75. MST Settings Window

2. To configure the settings in the MST Configuration Identification Settings section of the window, do the following:
 - a. Configure the settings in Table 30.

Table 30. MST Settings Window - MST Configuration Identification Settings

Field	Description
Configuration Name	Enter the name of the region where the bridge is a member. All the switches in an MSTP region need to have the same name.
Revision Level	Enter the region's revision level. All the switches in an MSTP region have to have the same revision level.

- b. Click the **Apply** button.
 - c. Go to step 6.
3. To map VLANs to MST instances, do the following:
 - a. Configure the settings in Table 31.

Table 31. MST Settings Window - MST Instance Settings

Field	Description
MSTI ID	Enter a unique ID for the new MST instance. The range is 1 to 31.
VID List	Enter VIDs of the VLANs to be associated with the MST instance.
Priority	Choose a priority value from the menu. The switch uses the priority to select the MSTI regional root. The range is 0 (zero) to 61,440 in increments of 4,096. The highest priority is 0. The default is 32768.

- b. Click the **Add** button. The new MSTI is added to the table.
 - c. Go to step 6.

4. To modify an MSTI, refer to Table 31 on page 173.
5. To delete an MSTI, click its **Delete** button in the **Action** column. The instance and its mapped VLAN associations are deleted.
6. Select **Save Settings to Flash** from the main menu to save your changes.

Configuring MST Port Settings

Perform the following procedure to configure MST port settings:

1. Select **Bridge > Spanning Tree > MST Port Settings** from the main menu. Refer to Figure 76.

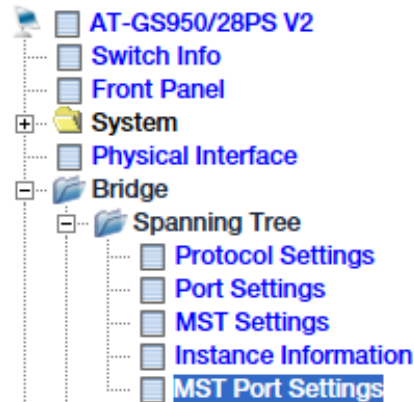


Figure 76. MST Port Settings

The MST Port Settings window is shown in Figure 77.

MST Port Settings

Select MST Port:

MSTI ID	Designated Bridge	Internal Path Cost	Admin Path Cost (0 = Auto)	Priority	Status	Role	Action
CIST	00:00:00:00:00:00:00	20000	0	<input type="text" value="128"/>	Disabled	Disabled	<input type="button" value="Apply"/>

Figure 77. MST Port Settings Page

2. Select a port to configure from the **Select MST Port** pull-down menu. You can configure only one port at a time.
3. Configure the settings in Table 32.

Table 32. MST Port Settings Window

Field	Description
MSTI ID	Displays the MSTP Instance of the port.
Designated Bridge	Displays the MAC address of the bridge that provides the least-cost path to the root bridge from a network segment.
Internal Path Cost	Displays the operating cost of the port when it is connected to a bridge in the same MSTP region.

Table 32. MST Port Settings Window (Continued)

Field	Description
Admin Path Cost	<p>Enter the cost of a port to the root. The range is 0 to 200,000,000. The 0 value, the default setting, activates the Auto setting, which sets the value according to port speed. Here are the MSTP port costs with the Auto setting when a port is not a member of a trunk.</p> <ul style="list-style-type: none"> - 10 Mbps - 2,000,000 - 100 Mbps - 200,000 - 1000 Mbps - 20,000 <p>Here are the MSTP port costs with the Auto setting when a port is a member of a trunk.</p> <ul style="list-style-type: none"> - 10 Mbps - 20,000 - 100 Mbps - 20,000 - 1000 Mbps - 2,000
Priority	<p>Select the spanning tree port priority from the menu. The priority is used as a tie breaker when two or more ports have equal costs to the regional root bridge. The range is 0 to 255 in increments of 16. The default value is 128.</p>
Status	<p>Displays the MSTP state of the port. The possible states are listed here:</p> <ul style="list-style-type: none"> - Disabled: The port has not established a link with a network device, or MSTP is disabled on the port or switch. - Discarding: The port is discarding received packets and is not submitting forwarded packets for transmission. - Learning: The port can receive but not forward packets. - Forwarding: Normal operations.

Table 32. MST Port Settings Window (Continued)

Field	Description
Role	<p>Displays the MSTP role of the port. The possible roles are listed here:</p> <ul style="list-style-type: none"> - Disabled: The port has not established a link with a network device, or MSTP is disabled on the port or switch. - Root - The port that is connected to the root switch, directly or through other switches, with the least path cost. - Alternate - The port offers an alternate path in the direction of the root switch. - Backup - The port on a designated switch that provides a backup for the path provided by the designated port. - Designated - The port on the designated switch for a LAN that has the least cost path to the root switch. This port connects the LAN to the root switch. - Master - Similar to the root port. When the port is a boundary port, the MSTI port roles follow the CIST port roles. The MSTI port role is called "master" when the CIST role is "root."

4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Displaying MST Instance Information

To display MST instance information, select **Bridge -> Spanning Tree -> Instance Information** from the main menu. Refer to Figure 78.

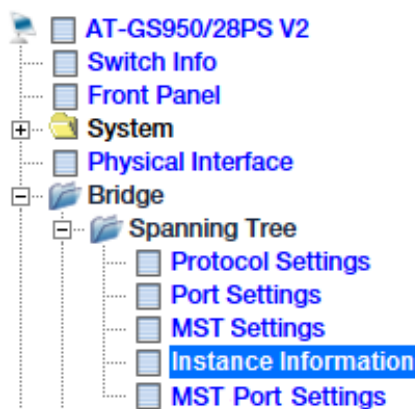


Figure 78. Instance Information Menu Selection

The Instance Information window is shown in Figure 79.

<i>Instance Information</i>					
MSTI ID	Internal Root Cost	Root Port	Regional Root Bridge	Designated Bridge	Instance Priority
CIST	0	0	00:00:00:00:00:00:00:00	00:00:00:00:00:00:00:00	32768

Figure 79. Instance Information Window

The columns are described in Table 33.

Table 33. Instance Information Window

Column	Description
MSTI ID	Displays the MSTP Instance.
Internal Root Cost	Displays the internal cost of the port leading to the root bridge in the MSTP region.
Root Port	Displays the port leading to the root bridge in the MSTP region.
Regional Root Bridge	Displays the MAC address of the root bridge of the MST instance.
Designated Bridge	Displays the MAC address of the bridge providing the least-cost path to the root bridge.

Table 33. Instance Information Window (Continued)

Column	Description
Instance Priority	Displays the priority value of the port leading to the root bridge.

Section III

Virtual LANs

This section contains the following chapters:

- ❑ Chapter 15, “Port-Based Virtual LANs” on page 183
- ❑ Chapter 16, “802.1Q Tagged Virtual LANs” on page 205
- ❑ Chapter 17, “GARP VLAN Registration Protocol” on page 225
- ❑ Chapter 18, “Private Virtual LAN” on page 235
- ❑ Chapter 19, “Voice VLAN” on page 241
- ❑ Chapter 20, “MAC Address Forwarding Table” on page 251

Chapter 15

Port-Based Virtual LANs

This chapter describes port-based Virtual Local Area Networks (VLANs) in the following sections:

- ❑ “VLAN Overview” on page 184
- ❑ “Port-Based VLAN Overview” on page 186
- ❑ “Adding or Viewing Port-Based VLANs” on page 190
- ❑ “Assigning the Port-Based VLAN Mode to Ports” on page 193
- ❑ “Modifying Port-Based VLANs” on page 196
- ❑ “Deleting Port-Based VLANs” on page 198
- ❑ “Selecting the MAC Address Forwarding Table Mode” on page 200
- ❑ “Viewing the MAC Addresses of Network Devices in the Forwarding Table” on page 202

VLAN Overview

A virtual LAN (VLAN) is a group of ports on an Ethernet switch that form a logical Ethernet segment. The ports of a VLAN form an independent traffic domain where the traffic generated by the nodes remain within the VLAN.

VLANs are commonly used to segment local area networks or group nodes with related functions into separate, logical, VLAN segments. These VLAN groupings can be based on similar data needs or security requirements. For example, you might add VLANs for the various department in your company, such as Sales, Accounting, and Engineering.

VLANs offer several benefits:

- ❑ Improved network performance

Network performance often suffers as networks grow in size and as data traffic increases. The more nodes on each LAN segment vying for bandwidth, the greater the likelihood overall network performance will decrease.

VLANs improve network performance because traffic stays within the separate, logical LAN segments of the VLANs. The nodes of a VLAN receive traffic only from nodes of the same VLAN. This reduces the need for nodes to handle traffic that is not destined for them and frees up bandwidth within all the logical work groups.

In addition, because each VLAN constitutes a separate broadcast domain, broadcast traffic remains within its VLAN and are not shared with other ports of the switch that are not members of that VLAN. This improves network performance for non-member ports.

- ❑ Increased security

Because data traffic generated by nodes in a VLAN are restricted only to nodes in the same VLAN, you can use VLANs to control the flow of packets in your network and prevent packets from being shared with unauthorized end nodes.

- ❑ Simplified network management

VLANs can simplify network management. Without VLANs, physical changes to the network have to been made at the switches in the wiring closets. For example, if an employee changed departments, changing the employee's LAN segment assignment might require changing the cabling of the switches.

With VLANS, you can reconfigure the LAN segment assignments of end nodes with the switch's management software. Also, you can change the VLAN memberships without moving the workstations

physically and change group memberships without moving cables from one port to another.

Virtual LANs can span more than one switch. End nodes of a VLAN do not need to be connected to the same switch and so are not restricted to being in the same physical location.

The GS950 PS V2 Switch supports the following types of VLANs:

- Port-Based VLANs
- Tagged VLANs
- Private VLAN
- Voice VLAN

This chapter describes Port-Based VLANs.

Port-Based VLAN Overview

As explained in the “VLAN Overview” on page 184, a VLAN consists of a group of ports that form an independent traffic domain on an Ethernet switch. The ports of a VLAN forward ingress traffic only to those ports that are part of the same VLAN. An interconnection device, such as a router or Layer 3 switch, is commonly added so that traffic can crossover to ports in different VLANs.

The switch supports several types of VLANs. This chapter described Port-Based VLANs. To add Port-Based VLANs to the switch, you simply identify the member ports and the switch automatically groups them together into VLANs. The switch supports up to 52 port-based VLANs at a time. Here are the components of Port-Based VLANs:

- VLAN index number
- VLAN name
- Ports
- Port mode

The components are described Table 34.

Table 34. Components of Port-Based VLANs

Item	Description
VLAN Index Number	A VLAN has to have a unique index number so that the switch can identify it. The range is 1 to 52.
Name	A VLAN has to have a name. VLANs will be easier to identify if you give them names that reflect the functions of the member devices, such as Sales, Production, or Engineering.
Ports	<p>A VLAN has to have at least one port. Here are the guidelines:</p> <ul style="list-style-type: none"> - A port-based VLAN can have any number of ports, up to all the ports on the switch. - Ports cannot be members of Port-Based and 802.1Q tagged, voice, or private VLANs at the same time. - Ports can belong to more than one Port-Based VLAN at a time. For example, if one VLAN contains ports 1 to 4 and another VLAN on the same switch has ports 4 to 9, port 4 would be a member of both VLANs.

Table 34. Components of Port-Based VLANs (Continued)

Item	Description
Mode	This component designates the VLAN type. Ports in this type of VLAN have to be assigned the Port-Based VLAN mode. You have to add ports to VLANs before you can set their VLAN modes.

Port-Based Example 1

Figure 80 illustrates an example of one switch with three Port-Based VLANs. (The default VLAN is not shown in the following examples.)

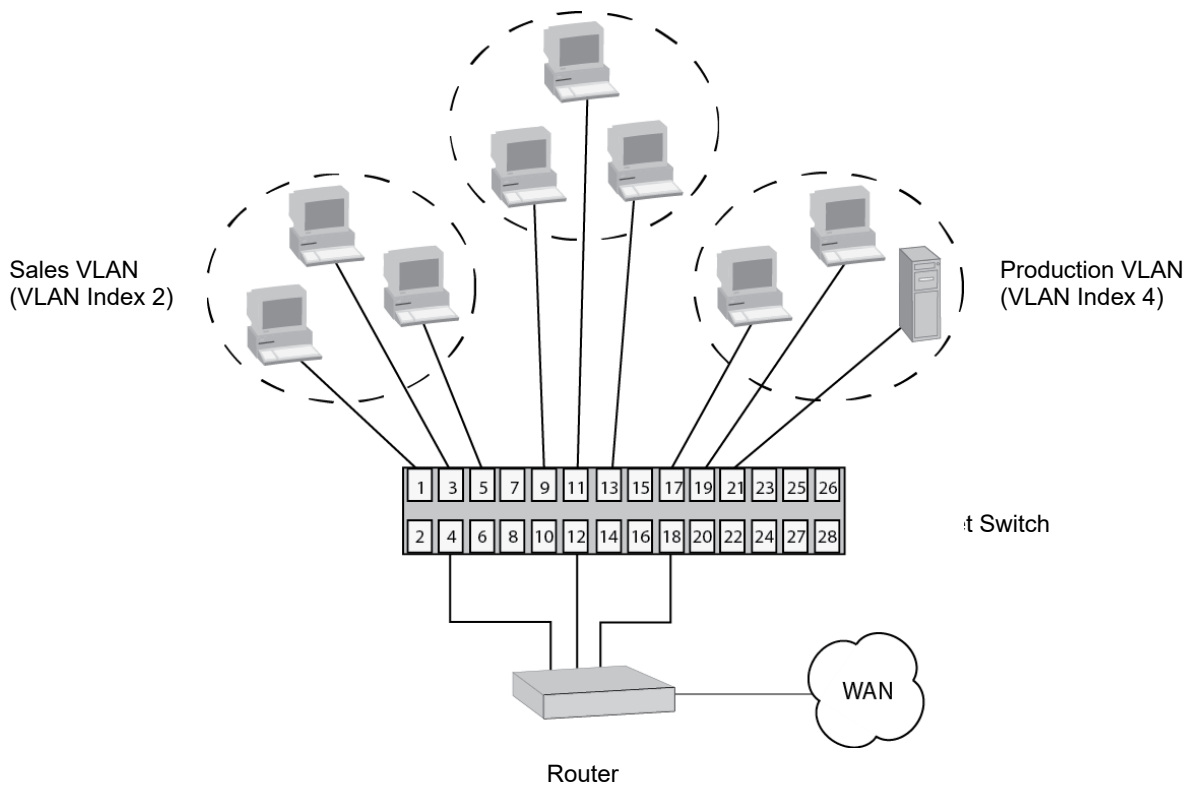


Figure 80. Port-Based VLAN - Example 1

Table 35 lists the port assignments of the VLANs.

Table 35. Example 1 of Port-Based VLANs

	Sales (VLAN Index 2)	Engineering (VLAN Index 3)	Production (VID Index 4)
Ethernet Switch	Ports 1, 3 - 5	Ports 9, 11 - 13	Ports 17 - 19, 21

The VLANs have unique VLAN index numbers that are assigned when you add the VLANs to the switch. Each VLAN has one port connected to the router, which interconnects the VLANs and acts as a gateway to the WAN.

Port-Based Example 2

The example in Figure 81 has two VLANs, Sales and Engineering, that span two switches.

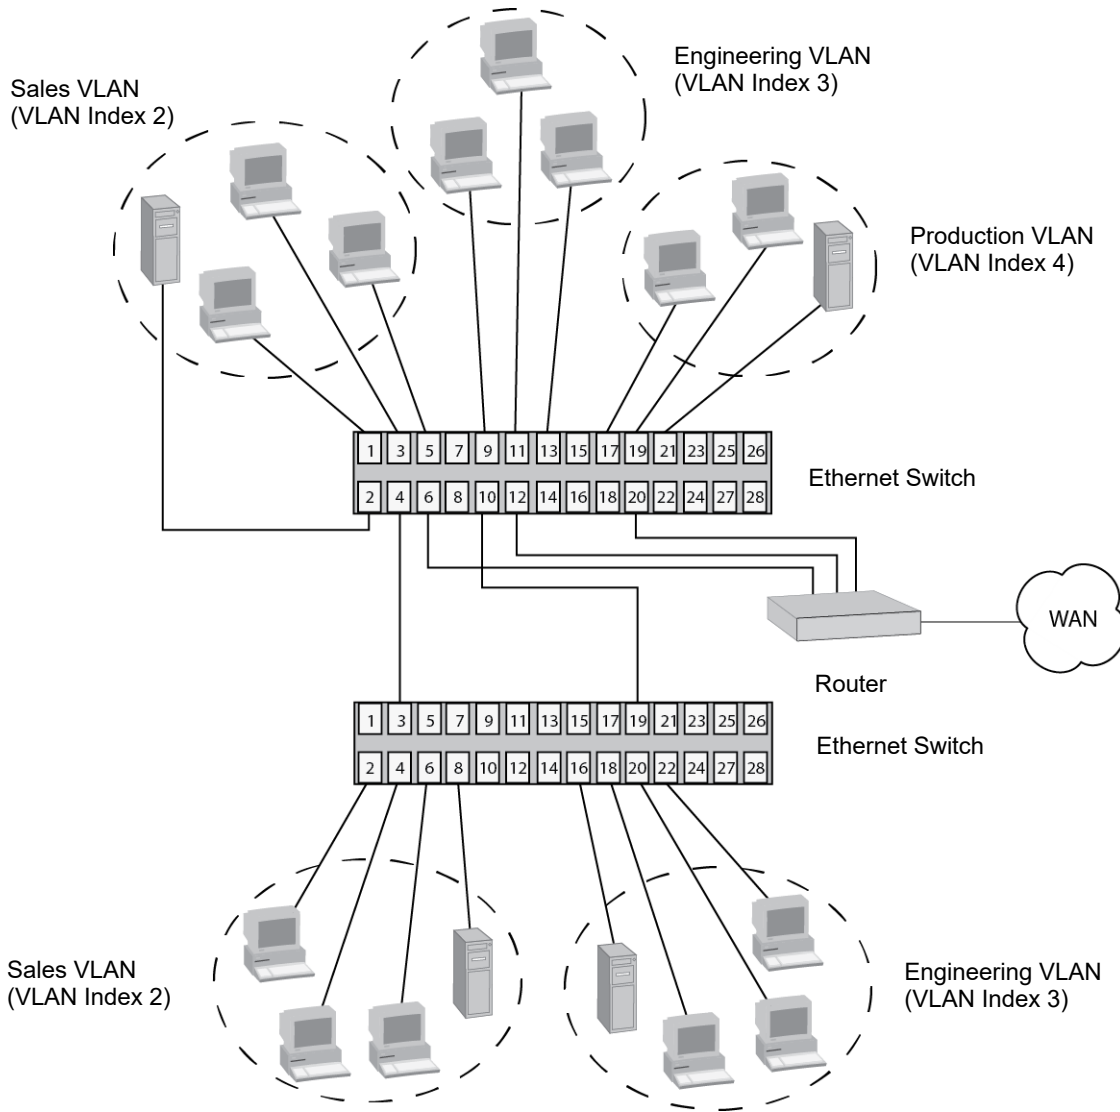


Figure 81. Port-Based VLAN - Example 2

Table 36 lists the port assignments for the Sales, Engineering, and Production VLANs on the switches:

Table 36. Example 2 of Port-Based VLANs

	Sales (VLAN Index 2)	Engineering (VLAN Index 3)	Production VLAN (VLAN Index 4)
Top Ethernet Switch	Ports 1 - 6	Ports 9 - 13	Ports 17, 19 - 21
Bottom Ethernet Switch	Ports 2 - 4, 6, 8	Ports 16, 18-20, 22	none

The VLANs are described here:

- ❑ Sales VLAN - This VLAN spans both switches. It has the VLAN index 2 and consists of six ports on the top switch and five ports on the bottom switch.

The two parts of the VLAN are connected by a direct link from port 4 on the top switch to port 3 on the bottom switch. This direct link allows the two parts of the Sales VLAN to function as one logical LAN segment.

Port 6 on the top switch connects to the router. This port allows the Sales VLAN to exchange Ethernet frames with the other VLANs and to access the WAN.

- ❑ Engineering VLAN - The workstations of this VLAN are connected to ports 9 to 13 on the top switch and ports 16, 18 to 20, and 22 on the bottom switch.

Because this VLAN spans multiple switches, it needs a direct connection between its various parts to provide a communications path. This is provided with a direct connection from port 10 on the top switch to port 19 on the bottom switch.

This VLAN uses port 12 on the top switch as a connection to the router and the WAN.

- ❑ Production VLAN - This is the third VLAN in the example. It has the VLAN of 4, and its ports have been assigned the VLAN index 4.

The nodes of this VLAN are connected only to the top switch. So the VLAN does not require a direct connection to the bottom switch. However, it uses port 20 as a connection to the router.

Adding or Viewing Port-Based VLANs

This section contains the procedure for adding or viewing Port-Based VLANs.

Note

After adding a Port-Based VLAN, you have to set the modes of the ports to Port-Based VLAN. Refer to “Assigning the Port-Based VLAN Mode to Ports” on page 193.

Perform the following procedure to add or view Port-Based VLANs:

1. Select **Bridge > VLAN > Port-Based VLAN** from the main menu. Refer to Figure 82.

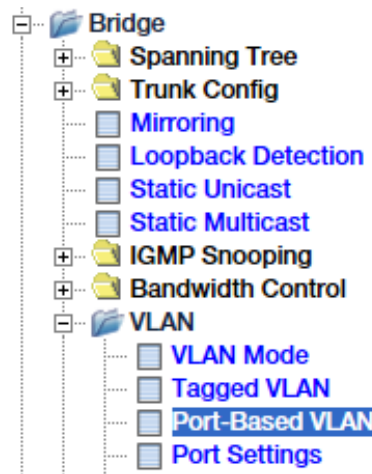


Figure 82. Port-Based VLAN Menu Selection

The Port-Based VLAN window is shown in Figure 83 on page 191.

Port-Based VLAN

VLAN Index: (1-52)

VLAN Name: (32 characters limit)

VLAN Member

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Not Member

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>

Port-Based VLAN Table

VLAN Index	VLAN Name	VLAN Member	Action
2	nnnn	1-5	<input type="button" value="Modify"/> <input type="button" value="Delete"/>

Page 1 / 1 Page

Figure 83. Port-Based VLAN Window

The table at the bottom of the window lists the current port-based VLANs on the switch.

- Configure the fields in Table 37.

Table 37. Port-Based VLAN Window

Field	Description
VLAN Index	<p>Enter a unique VLAN index number for the VLAN. Here are the guidelines:</p> <ul style="list-style-type: none"> - A VLAN can have only one index number. - The range is 1 to 52. - The VLAN index has to be unique from all other Port-Based VLANs on the switch. - Index numbers simply identify the VLANs on the switch. They are not the same as VID numbers for 802.1Q Tagged VLANs.

Table 37. Port-Based VLAN Window (Continued)

Field	Description
VLAN Index (Continued)	<ul style="list-style-type: none"> - You should assign VLANs that span multiple switches the same index number on all the switches. For example, if the Sales VLAN spans three switch, you should assign the same index number to the three parts of the VLAN.
VLAN Name	<p>Enter a name for the VLAN. Here are the guidelines:</p> <ul style="list-style-type: none"> - The name can be up to 32 characters. - Spaces are allowed, but not special characters. - The name should be unique from the names of all other VLANs on the switch. - The name of a VLAN that spans multiple switches should be the same on all the switches.
VLAN Member	Click the radio buttons of ports that are to be members of the new VLAN.
Not Member	Click the radio buttons of ports that are not to be members of the VLAN. This is the default setting. These ports are members of the Default VLAN or other VLANs.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.
5. Review the following:
 - Set the modes of the VLAN ports to Port-Based VLAN by performing “Assigning the Port-Based VLAN Mode to Ports” on page 193. This will remove the ports from any tagged VLAN assignments, including the Default VLAN.
 - Ports can belong to more than one Port-Based VLAN at a time. If the ports of the new VLAN already belong to other Port-Based VLANs and you want them to be members only of the new VLAN, you have to remove them from their current VLAN assignments. Refer to “Modifying Port-Based VLANs” on page 196.

Assigning the Port-Based VLAN Mode to Ports

This section contains the procedure for setting the VLAN modes of the switch ports. There are two VLAN modes. One mode is for ports that are members of 802.1Q Tagged VLANs and the other for Port-Based VLANs. The default port mode is 802.1Q Tagged VLAN. Here are the situations where you need to perform this procedure:

- ❑ You add a new Port-Based VLAN to the switch and some or all its ports are members of the default 802.1Q Tagged VLAN or another Tagged VLAN.
- ❑ You add ports to an existing Port-Based VLAN and the ports are members of the default 802.1Q Tagged VLAN or another Tagged VLAN.

Note

Ports have to be members of Port-Based VLANs before you can set their modes. Otherwise you will see an error message. Refer to “Adding or Viewing Port-Based VLANs” on page 190 or “Modifying Port-Based VLANs” on page 196.

Perform the following procedure to set the VLAN modes of ports to Port-Based VLAN:

1. Select **Bridge > VLAN > VLAN Mode** from the main menu. Refer to Figure 84.

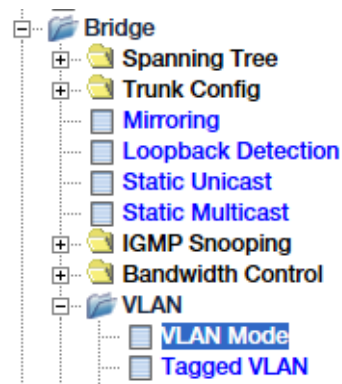


Figure 84. VLAN Mode Menu Selection

The VLAN Mode window is shown in Figure 85. The two rows in the window represent the port settings for the two port VLAN modes. The default setting for all ports is 802.1Q Tagged VLAN.

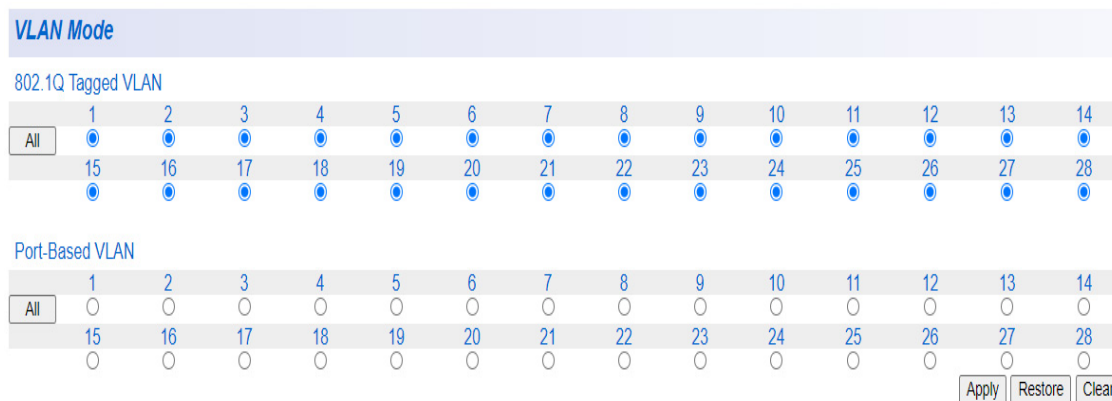


Figure 85. VLAN Mode Window

- In the Port-Based VLAN row, click the radio circles of ports that are members of Port-Based VLANs. In the example in Figure 86, ports 1 to 9, 14 to 19, and 28 are designated as members of Port-Based VLANs.

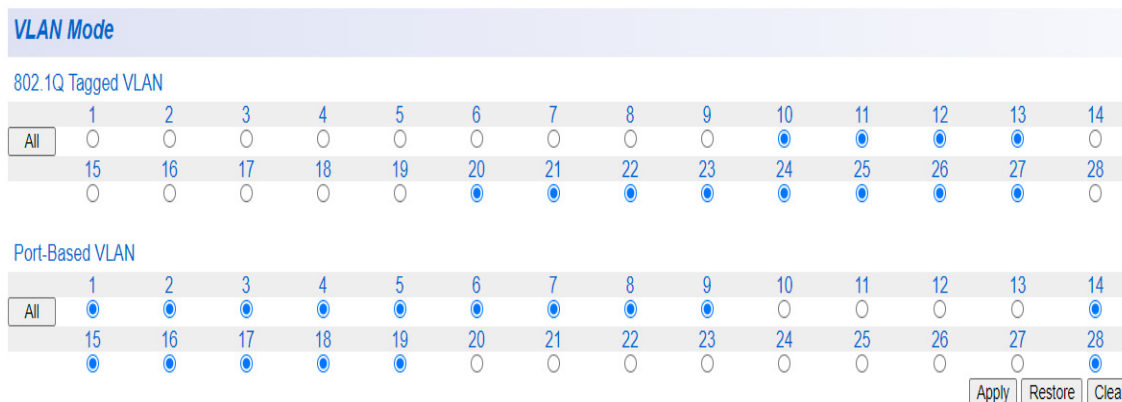


Figure 86. Example of the VLAN Mode Window

- Click the **Apply** button.

Note

The switch displays an error message if the selected ports are not already members of Port-Based VLANs. If you see the message, verify that you selected the correct ports in the VLAN Mode window. If necessary, add the ports to Port-Based VLANs by performing “Adding or Viewing Port-Based VLANs” on page 190 or “Modifying Port-Based VLANs” on page 196 and then repeat this procedure.

If you want to discard your changes to the window and have not yet clicked the Apply button, click the **Restore** button to restore the previous settings.

4. Select **Save Settings to Flash** from the main menu to save your changes.

Note the following:

- If you performed this procedure to add a new Port-Based VLAN to the switch, the VLAN is now operational.
- If you performed this procedure to add or remove ports from an existing new Port-Based VLAN, your changes are now operational.

Modifying Port-Based VLANs

This section contains the procedure for changing the names or ports of Port-Based VLANs. Review the following before performing the procedure:

- ❑ You cannot change the VLAN index number of a VLAN. Changing the number requires deleting a VLAN and adding it again with the new index number.
- ❑ When removing ports from a Port-Based VLAN that are not members of other Port-Based VLANs, you first have to change their VLAN mode to 802.1Q Tagged VLAN. This allows the switch to return the ports back to the default 802.1Q Tagged VLAN after the Port-Based VLAN is deleted. Refer to “Assigning the Port-Based VLAN Mode to Ports” on page 193. For example, to remove port 4 from a Port-Based VLAN and return it back to the default VLAN, you first go to “Assigning the Port-Based VLAN Mode to Ports” on page 193 and change its port mode to 802.1Q Tagged VLAN before performing this procedure.

Perform the following procedure to modify port-based VLANs:

1. Select **Bridge > VLAN > Port-Based VLAN** from the main menu. Refer to Figure 82 on page 190. The Port-Based VLAN window is shown in Figure 83 on page 191.
2. In the VLAN Action column, click the **Modify** button of the VLAN you want to modify. You can modify only one VLAN at a time. The Port-Based VLAN page is displayed with the settings of the VLAN. An example is shown in Figure 87.

Port-Based VLAN

VLAN Index:

VLAN Name: (32 characters limit)

VLAN Member

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Not Member

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>

Figure 87. Example of the Modify Port-Based VLAN Window

3. Modify the VLAN settings by referring to Table 37 on page 191.
4. Click the **Apply** button.

If you removed ports from the VLAN and the following message is displayed, you need to set the VLAN mode of the ports to 802.1Q Tagged VLAN mode so that the switch can return them to the default 802.1Q Tagged VLAN. Refer to “Assigning the Port-Based VLAN Mode to Ports” on page 193.

Please change delete port VLAN Mode to 802.1Q Tagged VLAN.

5. Select **Save Settings to Flash** from the main menu to save your changes.
6. Review the following:
 - If you added ports to the VLAN and the ports are members of the default 802.1Q Tagged VLAN or another tagged VLAN, go to “Assigning the Port-Based VLAN Mode to Ports” on page 193 and set the new ports to the Port-Based VLAN mode. This removes the ports from their current tagged VLAN assignments. To view all the VLANs on the switch, refer to “Viewing Port-based and 802.1Q Tagged VLANs” on page 223.
 - Ports can belong to more than one Port-Based VLAN at a time. Consequently, if you add ports to a VLAN and the ports already belong to another Port-Based VLAN, you should remove them from their other VLANs. Refer to “Modifying Port-Based VLANs” on page 196.

Deleting Port-Based VLANs

This section contains the steps to deleting Port-Based VLANs from the switch:

Step 1 Before deleting a Port-Based VLAN, consider the following:

- ❑ Do you want to assign its ports to another Port-Based VLAN? If so, add the new VLAN first with the ports and then delete the old VLAN. Refer to “Adding or Viewing Port-Based VLANs” on page 190.
- ❑ Do you want to return its ports to the default 802.1Q Tagged VLAN, perhaps prior to assigning them to another 802.1Q Tagged VLAN? If so, before deleting the Port-Based VLAN, change the VLAN modes of its ports to 802.1Q Tagged VLAN, as explained in “Assigning the Port-Based VLAN Mode to Ports” on page 193.

For example, assume you want to delete a Port-Based VLAN called Sales that had ports 1 to 8. Assume you want to assign ports 1 to 4 to a different Port-Based VLAN called Production and return ports 5 to 8 to the default 802.1Q Tagged VLAN. Here are the steps:

1. Add the new Production Port-Based VLAN with ports 1 to 4.
2. Change the port modes of ports 5 to 8 to 802.1Q Tagged VLAN so that the switch can return the ports to the default VLAN.
3. Delete the Sales VLAN.

Step 2 Perform the following procedure to delete Port-Based VLANs:

1. Select **Bridge > VLAN > Port-Based VLAN** from the main menu. Refer to Figure 82 on page 190. The Port-Based VLAN window is shown in Figure 83 on page 191.
2. In the Port-Based VLAN Table, click the **Delete** button in the Action column of the VLAN you want to delete. You can delete only one VLAN at a time.

If the following warning message is displayed, perform “Assigning the Port-Based VLAN Mode to Ports” on page 193 and set the modes of the VLAN ports that are not members of other Port-Based VLANs to the 802.1Q Tagged VLAN mode. This will allow the switch to return the ports to the default 802.1Q Tagged VLAN. Refer to the example in “Step 1” above.

Please change delete port VLAN Mode to 802.1Q Tagged VLAN.

3. Click the **Apply** button. The VLAN is delete from the switch. The ports of the VLAN are returned to the Default 802.1Q VLAN.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Selecting the MAC Address Forwarding Table Mode

The switch supports the following modes for learning MAC addresses on VLANs:

- ❑ Independent VLAN learning (IVL): The switch maintains separate MAC address forwarding tables for the VLANs. This is the default setting.
- ❑ Shared VLAN learning (SVL): The switch maintains one MAC address forwarding table for all the VLANs.

For more information, refer to the IEEE 802.1Q standard.

Perform the following procedure to select the MAC address learning mode:

1. Select **Bridge.> VLAN.> Forwarding Table Mode** from the main menu. Refer to Figure 88.

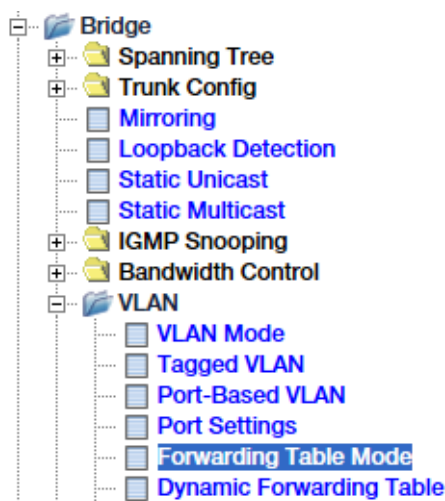


Figure 88. Forwarding Table Mode Menu Selection

The Forwarding Table Mode window is shown in Figure 89.



Figure 89. Forwarding Table Mode Window

2. Select **IVL** or **SVL** from the **Learning Mode** menu:

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Viewing the MAC Addresses of Network Devices in the Forwarding Table

This procedure explains how to view the MAC addresses the switch has stored in its forwarding table. You can view all the addresses or the addresses of specific ports.

Perform the following procedure to view the MAC addresses in the forwarding table:

1. Select **Bridge > VLAN > Dynamic Forwarding Table** in the main menu. Refer to Figure 90.

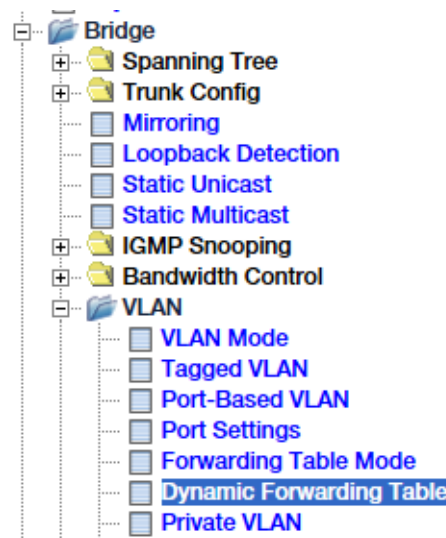


Figure 90. Dynamic Forwarding Table Menu Selection

An example of the Dynamic Forwarding Table window is shown in Figure 91.

Dynamic Forwarding Table

Port: All ▾ Refresh

ID	VID	Port	MAC Address	Type	VLAN Mode
1	1	23	00-21-CC-B5-75-9C	Dynamic	802.1Q
2	1	1	00-22-57-94-9D-D4	Dynamic	802.1Q
3	1	1	00-22-57-94-9D-D5	Dynamic	802.1Q

Page 1 / 1 First Page Previous Page Next Page Last Page Page Go

Figure 91. Dynamic Forwarding Table Window

The table columns are described in Table 38.

Table 38. Dynamic Forwarding Table Window

Column	Description
Index	Displays the index number of the address in the table.
VID	Displays one of the following: <ul style="list-style-type: none"> - If the MAC address is part of an 802.1Q Tagged VLAN, the column displays the VID. - If the MAC address is part of a Port-Based VLAN, the column displays 4094.
Port	Displays the port or trunk number. Port numbers with a "p" prefix indicate port trunks. For example, "p2" indicates port trunk 2.
MAC Address	Displays the MAC address from the Port menu.
Type	Displays one of the following: <ul style="list-style-type: none"> - Dynamic: The switched learned the MAC address. - Static: The MAC address was entered manually.
VLAN Mode	Displays one of the following: <ul style="list-style-type: none"> - 802.1Q: Belongs to IEEE 802.1Q VLAN learning entries. - PortBase: Belongs to Port-Based VLAN learning entries.

2. To view all the MAC addresses, select **All** from the Port menu. To view MAC addresses of a specific port, select the port from the Port menu.

Chapter 16

802.1Q Tagged Virtual LANs

This chapter describes 802.1Q Tagged VLANs in the following sections:

- ❑ “802.1Q Tagged VLAN Overview” on page 206
- ❑ “Adding or Viewing 802.1Q Tagged VLANs” on page 212
- ❑ “Assigning the 802.1Q Tagged VLAN Mode to Ports” on page 216
- ❑ “Configuring PVIDs and Filters for Tagged and Untagged Ports” on page 219
- ❑ “Modifying 802.1Q Tagged VLANs” on page 221
- ❑ “Deleting 802.1Q Tagged VLANs” on page 222
- ❑ “Viewing Port-based and 802.1Q Tagged VLANs” on page 223

802.1Q Tagged VLAN Overview

As explained in the “VLAN Overview” on page 184, a VLAN consists of a group of ports that form an independent traffic domain on an Ethernet switch. The ports of a VLAN forward ingress traffic only to those ports that are part of the same VLAN. An interconnection device, such as a router or Layer 3 switch, is commonly added to permit traffic between different VLANs.

The switch supports several types of VLANs. This chapter describes 802.1Q Tagged VLANs. Tagged VLANs commonly consist of two types of switch ports: tagged and untagged.

Tagged Ports

Tagged ports are connected to network devices that support 802.1Q tagged packets. Tagged packets contain in the headers, following the source and destination addresses, the VLAN identifiers (VIDs) of the virtual LANs to which the frames belong (IEEE 802.3ac standard). When the switch receives frames with VLAN tags, referred to as a tagged frame, on tagged ports, it forwards them only to those ports that share the same VID.

Network devices connected to tagged ports should be IEEE 802.1Q compliant. This is the standard that outlines the requirements and standards for tagging. The device must be able to process the tagged information on received frames and add tagged information to transmitted frames.

The benefit of tagged ports is that they can belong to more than one VLAN at one time, which can simplify the task of adding shared devices to the network. For example, an 802.1Q-compliant server connected to a tagged port on the switch can be configured to accept and return packets from different VLANs simultaneously.

Tagged VLANs are also useful where multiple VLANs span across switches. You can use one port per switch to connect the different parts of VLANs together.

The IEEE 802.1Q standard describes how this tagging information is used to forward the traffic throughout the switch. The handling of frames tagged with VIDs coming into a port is straightforward. If the incoming frame’s VID tag matches one of the VIDs of a VLAN of which the port is a tagged member, the frame is accepted and forwarded to the appropriate ports. If the frame’s VID does not match any of the VLANs in which the port is a member, the frame is discarded.

The VLAN information within an Ethernet frame is referred to as a *tag* and is contained in a *tagged header* for the frame. A tag, which follows the source and destination addresses in a frame, contains the VLAN ID of the VLAN to which the frame belongs (IEEE 802.3ac standard). This number

uniquely identifies each VLAN in a network.

A port that receives or transmits tagged frames is referred to as a *tagged port*. Any network device connected to a tagged port must be IEEE 802.1Q compliant. This is the standard that outlines the requirements and standards for VLAN tagging. The device must be able to process the tagged information on received frames and add tagged information to transmitted frames.

Untagged Ports

Untagged ports in Tagged VLANs are for network devices that do not support 802.1Q tagged packets. The packets that they transmit and receive do not contain VIDs in the headers. VLAN memberships of untagged ports are controlled by port VLAN identifiers (PVIDs). These are values that you assign to the individual switch ports that identify which VLANs the ports, and the devices connected to the ports, belong. Ports can have only one PVID.

The PVIDs of untagged ports should match the VID of their Tagged VLANs. For example, the untagged ports of a Tagged VLAN with the VID 56 should be assigned the PVID 56. Refer to “Configuring PVIDs and Filters for Tagged and Untagged Ports” on page 219.

802.1Q Tagged VLAN Components

The components of Tagged VLANs are described Table 39.

Table 39. Components of 802.1Q Tagged VLANs

Item	Description
VLAN Identifier	<p>A tagged VLAN has to have a unique identification number (VID). The VID has to match the VID in the tagged packets transmitted by the 802.1Q-compliant network devices in the VLAN. Tagged VLANs that span multiple switches should be assigned the same VID on each switch. The range is 2 to 4093. The VID 1 is reserved for the default Tagged VLAN.</p> <p>This number is not the same as the VLAN index number in Port-based VLANs. They are different numbering systems.</p>
Name	<p>A tagged VLAN has to have a name. VLANs will be easier to identify if you give them names that reflect the functions of the member devices, such as Sales, Production, or Engineering.</p>

Table 39. Components of 802.1Q Tagged VLANs (Continued)

Item	Description
Tagged Ports	<p>A tagged VLAN will have one or more tagged ports. These are ports that are connected to 802.1Q-compliant network devices that transmit tagged packets containing the VID that identifies their VLAN memberships. Here are the guidelines:</p> <ul style="list-style-type: none"> - A tagged VLAN can have any number of tagged ports, up to all the ports on the switch. - Tagged ports can belong to more than one tagged VLAN at a time. - Tagged ports cannot also be members of Port-based, voice, or private VLANs.
Untagged Ports	<p>A tagged VLAN can have one or more untagged ports. These are ports that are connected to network devices that do not support 802.1Q tagged packets. VLAN memberships of untagged ports are identified by the PVIDs assigned to the individual ports. Here are the guidelines:</p> <ul style="list-style-type: none"> - A tagged VLAN can have any number of untagged ports. - Untagged ports can belong to only one tagged VLAN at a time. - Untagged ports cannot also be members of Port-based, voice, or private VLANs. - Untagged ports should be assigned PVIDs that match the VID of their VLANs. For example, untagged ports in a tagged VLAN with the VID 11 should be assigned the PVID 11. Refer to “Configuring PVIDs and Filters for Tagged and Untagged Ports” on page 219.
VLAN Mode	<p>The VLAN mode identifies whether ports are members of Port-based or tagged VLANs. The default mode for ports is Tagged VLAN. You set the VLAN mode after adding ports to tagged or untagged VLANs. Refer to “Assigning the 802.1Q Tagged VLAN Mode to Ports” on page 216.</p>

Table 39. Components of 802.1Q Tagged VLANs (Continued)

Item	Description
PVID	PVIDs mainly apply to untagged ports. They identify the VLAN memberships of the untagged packets that arrive on untagged ports. However, PVIDs are applicable on tagged ports as well. As explained earlier, VLAN membership on tagged ports is typically controlled by the tag information in the header portion of the ingress frames themselves. But tagged ports do use PVIDs for any ingress untagged packets they might receive.

Guidelines to Adding Tagged VLANs

Here are the guidelines to adding tagged VLANs.

- ❑ Tagged VLANs can contain both tagged and untagged ports.
- ❑ Tagged ports are connected to network devices that support 802.1Q tagged packets.
- ❑ Untagged ports are connected to network devices that do not support 802.1Q tagged packets.
- ❑ A tagged VLAN needs to have a unique VID. If a VLAN spans multiple switches, you should assign the same VID to the various parts of the VLAN on the different switches.
- ❑ Tagged ports can be members of more than one tagged VLAN at a time.
- ❑ Untagged ports can be untagged members of only one VLAN at a time.
- ❑ You have to assign PVIDs to the untagged ports that match the VLAN's identifier. For example, untagged ports in the VLAN with the VID45 should be assigned PVID 45. Refer to "Configuring PVIDs and Filters for Tagged and Untagged Ports" on page 219.
- ❑ You should also assign PVIDs to tagged ports. The PVIDs will identify the VLAN memberships of any ingress untagged packets on tagged ports.

Tagged VLAN Example

Figure 92 on page 210 illustrates how tagged ports can be used to interconnect IEEE 802.1Q-based products.

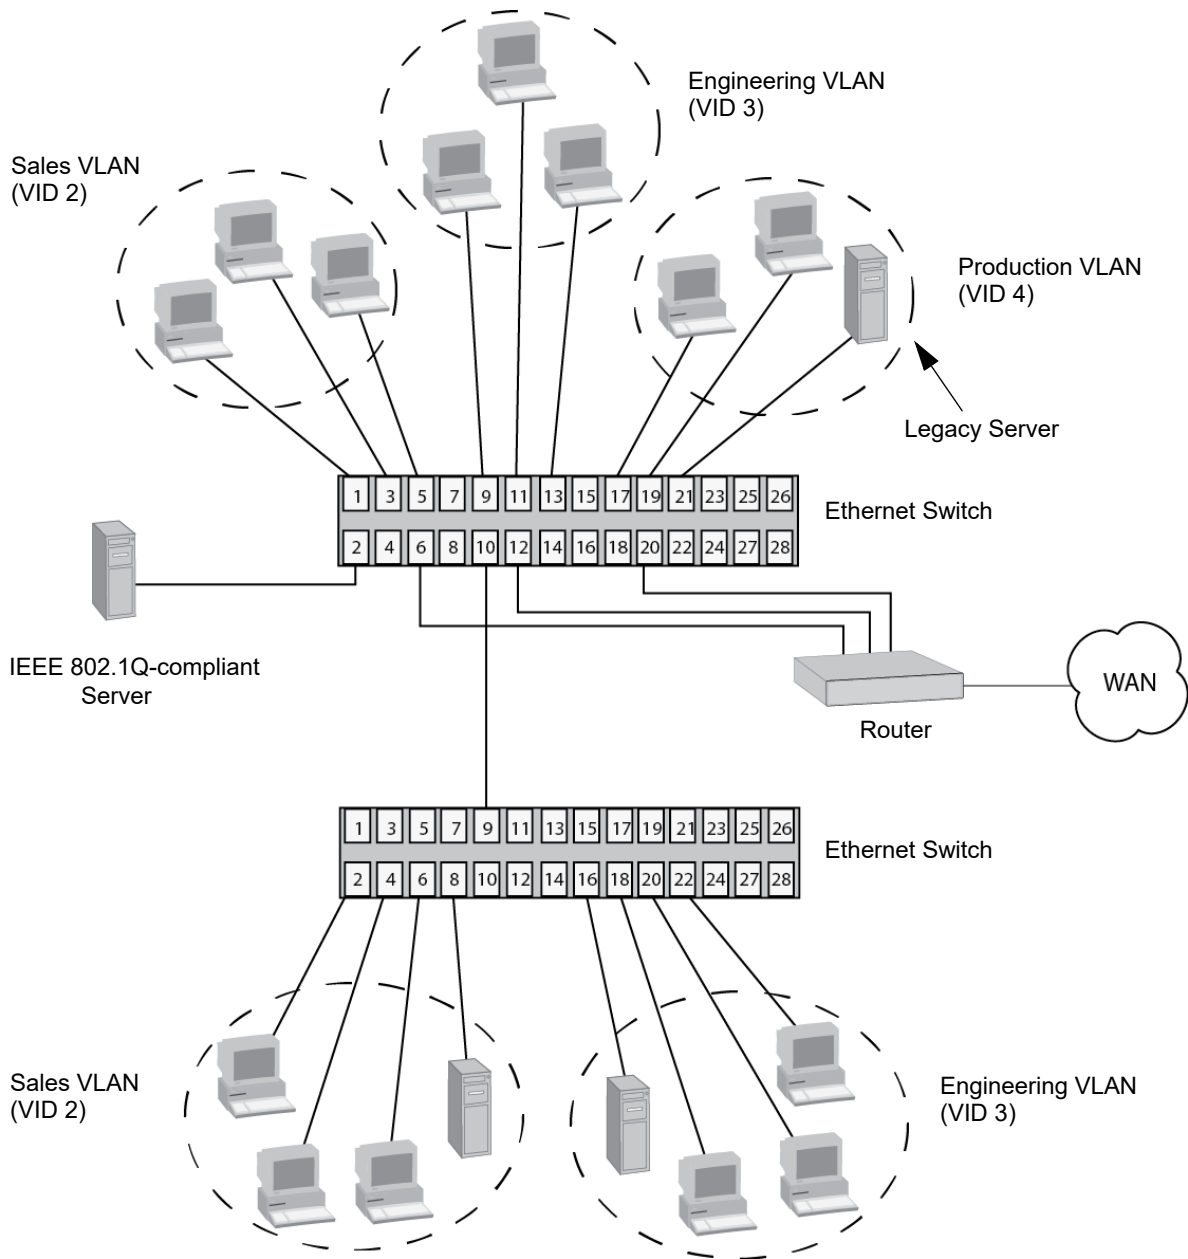


Figure 92. Example of a Tagged VLAN

The port assignments of the VLANs are described in Table 40 on page 211.

Table 40. Example of Tagged VLANs

	Sales VLAN (VID 2)		Engineering VLAN (VID 3)		Production VLAN (VID 4)	
	Untagged Ports	Tagged Ports	Untagged Ports	Tagged Ports	Untagged Ports	Tagged Ports
Top Ethernet Switch	1, 3 to 5 (PVID 2)	2, 10	9, 11 to 13 (PVID 3)	2, 10	17, 19 to 21 (PVID 4)	2
Bottom Ethernet Switch	2, 4, 6, 8 (PVID 2)	9	16, 18, 20, 22 (PVID 3)	9	none	none

The switches employ tagged ports to simplify network implementation and management. One of the tagged ports is port 2 on the top switch. It is a tagged member of the three VLANs. It is connected to an IEEE 802.1Q-compliant server, meaning the server can handle frames from multiple VLANs. The port makes it possible for the three VLANs to access the server without going through a router or other interconnection device. It is important to note that even though the server is accepting frames from and transmitting frames to more than one VLAN, data separation and security remain.

Two other tagged ports are used to simplify network design in the example. They are port 10 on the top switch and port 9 on the lower switch. These ports are tagged members of the Sales and Engineering VLANs so that they can carry traffic from both VLANs, simultaneously. They provide a common connection that enables different parts of the same VLAN to communicate with each other while maintaining data separation between the VLANs.

Adding or Viewing 802.1Q Tagged VLANs

This section contains the procedure for adding or viewing 802.1Q Tagged VLANs on the switch. Here are the basic steps:

1. Add the VLAN by defining its name, VID, and tagged and untagged ports. The procedure is explained in this section.
2. Set the VLAN modes of the ports to 802.1Q Tagged VLAN. You might be able to skip this step because Tagged VLAN is the default port settings. Refer to “Assigning the 802.1Q Tagged VLAN Mode to Ports” on page 216.
3. Set the PVIDs of the untagged ports to match the VID of the VLAN. Refer to “Configuring PVIDs and Filters for Tagged and Untagged Ports” on page 219.
4. Set the PVID on tagged ports to control the VLAN assignment of ingress untagged packets.

This section contains the first step to adding 802.1Q Tagged VLANs to the switch. A new Tagged VLAN requires the following information:

- VID number
- Name
- Tagged ports
- Untagged ports

Perform the following procedure to add a new 802.1Q Tagged VLAN:

1. Select **Bridge > VLAN > Tagged VLAN** from the main menu. Refer to Figure 93.

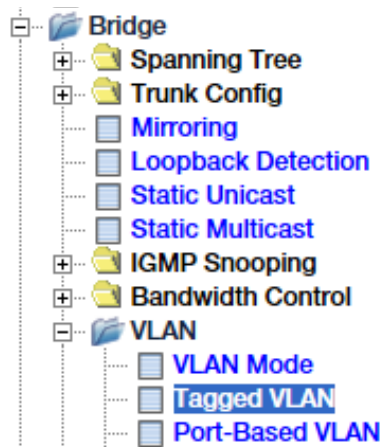


Figure 93. Tagged VLAN Menu Selection

The Tagged VLAN window is shown in Figure 94.

Tagged VLAN

VLAN ID: (2-4093)

VLAN Name: (32 characters limit)

Management VLAN:

Static Tagged

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Static Untagged

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Not Member

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>

VLAN ID	Name	VLAN Type	Management	Action
1		Permanent	Enabled	<input type="button" value="Modify"/>
2	test	Static	Disabled	<input type="button" value="Modify"/> <input type="button" value="Delete"/>

Note: If a port does not belong to any VLAN, its PVID will be changed to default VLAN ID and attends to default VLAN automatically.

Page 1 / 1 Page

Figure 94. Tagged VLAN Window

The table at the bottom of the window lists the current 802.1Q Tagged VLANs on the switch. However, the table does not include the VLAN ports. For that, refer to “Viewing Port-based and 802.1Q Tagged VLANs” on page 223. The switch comes with one default Tagged VLAN, with the VLAN ID 1. You cannot delete it or change its VID.

2. Configure the fields in Table 41 on page 214.

Table 41. Tagged VLAN Window

Field	Description
VLAN ID	Enter a VLAN ID. The guidelines are listed here: <ul style="list-style-type: none"> - A VLAN can have only one VID. - The VID has to be unique from all other Tagged VLANs on the switch. - The range is 2 to 4093. - VID 1 is reserved for the default Tagged VLAN. - A tagged VLAN that spans more than one switch should be assigned the same VID on all the switches. - VIDs are not the same as VLAN Index numbers of Port-based VLANs.
VLAN Name	Enter a name for the VLAN. The guidelines are listed here: <ul style="list-style-type: none"> - The name can be up to 32 characters. - Spaces are allowed, but not special characters. - The name should be unique from the names of all other VLANs on the switch. - A VLAN that spans more than one switch should be given the same name on all the switches.
Management VLAN	Choose one of the following: <ul style="list-style-type: none"> - Enabled: Allows you to manage the switch with the web browser and command line interfaces through the ports of the VLAN. - Disabled: Prevents you from managing the switch through the VLAN ports. This is the default setting. The guidelines are listed here: <ul style="list-style-type: none"> - There can be only one management VLAN. - VLAN 1 is the default management VLAN. - You can access management through the tagged port of all VLANs on which you have enabled management.

Table 41. Tagged VLAN Window (Continued)

Field	Description
Static Tagged	Click the radio buttons of ports that are to be tagged members of the VLAN. Ports can be tagged members of more than one tagged VLAN at a time. The devices connected to these ports should be 802.1Q-compliant. Refer to “Tagged Ports” on page 206.
Static Untagged	Click the radio buttons of ports that are to be untagged members of the VLAN. Untagged ports can be members of more than one VLAN at a time. These ports should be connected to network devices that do not supported tagged packets and 802.1Q. Refer to “Untagged Ports” on page 207.
Not Member	Click the radio buttons of ports that are not to be members of the VLAN. This is the default setting. These ports are members of the Default 802.1Q Tagged VLAN or other VLANs.

Note

The Management VLAN is always Enabled on the untagged ports of the Default VLAN. It cannot be disabled on the Default VLAN.

- Click the **Apply** button. The new VLAN is added to the switch.

Note

If your management session is interrupted, the new VLAN contains the port through which you were managing the switch, but it is not a Management VLAN. To resume managing the switch, access the switch through a port that is a member of a VLAN where Management VLAN is enabled, such as the Default VLAN.

- Select **Save Settings to Flash** from the main menu to save your changes.
- Do the following:
 - Set the VLAN modes of the ports to 802.1Q Tagged VLAN. Refer to “Assigning the 802.1Q Tagged VLAN Mode to Ports” on page 216.
 - Set the PVIDs of the tagged and untagged ports to match the VID of the VLAN. Refer to “Configuring PVIDs and Filters for Tagged and Untagged Ports” on page 219.

Assigning the 802.1Q Tagged VLAN Mode to Ports

After adding a new 802.1Q Tagged VLAN to the switch or adding ports to an existing VLAN, your next step is to verify and, if needed, set the VLAN Modes of the ports. The VLAN Mode defines whether ports are members of 801.2Q Tagged VLANs or Port-based VLANs. The default setting for ports is 802.1Q Tagged VLANs.

Note

Ports have to be members of Port-based or 802.1Q Tagged VLANs before you can set their VLAN modes. Refer to “Adding or Viewing 802.1Q Tagged VLANs” on page 212 or “Adding or Viewing Port-Based VLANs” on page 190.

Perform the following procedure to set port VLAN Modes:

1. Select **Bridge > VLAN > VLAN Mode** from the main menu. Refer to Figure 95.

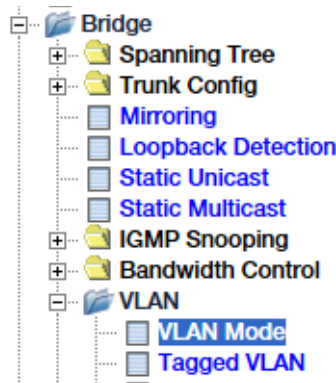


Figure 95. VLAN Mode Menu Selection

The VLAN Mode window is shown in Figure 96. The figure shows the ports at their default settings.

VLAN Mode														
802.1Q Tagged VLAN														
All	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
Port-Based VLAN														
All	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
												Apply	Restore	Clear

Figure 96. VLAN Mode Window

- Click the radio circles to indicate whether ports belong to 802.1Q Tagged VLANs or Port-Based VLAN. The default is 802.1Q Tagged VLAN.

The example in Figure 97 identifies ports 10 to 13 and 20 to 27 as members of 802.1Q Tagged VLANs and the other ports as members of Port-based VLANs.

VLAN Mode														
802.1Q Tagged VLAN														
All	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>
Port-Based VLAN														
All	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
												Apply	Restore	Clear

Figure 97. Example of the VLAN Mode Window

If you want to discard your changes to the window and have not yet clicked the Apply button, clicking the **Restore** button restores the previous settings.

- Click the **Apply** button.

4. Select **Save Settings to Flash** from the main menu to save your changes.
5. Go to “Configuring PVIDs and Filters for Tagged and Untagged Ports” on page 219.

Configuring PVIDs and Filters for Tagged and Untagged Ports

This section explains how to set the PVIDs on tagged and untagged ports in 802.1Q Tagged VLANs. The switch uses PVIDs to identify the VLAN memberships of ingress untagged packets on the ports. For background information, refer to “802.1Q Tagged VLAN Components” on page 207.

This section also explains how to set packet filters for tagged or untagged packets and enable or disable destination MAC address filtering.

Note

This procedure does not apply to Port-based, private, or Voice VLANs.

Perform the following procedure to set port PVID and filters:

1. Select **Bridge > VLAN > Port Settings** from the main menu. Refer to Figure 98.

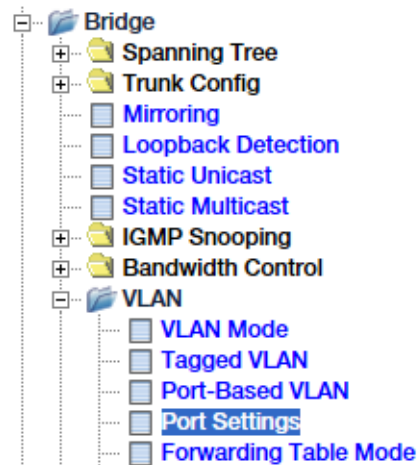


Figure 98. Port Settings Menu Selection

The Port Settings window is shown in Figure 99.

Port Settings				
Port	PVID	Acceptable Frame Types	Ingress Filtering	Action
All		Ignore	Ignore	Apply
1	1	All	Enabled	Apply
2	1	All	Enabled	Apply
3	1	All	Enabled	Apply
4	1	All	Enabled	Apply

Figure 99. Port Settings Window

Note

You cannot change these values on ports that are members of Port-based, private, or Voice VLANs.

2. Configure the fields in Table 42.

Table 42. Port Settings Window

Field	Description
PVID	<p>Enter a PVID value for the port. Here are the guidelines:</p> <ul style="list-style-type: none"> - Ports can have only one PVID. - PVIDs of untagged ports should be the same as the VID of their VLANs. For example, untagged ports in an 802.1Q Tagged VLAN with the VID 87 should be assigned PVID 87. - PVIDs of tagged ports should designate which VLAN will handle ingress untagged packets. - The default is PVID 1. <p>Refer to “802.1Q Tagged VLAN Components” on page 207.</p>
Acceptable Frame Types	<p>Select one of the following from the menu:</p> <ul style="list-style-type: none"> - All: The port accepts all ingress packets. This is the default. - Tagged: The port accepts only tagged packets and discards untagged packets. - Untagged and Priority Tagged: The port accepts untagged packets and those with priority tags, and discards tagged packets.
Ingress Filtering	<p>Select one of the following from the menu:</p> <ul style="list-style-type: none"> - Enabled: Enables destination MAC address ingress filtering on the port. Refer to “Destination MAC Address Filters Overview” on page 476. This is the default setting. - Disabled: Disables ingress filtering on the port.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying 802.1Q Tagged VLANs

Perform the following procedure to modify 802.1Q Tagged VLANs:

1. Select **Bridge > VLAN > Tagged VLAN** from the main menu. Refer to Figure 93 on page 212. An example of the Tagged VLAN window is shown in Figure 94 on page 213.
2. In the **VLAN Action** column of the table, click the **Modify** button of the VLAN you want to modify. You can modify only one VLAN at a time.

The current settings of the VLAN are displayed in the Modify VLAN window.

3. Modify the VLAN settings by referring to Table 41 on page 214. Review the following:
 - You cannot change the VID. Changing the VID requires deleting the VLAN and adding it again with the new VID.
 - Untagged ports removed from VLANs are automatically returned to the Default VLAN.
 - Tagged ports removed from VLANs retain their other VLAN assignments.
 - Your management session will be interrupted if you disable Management VLAN on the VLAN through which you are currently managing the switch. To resume managing the device, access the switch through a port that is a member of a VLAN where Management VLAN is enabled.
 - You cannot disable Management VLAN on the Default VLAN.
4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Deleting 802.1Q Tagged VLANs

This section contains the procedure for deleting 802.1Q Tagged VLANs from the switch. Review the following:

- You cannot delete the Default VLAN with the VID 1.
- Untagged ports of deleted VLANs are automatically returned to the Default VLAN
- Tagged ports of deleted VLANs retain their tagged and untagged assignments in other VLANs.

Perform the following procedure to delete 802.1Q Tagged VLANs:

1. Select **Bridge > VLAN > Tagged VLAN** from the main menu. Refer to Figure 93 on page 212. An example of the Tagged VLAN window is shown in Figure 94 on page 213.
2. In the **VLAN Action** column of the table, click the **Delete** button of the VLAN you want to delete. You can delete only one VLAN at a time.

The switch displays a confirmation prompt.

3. Click **OK** to delete the VLAN or **Cancel** to cancel the deletion.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Viewing Port-based and 802.1Q Tagged VLANs

To view all the Port-based and 802.1Q Tagged VLANs on the switch in one window, select **Bridge > VLAN > VLAN Current Database** from the main menu. Refer to Figure 100.

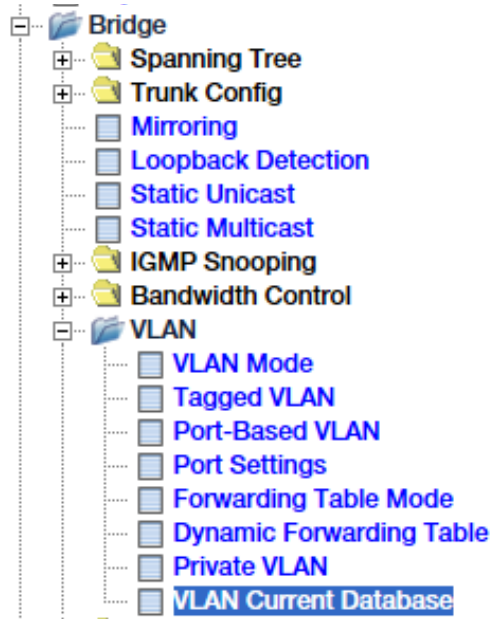


Figure 100. VLAN Current Database Menu Selection

Figure 101 is an example of the VLAN Current Database window.

VLAN Current Database					
802.1Q Tagged VLAN					
VLAN ID	VLAN Name	VLAN FDB ID	Member Ports	Untagged Ports	Status
1		1	1-28	1-28	permanent
Page 1 / 1 <input type="button" value="First Page"/> <input type="button" value="Previous Page"/> <input type="button" value="Next Page"/> <input type="button" value="Last Page"/> Page <input type="text" value=""/> <input type="button" value="Go"/>					
Port-Based VLAN					
VLAN Index	VLAN Name	VLAN Member			
< < Table is empty > >					
Page 0 / 0 <input type="button" value="First Page"/> <input type="button" value="Previous Page"/> <input type="button" value="Next Page"/> <input type="button" value="Last Page"/> Page <input type="text" value=""/> <input type="button" value="Go"/>					

Figure 101. VLAN Current Database Window

Refer to Table 43 for a description of the VLAN Current Database window

Table 43. VLAN Current Database Window

Column	Description
802.1Q Tagged VLAN	
VLAN ID	Displays the VLAN identification number.
VLAN Name	Displays the VLAN name.
VLAN FDB ID	Displays the VLAN FDB identification number.
Member Ports	Displays the tagged and untagged ports of the VLAN.
Untagged Ports	Displays only the untagged ports.
Status	Permanent (static) or dynamic.
Port-Based VLAN	
VLAN Index	Displays the index number of the VLAN.
VLAN Name	Displays the name of the VLAN.
VLAN Member	Displays the ports of the VLAN.

Chapter 17

GARP VLAN Registration Protocol

This chapter contains the following sections:

- “GVRP Overview” on page 226
- “Enabling or Disabling GVRP” on page 230
- “Configuring GVRP Port Settings” on page 231
- “Configuring GVRP Time Settings” on page 233

GVRP Overview

The GARP VLAN Registration Protocol (GVRP) allows network devices to share VLAN information and modify existing VLANs or add new VLANs, automatically. This makes managing VLANs that span multiple switches easier. Without GVRP, you have to manually configure the switches to ensure that the various parts of VLANs can communicate with each other across different switches. With GVRP, an application of the Generic Attribute Registration Protocol (GARP), this is done automatically.

The switch uses GVRP protocol data units (PDUs) to share VLAN information among GVRP-active devices. The PDUs contain the VID numbers of the VLANs on the switch. When the switch receives a GVRP PDU on a port, it examines the PDU to determine the VIDs of the VLANs on the device that sent it. It then does the following:

- ❑ If the PDU contains a VID of a VLAN that does not exist on the switch, it adds the designated VLAN and adds the port that received the PDU as a tagged member of the VLAN. A VLAN added by GVRP is called a dynamic GVRP VLAN.
- ❑ If the PDU contains a VID of a VLAN that already exists on the switch but the port is not a member of it, the switch adds the port as a tagged member of the VLAN. A port that has been added by GVRP to a static VLAN (that is, a VLAN added by the manager) is called a dynamic GVRP port.

Only GVRP can modify or delete dynamic GVRP VLANs. Dynamic GVRP VLANs exist only so long as there are active nodes in the VLANs. If all the nodes of a dynamic GVRP VLAN stop transmitting traffic and there are no active links, GVRP deletes it from the switch.

Dynamic GVRP ports in static VLANs remain members of the VLANs as long as there are active VLAN members. If all members of a VLAN become inactive or there are no active links, GVRP removes the dynamic ports from the VLAN, but does not delete the VLAN if it is a static VLAN.

Figure 102 on page 227 is an example of how GVRP works.




Figure 102. GVRP Example

The example consists of three switches. Switches #1 and #3 have the Sales VLAN, but switch #2 does not. Consequently, the end nodes of the two parts of the Sales VLANs cannot communicate with each other.

Without GVRP, you would have to manually add the Sales VLAN to switch #2. But with GVRP, the switch adds the VLAN automatically. Here is how GVRP resolves the example.

1. Port 1 on switch #1 sends to port 2 on switch #2 a PDU containing the VID of all the VLANs on the switch, including VID 11 for the Sales VLAN.
2. Switch #2 examines the PDU it receives on port 2 and notes that it does not have a VLAN with a VID 11. In response, it creates the VLAN as a dynamic GVRP VLAN, assigning it a VID 11 and the name GVRP_VLAN_11. (The name of a dynamic GVRP VLAN has the prefix "GVRP_VLAN_", followed by the VID number.) The switch then adds port 2, the port that received the PDU, as a tagged member of the VLAN.
3. Switch #2 sends a PDU from port 3 containing all the VIDs of the VLANs on the switch, including the new GVRP_VLAN_11 with its VID of 11. (Note that port 3 is not yet a member of the VLAN. Ports are added to VLANs when they receive PDUs from other network devices, not when they transmit PDUs.)
4. Switch #3 receives the PDU on port 4 and, after examining it, notes that one of the VLANs on switch #2 has the VID 11, which matches the VID of an already existing VLAN on the switch. So it does not add the VLAN because it already exists. It then determines whether the port that received the PDU, in this case port 4, is a member of the VLAN. If it is not a member, it automatically adds the port to the VLAN as a tagged dynamic GVRP port. If the port is already a member of the VLAN, then no change is made.

5. Switch #3 sends a PDU out port 4 to switch #2.
6. Switch #2 receives the PDU on port 3 and adds the port as a tagged dynamic GVRP port to the dynamic GVRP_VLAN_11 VLAN.

There is now a communications path for the end nodes of the Sales VLAN on switches #1 and #3. GVRP created the new GVRP_VLAN_11 dynamic GVRP VLAN with a VID of 11 on switch #2 and added ports 2 and 3 to the VLAN as tagged dynamic GVRP ports.

Guidelines

Here are the GVRP guidelines:

- GVRP is supported with STP or RSTP or without spanning tree.
- Both ports of a network link between switches have to be running GVRP.
- You cannot modify or delete dynamic GVRP VLANs.
- You cannot remove dynamic GVRP ports from static or dynamic VLANs.
- To be detected by GVRP, a VLAN needs to have at least one active node or have at least one port with a valid link to an end node. GVRP cannot detect VLANs that do not have any active nodes or valid port links.
- Resetting the switch erases all dynamic GVRP VLANs and dynamic GVRP port assignments. The dynamic assignments are relearned by the switch as PDUs arrive on the ports from other switches.
- GVRP has three timers: join timer, leave timer, and leave all timer. The values for these timers have to be set the same on all switches running GVRP. Timers with different values on different switches can result in GVRP compatibility problems.
- You can convert dynamic GVRP VLANs and dynamic GVRP port assignments to static VLANs and static port assignments.
- The default port settings on the switch for GVRP is active, meaning that the ports participate in GVRP. For security purposes, Allied Telesis recommends disabling GVRP on those ports that are connected to GVRP-inactive devices, meaning devices that do not feature GVRP.
- PDUs are transmitted from only those switch ports where GVRP is enabled.

GVRP and Network Security

GVRP should be used with caution because it can expose your network to unauthorized access. If a network intruder connects to a switch port running GVRP and transmits a bogus GVRP PDU containing VIDs of restricted VLANs, GVRP makes the port a member of the VLANs, giving the intruder access to restricted areas of your network.

Here are a suggestions on how to protect your network against this type of intrusion:

- ❑ Activating GVRP only on those switch ports connected to other GVRP devices. Do not activate GVRP on ports that are connected to non-GVRP devices.
- ❑ After the switches have used GVRP to form the VLANs and VLAN links, you should convert all dynamic GVRP VLANs and dynamic GVRP ports to static assignments, and then turn off GVRP on all switches. This preserves the new VLAN assignments while protecting against network intrusion.

GVRP-inactive Intermediate Switches

If two GVRP devices are separated by a non-GVRP switch, the GVRP devices may not be able to share VLAN information. There are two issues involved.

The first is whether the intermediate switch forwards GVRP PDUs that it receives from the GVRP switches. GVRP PDUs are management frames, intended for the switch's CPU. In all likelihood, a non-GVRP switch will discard the PDUs because it will not recognize them.

The second issue is that even if a non-GVRP switch forwards GVRP PDUs, it will not automatically add the VLANs. Consequently, even if GVRP switches receive the PDUs and add the necessary VLANs, an intermediate switch may block the VLAN traffic, unless you manually modify its VLANs and port assignments.

Enabling or Disabling GVRP

To enable or disable GVRP on the switch, perform the following procedure:

1. Select **Bridge > GVRP > GVRP Global Settings** from the main menu. Refer to Figure 103.

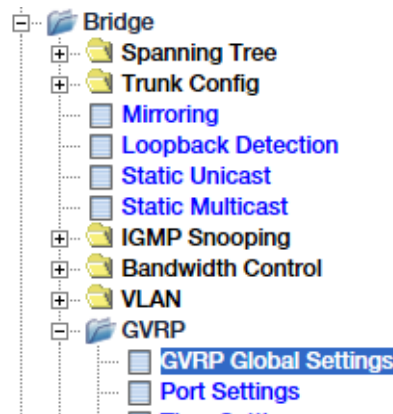


Figure 103. GVRP Global Settings Menu Selection

The GVRP Global Settings window is shown in Figure 104.




Figure 104. GVRP Global Settings Window

2. Select one of the following options from the GVRP Status menu:
 - Disabled** - Disables GVRP on the switch. This is the default setting.
 - Enabled** - Enables GVRP. The switch transmits GVRP PDUs from the ports and processes ingress PDUs.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Configuring GVRP Port Settings

Perform the following procedure to configure GVRP port settings:

1. Select **Bridge > GVRP > Port Settings** from the main menu. Refer to Figure 105.

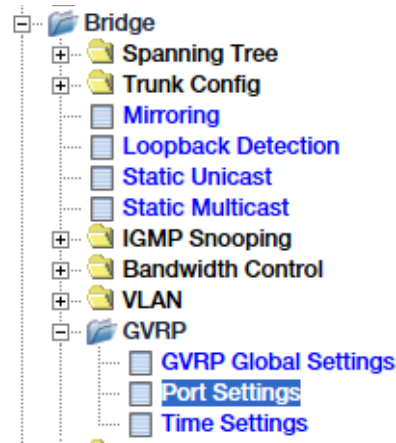


Figure 105. GVRP Port Settings Menu Selection

The GVRP Port Settings window is shown in Figure 106.

GVRP Port Settings			
Port	Dynamic VLAN Status	Restricted VLAN Registration	Action
All	Ignore ▾	Ignore ▾	Apply
1	Enabled ▾	Disabled ▾	Apply
2	Enabled ▾	Disabled ▾	Apply
3	Enabled ▾	Disabled ▾	Apply
4	Enabled ▾	Disabled ▾	Apply

Figure 106. GVRP Port Settings Window

2. Configure the fields in Table 44.

Table 44. GVRP Port Settings Window

Field	Description
Port	Lists the switch ports.

Table 44. GVRP Port Settings Window (Continued)

Field	Description
Dynamic VLAN status	Choose one of the following: <ul style="list-style-type: none"> - Enabled: Activates GVRP on a port. The switch transmits PDUs from it and processes ingress PDUs. This is the default setting. - Disabled: Deactivates GVRP on a port.
Restricted VLAN Registration	Choose one of the following: <ul style="list-style-type: none"> - Enabled: Adds a port only to static VLANs on the switch. When a port receives PDUs containing VLANs to which it is not a member, the switch adds it only if the VLANs are static VLANs. A port is not added to dynamic or unknown VLANs. - Disabled: Adds a port to static, dynamic, or unknown VLANs on the switch. When a port receives PDUs with VLANs to which it is not a member, the switch adds it to the VLANs regardless of VLAN type. This is the default setting.

3. Click the **Apply** button for the port.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Configuring GVRP Time Settings

Perform the following procedure to configure the GVRP time settings:

1. Select **Bridge > GVRP > Time Settings** from the main menu. Refer to Figure 107.

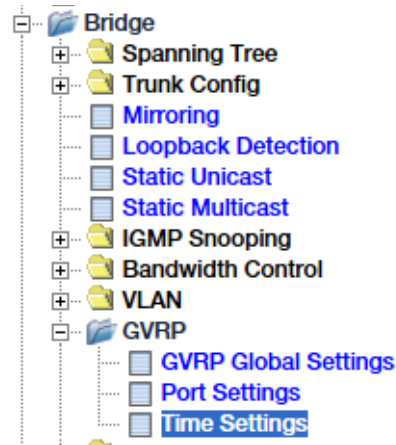


Figure 107. GVRP Time Settings Menu Selection

The GVRP Time Settings window is shown in Figure 108.

GVRP Time Settings						
Port	JoinTime(10 ~ 2 ³⁰ -14) msec	LeaveTime (30 ~ 2 ³¹ -18) msec	LeaveAllTime(40 ~ 2 ³¹ -8) msec	Action		
All	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="button" value="Apply"/>		
1	<input type="text" value="200"/>	<input type="text" value="600"/>	<input type="text" value="10000"/>	<input type="button" value="Apply"/>		
2	<input type="text" value="200"/>	<input type="text" value="600"/>	<input type="text" value="10000"/>	<input type="button" value="Apply"/>		
3	<input type="text" value="200"/>	<input type="text" value="600"/>	<input type="text" value="10000"/>	<input type="button" value="Apply"/>		
4	<input type="text" value="200"/>	<input type="text" value="600"/>	<input type="text" value="10000"/>	<input type="button" value="Apply"/>		

Figure 108. GVRP Time Settings Window

Note

You should not change the timers if you are unfamiliar with their functions. Refer to IEEE 802.1p standard for definitions.

2. Configure the fields in Table 45.

Table 45. GVRP Time Settings Window

Column	Description
Port	Lists the switch ports.
JoinTime	Enter the GARP Join Timer. The range is 10 to 1073741810 milliseconds.
LeaveTime	Enter the GARP Leave Timer. The range is 30 to 2147483630 milliseconds. The timer has to be set in relation to the GVRP Join Timer according to the following equation: $\text{GARPLeaveTimer} \geq (\text{GARPJoinTimer} \times 2) + 10$
LeaveAllTime	Enter the GARP Leave All Timer. The range is 40 to 2147483640 milliseconds. The timer must be set in relation to the GVRP Leave Timer according to the following equation: $\text{GARPLeaveAllTimer} > (\text{GARPLeaveTimer} + 10)$

Review the following:

- The GARPLeaveTimer has to be greater than $(\text{GARPJoinTimer} \times 2 + 10)$ and the GARPLeaveAllTimer must be greater than $(\text{GARPLeaveTimer} + 10)$. The acceptable input values are multiples of 10. The switch rounds down values that are not multiples of 10.
- To ensure compatibility between network devices, you have to configure the same values for the GARP Join Timer, GARP Leave Timer, and GARP Leave All Timer on all participating GVRP devices in your network.

3. Click the **Apply** button.4. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 18

Private Virtual LAN

This chapter describes the private VLAN in the following sections:

- ❑ “Private VLAN Overview” on page 236
- ❑ “Configuring the Private VLAN” on page 237
- ❑ “Disabling the Private VLAN” on page 240

Private VLAN Overview

The private VLAN creates a special broadcast domain on the switch in which the traffic of its member ports are restricted to a single source port. Ports in the VLAN are only allowed to forward traffic to and receive traffic from the designated source port, and are prohibited from forwarding traffic to each other.

An example application of the private port VLAN would be user booths in a library. The private VLAN could enable the computers in the booths to access the Internet or a library server through the source port, but would block communications between the computers.

Another application for the private port VLAN is to simplify IP address assignments. Ports can be isolated from each other while still belonging to the same subnet.

The switch supports only one private VLAN, The switch can support private, port-based, and tagged VLANs simultaneously.

Source Port

The private VLAN has one source port. The switch permits the other ports in the VLAN to forward traffic only to the source port, and not to each other. Here are the source port guidelines:

- ❑ The private VLAN can have only one source port.
- ❑ The source port can be any port on the switch.
- ❑ It cannot be a member of a static or LACP trunk.
- ❑ It cannot be a member of another VLAN.
- ❑ It is an untagged VID port and should be connected to a device that sends untagged packets.

Forwarding Ports

The other ports of the private VLAN are referred to as forwarding ports. The switch allows them to forward traffic only to the source port. Here are the guidelines:

- ❑ The private VLAN can have any number of forwarding ports, up to all the switch ports minus the source port.
- ❑ They cannot be members of static or LACP trunks.
- ❑ They cannot be members of other VLANs.
- ❑ They are untagged ports. They should be connected to devices that send untagged packets. To set the PVID, refer to “Configuring PVIDs and Filters for Tagged and Untagged Ports” on page 219.

Configuring the Private VLAN

To configure the private VLAN, perform the following procedure:

1. Select **Bridge > VLAN > Private VLAN** from the main menu. See Figure 109.

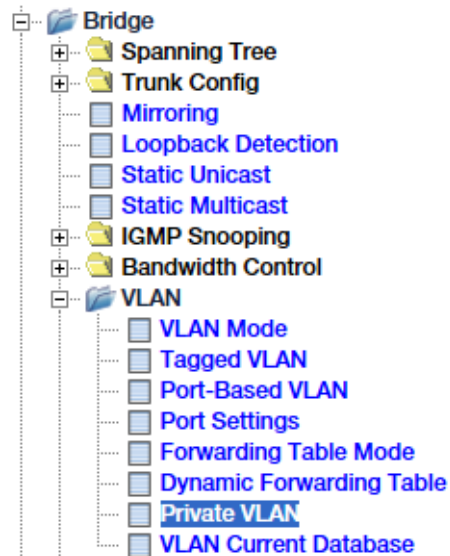


Figure 109. Private VLAN Menu Selection

The Private VLAN window in Figure 110 on page 238 shows the default settings. The feature is disabled. If enabled without any changes, port 1 is the source port and all other ports are forwarding ports of the VLAN.

Private VLAN

State: Enabled ▾

Apply

Source Port: 1 ▾

Forwarding Ports

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Clear	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Apply

Port List

Port	Port Map
1	1-28
2	1-28
3	1-28
4	1-28

Figure 110. Private VLAN Window

- Configure the private VLAN parameters in Table 46.

Table 46. Private VLAN Window

Parameter	Description
State	Select one of the following from the pull-down menu: <ul style="list-style-type: none"> - Enabled: Activates the private VLAN. You have to enable the VLAN to configure the settings. - Disabled: Disables the private VLAN. This is the default setting.
Source Port	Select the source port for the private VLAN from the pull-down menu. You can select only one source port. The default is port 1.

Table 46. Private VLAN Window (Continued)

Parameter	Description
Forwarding Ports	<p>Designate the forwarding ports of the VLAN by clicking on the check boxes. Here are the guidelines:</p> <ul style="list-style-type: none"> - Ports with check marks are forwarding ports of the VLAN. - The default is all ports except the source port are forwarding ports. - You cannot designate the source port as also a forwarding port.
Port List	<p>Lists the ports on the switch and indicates whether they are members of the private VLAN. The columns in the table are defined here:</p> <ul style="list-style-type: none"> - Port: Lists the ports on the switch. - Port Map: Indicates the source and forwarding ports of the private VLAN. The source port is listed first, followed by a comma and the forwarding ports. For example, if port 2 is the source port, and ports 6 to 13 are forwarding ports, the port map would be 2,6-13. - This column is empty for ports that are not members of the private VLAN.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Disabling the Private VLAN

To disable the private VLAN, perform the following procedure:

1. Select **Bridge > VLAN > Private VLAN** from the main menu. Refer to Figure 109 on page 237. The Private VLAN window is shown in Figure 110 on page 238.
2. Select **Disabled** from the State pull-down menu.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 19

Voice VLAN

This chapter describes the voice VLAN feature in the following sections:

- ❑ “Voice VLAN Overview” on page 242
- ❑ “General Guidelines” on page 244
- ❑ “Configuring the Voice VLAN” on page 245
- ❑ “Managing the OUI Table” on page 248

Voice VLAN Overview

The voice VLAN is intended for Voice over IP (VoIP) phones. It enables the switch to provide high-quality, uninterrupted voice traffic from phone users on its ports. The voice VLAN provides these basic functions:

- ❑ Identify IP phone traffic on an 802.1Q tagged VLAN by Organization Unique Identifiers (OUIs).
- ❑ Assign traffic a high Class of Service value.
- ❑ Direct traffic to a high traffic queue on the egress port.

These actions help minimize interruptions to phone conversations and increases voice quality.

802.1Q VLAN

The first step to implementing the voice VLAN is to add an 802.1Q tagged VLAN to the switch. The VLAN will form the base of the voice VLAN. The VLAN has to have one or more tagged or untagged ports that will serve as the voice VLAN ports. The VLAN has an upstream port leading to the phone system and downstream ports leading to the individual IP phones or other voice VLAN network nodes, such as other Ethernet switches or a DHCP server. The voice VLAN can have only one 802.1Q VLAN. Refer to Chapter 16, “802.1Q Tagged Virtual LANs” on page 205.

Organization Unique Identifier (OUI)

The switch identifies IP phone traffic by examining the OUIs in the source MAC addresses of the packets. Each IP phone manufacturer has one or more unique OUIs. An OUI is three bytes long and is usually expressed in hexadecimal format. It is incorporated into the first part of the MAC address of an Ethernet network device. You can find the OUI of an IP phone in the first three complete bytes of its MAC address. IP phones from the same manufacturer typically have the same OUI.

To identify voice data packets, the switch compares the OUI information in the source MAC addresses in the packets against entries in its OUI table. To configure the table, you enter the complete MAC address of one of your IP phones. An “OUI Mask” is automatically generated and applied by the switch to produce the manufacturer’s OUI. The switch supports up to ten OUI entries.

If all your phones have the same OUI, then you only have to enter the MAC address from one phone in the OUI table. If your phones have more than one OUI, then you have to enter one MAC address for each OUI.

CoS with Voice VLAN

The Voice VLAN has a CoS parameter to maintain the voice quality between the ingress and egress ports of the switch. You have to enable CoS for the Voice VLAN CoS priority to take effect. The CoS priority level you designate is applied to voice traffic on all ports of the voice VLAN.

Normally, most (non-Voice) Ethernet traffic traverses the switch through lower-order egress queues. To avoid delays and interruptions in the voice data flow, you should assign a high-order queue to the CoS priority level of the voice VLAN, and set the scheduling algorithm to Strict Priority. These settings ensure that the voice data packets are processed before other types of data so that the voice quality is maintained as the switch forwards voice data. For more information, refer to “Mapping CoS Priorities to Egress Queues” on page 304 and “Setting the Queue Scheduling Algorithm” on page 309.

Dynamic Port Auto-Detection

The switch has two ways to learn which ports are connected to IP phones in the voice VLAN. One way is for you to identify the ports yourself when adding the 802.1Q tagged VLAN that serves as the base of the voice VLAN. Another way is to let the switch discover them with its auto-detection feature. When the switch detects an ingress packet with an OUI on an ingress port that is not already part of the voice VLAN, it moves the port into the VLAN. The switch can also remove ports from the voice VLAN after a predefined timeout period, when IP phones are inactive.

Here are the guidelines for the auto-detection feature:

- ❑ Ports cannot be members of any VLAN. They have to be designated as “Not Member” ports in all VLANs. The switch will not examine packets for OUIs on ports that are members of other VLANs. Refer to Chapter 15, “Port-Based Virtual LANs” on page 183 and Chapter 16, “802.1Q Tagged Virtual LANs” on page 205.
- ❑ IP phones have to support 802.1Q VID packets with imbedded VLAN ID tags. You have to manually set their VLAN IDs to be the same as the voice VLAN ID. This enables the switch to identify the packets from the IP phones as part of the voice VLAN, and so automatically move the ports into the VLAN. For example, if the voice VLAN has the VID 20, you have to set the VIDs on the IP phones to 20.
- ❑ Auto-detection is not supported on switch ports connected to IP phones that do not support 802.1Q VIDs. Those ports have to be configured as static tagged ports in the voice VLAN.

Note

The port VLAN IDs (PVID) of static tagged members of the voice VLAN have to be the same as the voice VLAN ID. This is to insure that all untagged packets entering the ports are switched within the voice VLAN as the voice data pass through the switch. Refer to “Configuring PVIDs and Filters for Tagged and Untagged Ports” on page 219.

Note

The switch does not support Link Layer Discovery Protocol for Media Endpoint Devices (LLDP-MED). It cannot set the VIDs on IP phones that support 802.1Q VID.

General Guidelines

Here are the voice VLAN guidelines:

- ❑ The switch supports one voice VLAN.
- ❑ The voice VLAN has to be an 802.1Q tagged VLAN. The voice VLAN is not supported on port-based VLANs or the private VLAN.
- ❑ The switch port connected to the phone system at the network core has to be a static member of the voice VLAN.
- ❑ You have to configure the port VLAN IDs (PVID) of static tagged members of the voice VLAN to be the same as the voice VLAN ID. Refer to “Configuring PVIDs and Filters for Tagged and Untagged Ports” on page 219.
- ❑ You have to enter the OUIs of the IP phones into the OUI table. Refer to “Managing the OUI Table” on page 248.
- ❑ Only one MAC address for each unique OUI is required.
- ❑ The switch supports up to ten OUIs.
- ❑ The switch does not support Link Layer Discovery Protocol for Media Endpoint Devices (LLDP-MED).

Configuring the Voice VLAN

The voice VLAN requires an 802.1Q tagged VLAN. If you have not yet added the VLAN to the switch, perform “Adding or Viewing 802.1Q Tagged VLANs” on page 212.

Perform the following procedure to configure a voice VLAN on the switch:

1. Select **Voice VLAN > Voice VLAN Settings** from the main menu. Refer to Figure 111.




Figure 111. Voice VLAN Settings Menu Selection

The Voice VLAN Setting window is shown in Figure 112.

Voice VLAN Settings

Voice VLAN Disabled ▾

Note: Disabling will turn off the function and return all values to default.

Voice VLAN Global Settings

VLAN ID: ▾

Aging Time: (1-120 hours)

CoS: 7 ▾

Port	Auto Detection	Status	Action
All	Ignore ▾	-	<input type="button" value="Apply"/>
1	Disabled ▾	None	<input type="button" value="Apply"/>
2	Disabled ▾	None	<input type="button" value="Apply"/>
3	Disabled ▾	None	<input type="button" value="Apply"/>
4	Disabled ▾	None	<input type="button" value="Apply"/>

Figure 112. Voice VLAN Settings Window

2. Configure the settings in Table 47.

Table 47. Voice VLAN Settings Window

Field	Description
Voice VLAN	Select one of the following from the menu: - Enabled: Enables voice VLAN on the switch. You have to enable voice VLAN to configure its settings. - Disabled: Disables the voice VLAN. This is the default setting.
VLAN ID	Select the VID of the 802.1Q tagged VLAN to be the voice VLAN. You can select only one VLAN. If you have not added the VLAN, go to “Adding or Viewing 802.1Q Tagged VLANs” on page 212.
Aging Time	Enter the amount of elapsed time, in hours, after which the switch removes ports with inactive IP phones from the voice VLAN. This applies only to ports discovered with auto-detection. For example, when the Aging Time is set to 1 hour, the default setting, the switch removes ports from the voice VLAN when IP phones discovered with auto-detection are inactive for one hour. The range is 1 to 120 hours.
CoS	Select from the menu the CoS priority level the switch is to assign to voice data packets received on voice VLAN ports. The range is 0 to 7. The default is 7. You need to enable CoS on the switch and map this CoS value to the CoS “Highest” egress queue on the ports. Refer to “Mapping CoS Priorities to Egress Queues” on page 304. You should also set the scheduling algorithm to Strict Priority, as explained in “Setting the Queue Scheduling Algorithm” on page 309.

3. Click the **Apply** button.

To configure ports for auto-detection, explained in “Dynamic Port Auto-Detection” on page 243, continue with the next step. Otherwise, go to step 6.

4. To enable or disable auto-detection, select one of the following from the **Auto Detection** column for a port in the table:
 - Enabled** - Enables the voice VLAN auto-detection feature on the port.
 - Disabled** - Disables the auto-detection feature.
5. Click the **Apply** button in the **Action** column.
6. Select **Save Settings to Flash** from the main menu to save your changes.

Managing the OUI Table

The OUI table contains the OUIs of your IP phones. The switch uses the table to identify voice packets on its ports. You have to enter the MAC address from one phone for each unique OUI. Refer to “Organization Unique Identifier (OUI)” on page 242.

Adding OUI Entries

Perform the following procedure to add up to ten OUI entries of phone manufacturers to the table:

1. Select **Voice VLAN > Voice VLAN OUI Settings** from the main menu. Refer to Figure 113.




Figure 113. Voice VLAN OUI Settings Menu Selection

The Voice VLAN OUI Settings window is shown in Figure 114.

Voice VLAN OUI Settings

Description

Telephony OUI : : : : : (e.g. XX:XX:XX:XX:XX:XX)

(Maximum user defined OUI : 10)

Free Policies : 188

ID	Description	Telephony OUI	OUI Mask	Action
<< Table is empty >>				

Figure 114. Voice VLAN OUI Settings Window

2. Configure the settings in Table 48.

Table 48. Voice VLAN OUI Settings Window

Field	Description
Description	Enter the name of the phone manufacturer and phone model. The description can be up to 20 characters.
Telephony OUI	Enter the MAC address of one of the phones.

3. Click the **Add** button.

4. To add more OUI entries, repeat steps 2 and 3. The switch supports ten entries.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying OUI Entries

You cannot modify OUIs in the table. If you need to change an OUI, you have to delete it and enter it again.

Deleting OUI Entries

Perform the following procedure to delete OUI entries:

1. Select **Voice VLAN > Voice VLAN OUI Settings** from the main menu. Refer to Figure 113 on page 248 and Figure 114 on page 248.
2. Click the entry's **Delete** button in the Action column of the table.
3. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 20

MAC Address Forwarding Table

This chapter describes the MAC address forwarding table in the following sections:

- ❑ “Viewing the MAC Addresses of Network Devices in the Forwarding Table” on page 252
- ❑ “Selecting the MAC Address Forwarding Table Mode” on page 254

Viewing the MAC Addresses of Network Devices in the Forwarding Table

This procedure explains how to view the MAC addresses the switch has stored in its forwarding table. You can view all the addresses or the addresses of specific ports.

Perform the following procedure to view the MAC addresses in the forwarding table:

1. Select **Bridge > VLAN > Dynamic Forwarding Table** in the main menu. Refer to Figure 115.

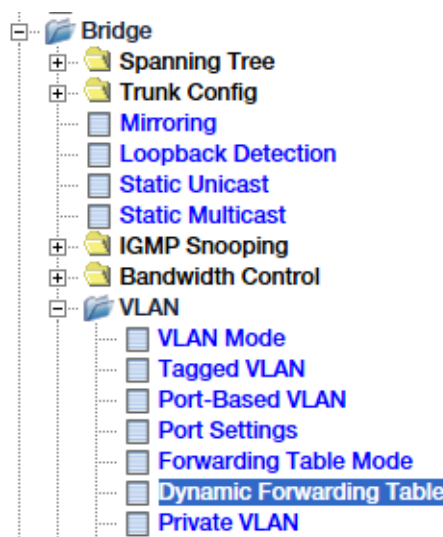


Figure 115. Dynamic Forwarding Table Menu Selection

An example of the Dynamic Forwarding Table window is shown in Figure 116.

Dynamic Forwarding Table

Port: All ▾

Refresh

Dynamic Forwarding Table

Index	VID	Port	MAC Address	Type
1	34	2	nn - nn - nn - nn - nn - nn	Dynamic
2	34	3	nn - nn - nn - nn - nn - nn	Dynamic
3	34	4	nn - nn - nn - nn - nn - nn	Dynamic

Figure 116. Dynamic Forwarding Table Window

The table columns are described in Table 49.

Table 49. Dynamic Forwarding Table Window

Column	Description
Index	Displays the index number of the address in the table.
VID	Displays one of the following: <ul style="list-style-type: none"> - If the MAC address is part of an 802.1Q Tagged VLAN, the column displays the VID. - If the MAC address is part of a Port-based VLAN, the column displays 4094.
Port	Displays the port or trunk number. Port numbers with a “p” prefix indicate port trunks. For example, “p2” indicates port trunk 2.
Type	Displays one of the following: <ul style="list-style-type: none"> - Dynamic: The switched learned the MAC address. - Static: The MAC address was entered manually.

2. To view all the MAC addresses, select **All** from the Port menu. To view MAC addresses of a specific port, select the port from the Port menu.

Selecting the MAC Address Forwarding Table Mode

The switch supports the following modes for learning MAC addresses on VLANs:

- ❑ Independent VLAN learning (IVL): The switch maintains separate MAC address forwarding tables for the VLANs. This is the default setting.
- ❑ Shared VLAN learning (SVL): The switch maintains one MAC address forwarding table for all the VLANs.

For more information, refer to the IEEE 802.1Q standard.

Perform the following procedure to select the MAC address learning mode:

1. Select **Bridge.> VLAN.> Forwarding Table Mode** from the main menu. Refer to Figure 117.




Figure 117. Forwarding Table Mode Menu Selection

The Forwarding Table Mode window is shown in Figure 118.



Figure 118. Forwarding Table Mode Window

2. Select **IVL** or **SVL** from the **Learning Mode** menu:

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Section IV

Network Traffic Management

This section contains the following chapters:

- ❑ Chapter 21, “Traffic Statistics and Charts” on page 259
- ❑ Chapter 22, “Static Unicast MAC Addresses” on page 269
- ❑ Chapter 23, “Static Multicast MAC Addresses” on page 279
- ❑ Chapter 24, “IGMP Snooping” on page 287
- ❑ Chapter 25, “Quality of Service and Class of Service” on page 299
- ❑ Chapter 26, “Access Control Policies” on page 313
- ❑ Chapter 27, “Bandwidth Control” on page 337
- ❑ Chapter 28, “RMON” on page 345

Chapter 21

Traffic Statistics and Charts

This chapter describes traffic charts in the following sections:

- ❑ “Network Statistics Overview” on page 260
- ❑ “Displaying Port Traffic Comparison Statistics” on page 261
- ❑ “Displaying Port Error Statistics” on page 265
- ❑ “Displaying Historical Statistics” on page 266

Network Statistics Overview

Statistics provide important information for troubleshooting switch problems at the port level. The switch provides a versatile set of statistics charts you can customize for your needs, including (depending upon the chart) the ports whose statistics you want to view and the chart colors.

Note

Your web browser has to have the Java SSV Helper plug-in to display the network traffic statistics charts.

There are three types of statistics charts:

- ❑ **Traffic Comparison:** This chart displays traffic statistics per port. You can select from 24 statistics and 12 chart colors. Refer to “Displaying Port Traffic Comparison Statistics” on page 261.
- ❑ **Error Group:** This chart displays discard and error counts for ports. Refer to “Displaying Port Error Statistics” on page 265.
- ❑ **Historical Status:** This chart displays the historical records of port statistics. Refer to “Displaying Historical Statistics” on page 266.

Displaying Port Traffic Comparison Statistics

The Traffic Comparison chart displays traffic statistics per port. You can select from 24 statistics and 12 colors.

Perform the following procedure to display traffic comparison statistics:

1. Select **Statistics Chart > Traffic Comparison** from the main menu. Refer to Figure 119.



Figure 119. Traffic Comparison Menu Selection

The Traffic Comparison Chart window is shown in Figure 120.

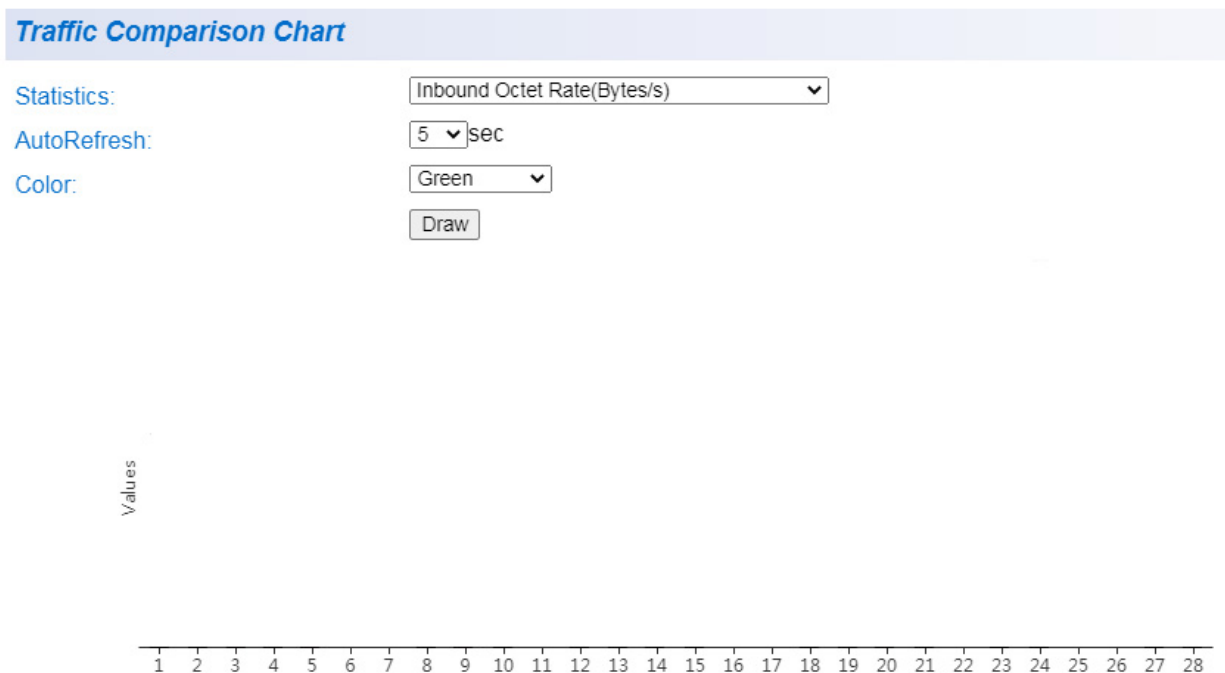


Figure 120. Traffic Comparison Chart Window

2. Select a traffic statistic from the **Statistics** menu. Refer to Table 50 on page 262.

Table 50. Traffic Comparison Chart Statistics

Statistics	Definition
Inbound Octet Rate (Bytes/s)	Displays the rate of inbound octet bits in bytes per second.
Inbound Unicast Packet Rate (Pkts/s)	Displays the rate of inbound unicast packets in packets per second.
Inbound Non-unicast Packet Rate (Pkts/s)	Displays the rate of inbound non-unicast packets in packets per second.
Inbound Discard Rate (Pkts/s)	Displays the rate of inbound discarded packets in packets per second.
Inbound Error Rate (Pkts/s)	Displays the rate of inbound errors in packets per second.
Outbound Octet Rate (Bytes/s)	Displays the rate of outbound octet bits in bytes per second.
Outbound Unicast Packet Rate (Pkts/s)	Displays the rate of outbound unicast packets in packets per second.
Outbound Non-unicast Packet Rate (Pkts/s)	Displays the rate of outbound non-unicast packets in packets per second.
Outbound Discard Rate (Pkts/s)	Displays the rate of outbound discarded packets in packets per second.
Outbound Error Rate (Pkts/s)	Displays the rate of outbound errors in packets per second.
Ethernet Undersize Packet Rate (Pkts/s)	Displays the rate of undersized Ethernet packets in packets per second.
Ethernet Oversize Packet Rate (Pkts/s)	Displays the rate of oversized Ethernet packets in packets per second.
Inbound Octets (Bytes)	Displays the number of inbound octet bits in bytes per second.
Inbound Unicast Packets (Pkts)	Displays the number of inbound unicast packets in packets per second.
Inbound Non-unicast Packets (Pkts)	Displays the number of inbound non-unicast packets (such as broadcast and multicast packets) in packets per second.
Inbound Discards (Pkts)	Displays the number of inbound discarded packets in packets per second.

Table 50. Traffic Comparison Chart Statistics (Continued)

Statistics	Definition
Inbound Errors (Pkts)	Displays the number of inbound packets with errors per second.
Outbound Octets (Bytes)	Displays the number of outbound octet bits in bytes per second.
Outbound Unicast Packets (Pkts)	Displays the number of outbound unicast packets in packets per second.
Outbound Non-unicast Packets (Pkts)	Displays the number of outbound non-unicast (such as broadcast and multicast packets) packets in packets per second.
Outbound Discards (Pkts)	Displays the number of outbound discarded packets in packets per second.
Outbound Errors (Pkts)	Displays the number of outbound error packets in packets per second.
Ethernet Undersize Packets (Pkts)	Displays the number of undersized Ethernet packets in packets per second.
Ethernet Oversize Packets (Pkts)	Displays the number of oversized Ethernet packets in packets per second.

3. Select a refresh time from the **AutoRefresh** menu. The range is 5 to 30 seconds. The default is 5 seconds.
4. Select a chart color from the **Color** menu. Refer to Table 51. The default is green.

Table 51. Graph Colors

Green	Purple	Gray	Light Green
Blue	Yellow	Light Red	Light Yellow
Red	Orange	Light Blue	Light Gray

5. Click **Draw** to generate the Traffic Comparison chart. Refer to Figure 121 for an example.

Traffic Comparison Chart

Statistics:

AutoRefresh: sec

Color:




Figure 121. Traffic Comparison Window Example

6. Select **Save Settings to Flash** from the main menu to save your changes.

Displaying Port Error Statistics

Perform the following procedure to display port error statistics:

1. Select **Statistics Chart > Error Group** from the main menu. Refer to Figure 122.




Figure 122. Error Group Menu Selection

The Error Group Chart window is shown in Figure 123.




Figure 123. Error Group Chart Window

2. Select a port number from the **Port** menu.
3. Select a refresh time from the **AutoRefresh** menu. The range is 5 to 30 seconds. The default is 5 seconds.
4. Select a chart color from the **Color** menu. Refer to Table 51 on page 263. The default is green.
5. Click the **Draw** button to generate the chart.
6. Select **Save Settings to Flash** from the menu to save your changes.

Displaying Historical Statistics

Perform the following procedure to display port historical statistics:

1. Select **Statistics Chart > Historical Status** from the main menu. Refer to Figure 124.




Figure 124. Historical Status Menu Statistics

The Historical Status Chart window is shown in Figure 125.

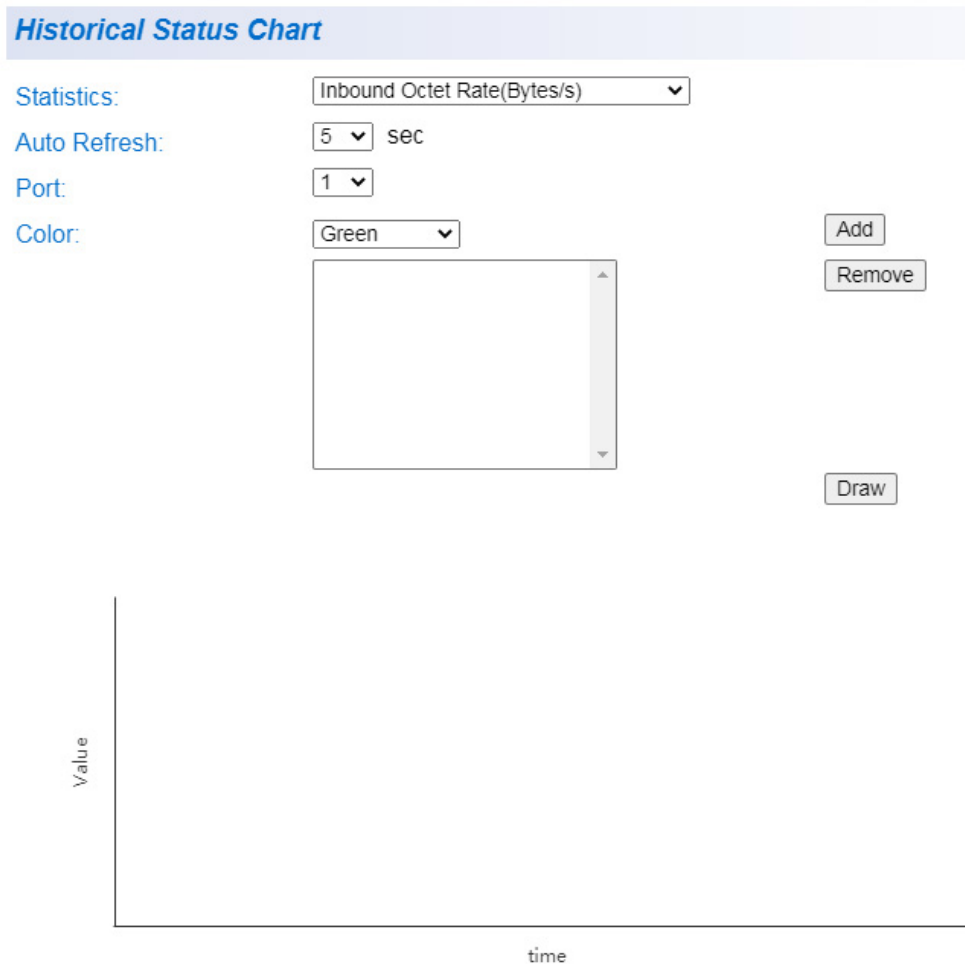


Figure 125. Historical Status Chart Window

2. Select a statistics from the **Statistics** menu. Refer to Table 52.

Table 52. Statistics in the Historical Status Chart Window

Statistics	Definition
Inbound Octet Rate (Bytes/second)	Displays the number of inbound octet bits in bytes per second.
Inbound Unicast Packets (Pkts)	Displays the number of inbound unicast packets in packets per second.
Inbound Non-unicast Packets (Pkts)	Displays the number of inbound non-unicast packets (such as broadcast and multicast packets) in packets per second.
Inbound Discards (Pkts)	Displays the number of inbound discarded packets in packets per second.
Inbound Errors (Pkts)	Displays the number of inbound errors in packets per second.
Outbound Octets (Bytes)	Displays the number of outbound octet bits in bytes per second.
Outbound Unicast Packets (Pkts)	Displays the number of outbound unicast packets in packets per second.
Outbound Non-unicast Packets (Pkts)	Displays the number of outbound non-unicast (such as broadcast and multicast packets) packets.
Outbound Discards (Pkts)	Displays the number of outbound discarded packets.
Outbound Errors (Pkts)	Displays the number of outbound error packets.
Ethernet Undersize Packets (Pkts)	Displays the number of undersized Ethernet packets.
Ethernet Oversize Packet Rate (Pkts)	Displays the number of oversized Ethernet packets.

3. Select a refresh time from the **AutoRefresh** menu. The range is 5 to 30 seconds. The default is 5 seconds.
4. Select a port from the **Port** menu.
5. Select a chart color from the **Color** menu. Refer to Table 51 on page 263. The default is green.
6. Click the **Add** button.

7. Repeat steps 4 to 6 to add more ports.
8. Click the **Draw** button to start the Historical Status Chart.
9. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 22

Static Unicast MAC Addresses

This chapter describes static unicast MAC addresses in the following sections:

- ❑ “Static MAC Addresses Overview” on page 270
- ❑ “Adding Static Unicast MAC Addresses” on page 272
- ❑ “Modifying Static Unicast MAC Addresses” on page 275
- ❑ “Deleting Static Unicast MAC Addresses” on page 278

Static MAC Addresses Overview

The switch has a MAC address table for storing the MAC addresses of the network devices connected to its ports. Each entry in the table consists of a MAC address, a port number where an address was learned by the switch, and an ID number of a VLAN where a port is a member.

The switch learns the MAC addresses of the network devices by examining the source addresses in the packets as they arrive on the ports. When the switch receives a packet that has a source address that is not already in the table, it adds the address, along with the port number where the packet was received and the ID number of the VLAN where the port is a member. The result is a table that contains the MAC addresses of all the network devices that are connected to the switch's ports.

The purpose of the table is to allow the switch to forward packets more efficiently. When a packet arrives on a port, the switch examines the destination address in the packet and refers to its MAC address table to determine the port where the destination node of that address is connected. It then forwards the packet to that port and on to the network device.

If the switch receives a packet with a destination address that is not in the MAC address table, it floods the packet to all its ports, excluding the port where the packet was received. If the ports are grouped into virtual LANs, the switch floods the packet only to those ports that belong to the same VLAN from which the packet originated. This prevents packets from being forwarded to inappropriate LAN segments and increases network security. When the destination node responds, the switch adds its MAC address and port number to the MAC address table.

If the switch receives a packet with a destination address that is on the same port on which the packet was received, it discards the packet without forwarding it on to any port. There is no reason for the switch to forward the packet because the source and destination nodes are located on the same port on the switch. This too increases network performance by preventing frames from being forwarded unnecessarily to other network devices.

MAC addresses learned by the switch are referred to as dynamic addresses. Dynamic MAC addresses are not stored indefinitely in the MAC address table. They are automatically deleted when they are inactive. A MAC address is considered inactive if the switch does not receive any frames from the network device after a specified period of time. The switch assumes that the node with that MAC address is no longer active and that its MAC address can be deleted from the table. This prevents the MAC address table from becoming filled with addresses of inactive nodes.

The period of time the switch waits before deleting inactive dynamic MAC addresses is called the aging time. This value is adjustable on the switch. The default value is 300 seconds (5 minutes).

You can also enter addresses manually into the table. These addresses are referred to as static addresses. Static MAC addresses remain in the table indefinitely and are never deleted, even when the network devices are inactive. Static MAC addresses are useful for addresses that the switch might not learn through its normal learning process or for addresses that you want the switch to retain, even when the end nodes are inactive.

Adding Static Unicast MAC Addresses

The procedure in this section explains how to add static unicast MAC addresses to the MAC address tables for 802.1Q tagged and port-based VLANs. New static unicast MAC addresses require the following information:

- ❑ What is the static unicast MAC address?
- ❑ For a static MAC address for a port in a 802.1Q Tagged VLAN, what is the VID of the VLAN.
- ❑ Which switch port is the device connected to?

Here are the guidelines:

- ❑ You can assign a static unicast MAC address to only one switch port.
- ❑ Before assigning a MAC address to a port in an 802.1Q tagged VLAN, you have to add the VLAN to the switch and set the VLAN mode of the port to 802.1Q Tagged VLAN. Refer to “Adding or Viewing 802.1Q Tagged VLANs” on page 212 and “Assigning the Port-Based VLAN Mode to Ports” on page 193.
- ❑ Before assigning a MAC address to a port in a port-based VLAN, you have to add the VLAN to the switch and set the VLAN mode of the port to Port-based VLAN. Refer to “Adding or Viewing Port-Based VLANs” on page 190 and “Assigning the Port-Based VLAN Mode to Ports” on page 193.
- ❑ The switch has a maximum capacity of 255 static unicast addresses.

Perform the following procedure to add static unicast addresses to the switch:

1. Select **Bridge -> Static Unicast** from the main menu. Refer to Figure 126.

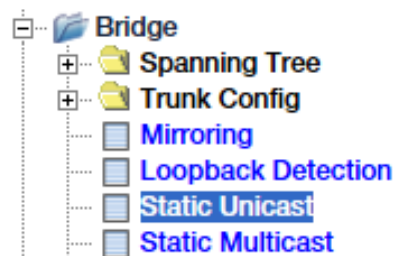


Figure 126. Static Unicast Menu Selection

The Static Unicast Address Table window is shown in Figure 127.

Static Unicast Address Table

802.1Q VLAN: (1-4093)

Port-Based VLAN

MAC Address: : : : : :

Port Member Settings

1	2	3	4	5	6	7	8	9	10	11	12	13	14
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
15	16	17	18	19	20	21	22	23	24	25	26	27	28
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

802.1Q VLAN (Free Entries: 256, Total Entries: 0)

VLAN ID	MAC Address	Port Members	Action
<< Table is empty >>			

Page 0 / 0 Page

Port-Based VLAN (Free Entries: 256, Total Entries: 0)

VLAN Index	MAC Address	Port Members	Action
<< Table is empty >>			

Page 0 / 0 Page

Figure 127. Static Unicast Address Table Window

The two tables in the bottom of the window display the current static unicast MAC addresses:

- Static Unicast Table: Displays the static MAC addresses of network devices connected to ports in 802.1Q tagged VLANs.
- Port-based VLAN. Displays the static MAC addresses of network devices connected to ports in port-based VLANs.

The columns in the tables are described in Table 53.

Table 53. Static Unicast MAC Address Tables

Column	Description
VLAN ID	Displays the VLAN ID of the 802.1Q tagged VLAN of the MAC address.

Table 53. Static Unicast MAC Address Tables (Continued)

Column	Description
VLAN Index	Displays N/A. The VLAN Index numbers of Port-based VLANs do not apply to static unicast MAC addresses.
MAC Address	Displays the static MAC address.
Port Members	Displays the switch port number of the address.
Action	Contains Modify and Delete buttons.

2. To add a static unicast MAC address to an 802.1Q Tagged VLAN, do the following:
 - a. Click the **802.1Q VLAN** radio button. This is the default setting.
 - b. Enter the VID of the tagged VLAN containing the port to which the corresponding network device is connected. You can enter only one VID. The VLAN has to already exist on the switch.
 - c. Go to step 4.
3. To add a static unicast MAC address to a Port-based VLAN, do the following:
 - a. Click the **Port-Based VLAN** radio button.
 - b. Go to step 4.
4. Enter the static unicast MAC address in the **MAC Address** fields. You can enter only one address.
5. Designate the port of the MAC address by clicking the corresponding radio button in Port Member Settings. You can select only one port.

You have to select a port whose VLAN mode matches the type of VLAN of the static unicast MAC address. Radio buttons of ports whose VLAN mode does not match the VLAN type will be grayed out. For example, when adding static unicast MAC addresses to 802.1Q Tagged VLANs, you can only select ports set to the 802.1Q Tagged VLAN mode, the default mode. Or, when adding static unicast MAC addresses to Port-based VLANs, you can only select ports set to the Port-based VLAN mode. Refer to “Assigning the Port-Based VLAN Mode to Ports” on page 193.

6. Click the **Apply** button.
7. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying Static Unicast MAC Addresses

This procedure explains how to change the port assignments of unicast MAC addresses. Here are the guidelines:

- ❑ You cannot change the VLAN or MAC address of a static address with this procedure. Instead, you have to delete the entry from the switch and enter it again with the new VLAN or MAC address. Refer to “Deleting Static Unicast MAC Addresses” on page 278.
- ❑ For you to assign a MAC address to a port in an 802.1Q Tagged VLAN, the VLAN has to already exist on the switch and the VLAN mode of the port has to be set to 802.1Q Tagged VLAN. Refer to “Adding or Viewing 802.1Q Tagged VLANs” on page 212 and “Assigning the 802.1Q Tagged VLAN Mode to Ports” on page 216.
- ❑ For you to assign a MAC address to a port in a port-based VLAN, the VLAN has to already exist on the switch and the VLAN mode of the port has to be set to Port-based VLAN. Refer to “Adding or Viewing Port-Based VLANs” on page 190 and “Assigning the Port-Based VLAN Mode to Ports” on page 193.

Perform the following procedure to modify the port assignment of a unicast MAC address:

1. Select **Bridge** -> **Static Unicast** from the main menu. Refer to Figure 126 on page 272. The Static Unicast Address Table window is shown in Figure 127 on page 273.
2. Click the **Modify** button on the Action column of the static MAC address you want to change. You can modify only one address at a time. The Modify Static Unicast Address window changes to display the details of the address. The example in Figure 128 on page 276 is for a MAC address assigned to an 802.1Q Tagged VLAN.

Port-Based VLAN

MAC Address: : : : : :

Port Member Settings

1	2	3	4	5	6	7	8	9	10	11	12	13
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
15	16	17	18	19	20	21	22	23	24	25	26	27
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

802.1Q VLAN (Free Entries: 256, Total Entries: 0)

VLAN ID	MAC Address	Port Members	Action
<< Table is empty >>			

Page 0 / 0 Page

Port-Based VLAN (Free Entries: 256, Total Entries: 0)

VLAN Index	MAC Address	Port Members	Action
<< Table is empty >>			

Figure 128. Modifying a Static Unicast Address in an 802.1Q Tagged VLAN

The example in Figure 129 is for a MAC address assigned to a Port-based VLAN.

Static Unicast Address Table

802.1Q VLAN: (1-4093)

Port-Based VLAN

MAC Address: : : : : :

Port Member Settings

1	2	3	4	5	6	7	8	9	10	11	12	13	14
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
15	16	17	18	19	20	21	22	23	24	25	26	27	28
<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

802.1Q VLAN (Free Entries: 256, Total Entries: 0)

VLAN ID	MAC Address	Port Members	Action
<< Table is empty >>			

Page 0 / 0 Page

Port-Based VLAN (Free Entries: 256, Total Entries: 0)

VLAN Index	MAC Address	Port Members	Action
<< Table is empty >>			

Figure 129. Modifying a Static Unicast Address in a Port-Based VLAN

The Port-Based VLAN field displays N/A because the Port-based VLAN Index number does not apply to static multicast addresses.

3. Select the new port of the MAC address by clicking the corresponding radio button in Port Member Settings. You can select only one port.

You can only select a port whose VLAN mode matches the type of VLAN of the static unicast MAC address. Radio buttons of ports whose VLAN mode does not match the VLAN type will be grayed out. For example, when adding a static unicast MAC addresses to an 802.1Q VLAN, you can only select a port set to the 802.1Q Tagged VLAN mode, the default mode. Or, when adding a static unicast MAC address to a Port-based VLAN, you can select only a port set to the Port-based VLAN mode. Refer to “Assigning the Port-Based VLAN Mode to Ports” on page 193.

4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Deleting Static Unicast MAC Addresses

Perform the following procedure to delete static MAC addresses from the switch:

1. Select **Bridge -> Static Unicast** from the main menu. Refer to Figure 126. The Static Unicast Address Table window is shown in Figure 127.
2. Do one of the following:
 - To delete a single entry, click its **Delete** button in the Action column.
 - To delete all 802.1q address entries, click the **Delete All** button above the 802.1Q VLAN table.
 - To delete all port-based VLAN address entries, click the **Delete All** button above the Port-Based VLAN table.
3. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 23

Static Multicast MAC Addresses

Static multicast addresses are described in the following sections:

- ❑ “Adding Static Multicast MAC Addresses” on page 280
- ❑ “Modifying Static Multicast MAC Addresses” on page 283
- ❑ “Deleting Static Multicast Addresses” on page 286

Adding Static Multicast MAC Addresses

This section contains the procedure for adding static multicast MAC addresses to the ports on the switch. The switch supports up to 256 static multicast addresses. Static multicast MAC addresses require the following information:

- ❑ What is the static multicast MAC address?
- ❑ For a static multicast MAC address for ports in a 802.1Q tagged VLAN, what is the VID of the VLAN.
- ❑ What are the switch ports that are connected to the host nodes and multicast router of the multicast address?

Here are the guidelines:

- ❑ You assign static multicast MAC addresses to hosts ports and multicast router ports.
- ❑ To assign a multicast MAC address to ports in an 802.1Q tagged VLAN, you have to add the VLAN to the switch and set the VLAN modes of the ports to 802.1Q Tagged VLAN. Refer to “Adding or Viewing 802.1Q Tagged VLANs” on page 212 and “Assigning the Port-Based VLAN Mode to Ports” on page 193.
- ❑ To assign a multicast MAC address to ports in a port-based VLAN, you have to add the VLAN to the switch and set the VLAN modes of the ports to Port-based VLAN. Refer to “Adding or Viewing Port-Based VLANs” on page 190 and “Assigning the Port-Based VLAN Mode to Ports” on page 193.
- ❑ The switch has a maximum capacity of 255 static multicast addresses.

Perform the following procedure to add static multicast addresses to the switch:

1. Select **Bridge** -> **Static Multicast** from the main menu. Refer to Figure 130.

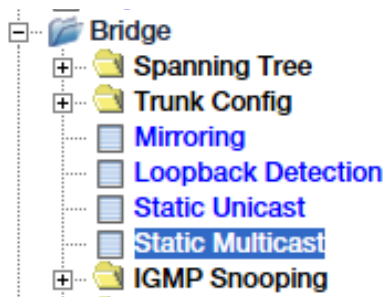


Figure 130. Static Multicast Menu Selection

The Static Multicast Address Table window is shown in Figure 131.

Static Multicast Address Table

802.1Q VLAN: (1-4093)

Port-Based VLAN

Group MAC Address: : : : : :

Group Member:

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

802.1Q VLAN (Free Entries: 256, Total Entries: 0)

VLAN ID	MAC Address	Port Members	Action
<< Table is empty >>			

Page 0 / 0 Page

Port-Based VLAN (Free Entries: 256, Total Entries: 0)

VLAN Index	MAC Address	Port Members	Action
<< Table is empty >>			

Page 0 / 0 Page

Figure 131. Static Multicast Address Window

The two tables at the bottom of the window display the current static multicast MAC addresses:

- 802.1Q VLAN: Displays the static MAC addresses of network devices in tagged VLANs.
- Port-based VLAN. Displays the static MAC addresses of network devices in port-based VLANs.

The columns in the table are explained in Table 54.

Table 54. Static Multicast Address Table

Column	Description
VLAN ID	Displays the VLAN ID of the tagged VLAN of the MAC address.
VLAN Index	Displays N/A. The VLAN Index numbers of Port-based VLANs does not apply to static addresses.
MAC Address	Displays the static multicast MAC address.

Table 54. Static Multicast Address Table (Continued)

Column	Description
Port Members	Displays the port numbers of the host and router ports where the address is assigned.
Action	Contains Modify and Delete buttons.

2. To add a static multicast MAC address to an 802.1Q Tagged VLAN, do the following:
 - a. Click the **802.1Q VLAN** radio button. This is the default setting.
 - b. Enter the VID of the tagged VLAN containing the host and router ports of the multicast address. You can enter only one VID. The VLAN has to already exist on the switch.
 - c. Go to step 4.
3. To add a static multicast MAC address to a Port-based VLAN, do the following:
 - a. Click the **Port-Based VLAN** radio button.
 - b. Go to step 4.
4. Enter a multicast MAC address in the **Group MAC Address** field. The range is 01:00:5E:00:01:00 to 01:00:5E:7F:FF:FF.
5. Designate the ports of the multicast group by clicking the corresponding radio boxes in **Group Member Settings**. The ports should include host ports and the multicast router port.

You can only select ports whose VLAN mode matches the type of VLAN of the static multicast MAC address. Radio buttons of ports whose VLAN mode does not match the VLAN type will be grayed out. For example, when adding static multicast MAC addresses to 802.1Q Tagged VLANs, you can only choose ports set to the 802.1Q Tagged VLAN mode, the default mode. Or, when adding static multicast MAC addresses to Port-based VLANs, you can only select ports set to the Port-based VLAN mode. Refer to “Assigning the Port-Based VLAN Mode to Ports” on page 193.

6. Click the **Apply** button.
7. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying Static Multicast MAC Addresses

This procedure explains how to change the port assignments of multicast MAC addresses. Here are the guidelines:

- ❑ You cannot change the VLAN or MAC address of a multicast address with this procedure. Instead, you have to delete the entry from the switch and enter it again with the new VLAN or MAC address. Refer to “Deleting Static Multicast Addresses” on page 286.
- ❑ For you to assign a MAC address to a port in an 802.1Q VLAN, the VLAN has to already exist on the switch and the VLAN mode of the port has to be set to 802.1Q Tagged VLAN. Refer to “Adding or Viewing 802.1Q Tagged VLANs” on page 212 and “Assigning the Port-Based VLAN Mode to Ports” on page 193.
- ❑ For you to assign a MAC address to a port in a port-based VLAN, the VLAN has to already exist on the switch and the VLAN mode of the port has to be set to Port-based VLAN. Refer to “Adding or Viewing Port-Based VLANs” on page 190 and “Assigning the Port-Based VLAN Mode to Ports” on page 193.

Perform the following procedure to modify the port assignments of a multicast MAC address:

1. Select the **Bridge -> Static Multicast** from the main menu. Refer to Figure 130 on page 280. The Static Multicast window is shown in Figure 131 on page 281.
2. Click the **Modify** button on the Action column of the static MAC address you want to change. You can modify only one address at a time. The details of the address are displayed in the Modify Static Multicast Address window. The example in Figure 132 is for a MAC address assigned to the default 802.1Q Tagged VLAN with the VID 1.

Static Multicast Address Table

802.1Q VLAN: (1-4093)

Port-Based VLAN

Group MAC Address : : : : : :

Group Member:

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

802.1Q VLAN (Free Entries: 256, Total Entries: 0)

VLAN ID	MAC Address	Port Members	Action
<< Table is Empty >>			

Page 0 / 0 Page

Port-Based VLAN (Free Entries: 256, Total Entries: 0)

VLAN Index	MAC Address	Port Members	Action
<< Table is Empty >>			

Page 0 / 0 Page

Figure 132. Static Multicast Address Table Window for an 802.1Q Tagged VLAN

The example in Figure 133 is for a MAC address assigned to a Port-based VLAN.

Static Multicast Address Table

802.1Q VLAN: (1-4093)

Port-Based VLAN

Group MAC Address : : : : : :

Group Member:

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

802.1Q VLAN (Free Entries: 256, Total Entries: 0)

VLAN ID	MAC Address	Port Members	Action
<< Table is Empty >>			

Page 0 / 0 Page

Port-Based VLAN (Free Entries: 256, Total Entries: 0)

VLAN Index	MAC Address	Port Members	Action
<< Table is Empty >>			

Page 0 / 0 Page

Figure 133. Static Multicast Address Table Window for a Port-Based VLAN

The 802.1Q VLAN field displays 0 (zero) because it does not apply to MAC addresses assigned to ports in Port-based VLANs.

3. Add or delete ports from the MAC address by clicking the corresponding radio buttons under Group Member.

When adding ports, you can only select ports whose VLAN mode matches the type of VLAN of the static unicast MAC address. Radio

buttons of ports whose VLAN mode does not match the VLAN type will be grayed out. For example, when adding static unicast MAC addresses to 802.1Q Tagged VLANs, you can only select ports set to the 802.1Q Tagged VLAN mode, the default mode. Or, when adding static unicast MAC addresses to Port-based VLANs, you can only choose ports set to the Port-based VLAN mode. Refer to “Assigning the Port-Based VLAN Mode to Ports” on page 193 or “Assigning the 802.1Q Tagged VLAN Mode to Ports” on page 216.

4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Deleting Static Multicast Addresses

Perform the following procedure to delete static multicast MAC addresses from the MAC address table:

1. Select the **Bridge -> Static Multicast** from the main menu. Refer to Figure 130 on page 280. The Static Multicast window is shown in Figure 131 on page 281.
2. Do one of the following:
 - To delete a single entry, click its **Delete** button in the Action column.
 - To delete all 802.1Q address entries, click the **Delete All** button above the 802.1Q VLAN table.
 - To delete all Port-based VLAN address entries, click the **Delete All** button above the Port-Based VLAN table.
3. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 24

IGMP Snooping

This chapter explains IGMP snooping in the following sections:

- ❑ “IGMP Snooping Overview” on page 288
- ❑ “Configuring IGMP Snooping” on page 290
- ❑ “Adding or Deleting Static Multicast Router Ports” on page 293
- ❑ “Viewing Multicast Addresses and Host Ports” on page 297

IGMP Snooping Overview

The switch uses IGMP snooping to improve network performance of multicast groups. The switch directs multicast traffic to only those ports in VLANs that have multicast hosts. The switch monitors the transmissions of queries from multicast router, and reports and leave requests from hosts to build its own multicast membership lists. It then uses the lists to forward multicast packets only to ports with host nodes. Without IGMP snooping, the switch has to flood multicast traffic to all VLAN ports, including ports without host nodes.

IPv4 multicast routers use IGMP to create lists of nodes that are members of multicast groups. (A group of end nodes that receive multicast packets from a multicast application is defined as a multicast group.) The router creates a multicast membership list by periodically sending out queries to the local area networks connected to its ports.

Nodes that want to become members of multicast groups respond to queries by sending *reports*. Nodes that join multicast groups are referred to as *host nodes*. After joining multicast groups, host nodes have to periodically issue new reports to remain group members.

After the router has received a report from a host node, it notes the multicast group that the host node wants to join and the port on the router where the node is located. Any multicast packets belonging to that multicast group are then forwarded by the router out the port. If a particular port on the router has no nodes that want to be members of multicast groups, the router does not send multicast packets from the port. This improves network performance by restricting multicast packets only to router ports where host nodes are located.

There are three versions of IGMP— versions 1, 2, and 3. A difference between the versions is how a host node signals that it no longer wants to be a member of a multicast group. In version 1, it stops sending reports. If a router does not receive a report from a host node after a pre-defined length of time, referred to as a *time-out value*, it assumes that the host node no longer wants to receive multicast frames and removes it from the membership list of the multicast group. If there are no additional member hosts on the port, the switch stops forwarding multicast packets from it.

In version 2, a host node exits from a multicast group by sending a *leave request*. After receiving a leave request from a host node, the router removes the node from the appropriate membership list. It also stops sending multicast packets from the port if it determines there are no additional host nodes on the port.

Version 3 adds the ability of host nodes to join or leave specific sources in multicast groups. It also enables the switch to build membership reports of the host nodes, and to send the reports in lieu of individual report and

leave requests to multicast routers, thereby reducing traffic overhead.

Note

The GS950 PS V2 Switch do not support IGMPv3 snooping.

The IGMP snooping feature on the GS950 PS V2 Switch supports IGMP versions 1 and 2. The switch monitors the flow of queries from routers and checks report and leave messages from host nodes to build its own multicast membership lists. It uses the lists to forward multicast packets only to ports where there are host nodes that are members of multicast groups. This improves switch performance and network security by restricting the flow of multicast packets only to those ports connected to host nodes. The switch can support up to 256 multicast addresses.

Without IGMP snooping, a switch has to flood multicast packets on all VLAN ports, except the port on which it received the packet. Such flooding of packets can reduce network performance.

The switch also has IGMP snooping querier. You can use this feature if your network has multicast traffic, but no multicast router. When the feature is enabled, the switch sends out queries, searching for host nodes in VLANs on its ports wanting to join multicast group.

The GS950 PS V2 Switch maintains a list of multicast groups through an adjustable time-out value, which controls how frequently it expects to see reports from end nodes that want to remain members of multicast groups, and by processing leave requests.

Note

The default setting for IGMP snooping is disabled.

Configuring IGMP Snooping

Perform the following procedure to enable or disable IGMP snooping and configure its settings:

1. Select **Bridge -> IGMP Snooping -> IGMP Snooping Settings** from the main menu. Refer to Figure 134.

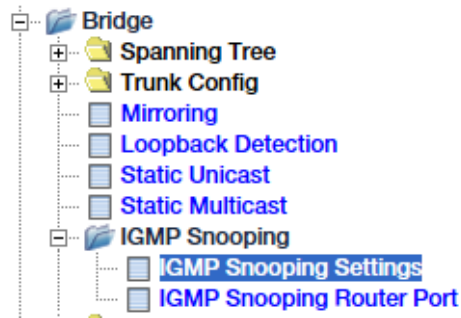


Figure 134. IGMP Snooping Settings Menu Selection

The IGMP Snooping Settings window is shown in Figure 135.

IGMP Snooping Settings

Status: ▾

Age-Out Timer: Sec. (130-153025)

Querier Status: ▾

Query Interval: Sec. (60-600)

Max Response Time: Sec. (10-25)

Robustness Variable: Times. (2-255)

Last Member Query Interval: Sec. (1-25)

Router Timeout: Sec. (120-1200)

Note: The Host Timeout will be computed automatically in Querier Enabled by (Robustness Variable * Query Interval + Max Response Time).

Show result per 802.1Q VLAN ID
Show result per port-based VLAN Index

802.1Q VLAN (Free entries: 256 Total entries: 0)

VLAN ID	Multicast Group Address	Member Ports
<< Table is empty >>		

Page 0 / 0

Page

Figure 135. IGMP Snooping Settings Window

Note

The table at the bottom of the window is explained in “Viewing Multicast Addresses and Host Ports” on page 297.

2. Configure the IGMP snooping parameters in Table 55.

Table 55. IGMP Snooping Settings Window

Setting	Description
Status	Select one of the following options: <ul style="list-style-type: none"> - Enabled: Enables IGMP snooping on the switch. You have to enable IGMP snooping to configure the settings. - Disabled: Disables IGMP snooping. This is the default setting.
Age-out Timer	Enter the maximum time in seconds the switch should wait for reports from hosts wanting to remain members of multicast groups. Hosts are removed from multicast groups if they do not send reports before the age-out timer expires. The range is 130 to 153025 seconds. The default is 260 seconds.
Querier Status	Select the status of the IGMP snooping querier on the switch. The querier can be used in networks where there is multicast traffic but no multicast router. The switch uses the querier to send its own multicast queries to search for host nodes. The settings are: <ul style="list-style-type: none"> - Enabled: The multicast querier is enabled. The switch transmits its own queries for host nodes. - Disabled: The querier is disabled. This is the default setting.

Table 55. IGMP Snooping Settings Window (Continued)

Setting	Description
Query Interval	<p>Enter the interval in seconds for the transmission of multicast queries for host nodes by the switch. Here are the guidelines:</p> <ul style="list-style-type: none"> - Querier Status has to be enabled. - The query interval has to be less than the age-out timer. Otherwise, hosts might not receive multicast traffic. - The range is 60 to 600 seconds. The default is 125 seconds.
Max Response Time	<p>Enter the maximum response time the switch adds to the query packets that it transmits to hosts. This is the maximum amount of time hosts have to respond to the queries. Here are the guidelines.</p> <ul style="list-style-type: none"> - Querier Status has to be enabled. - The range is 10 to 25 seconds. The default is 10 seconds.
Robustness Variable	<p>Enter the number of times hosts can fail to respond to consecutive snooping queries before the switch deletes them from multicast groups. For example, when set to 4, the switch deletes hosts from multicast groups if they fail to respond to four consecutive queries. The range is 2 to 255 times. The default is 2 times.</p>
Last Member Query Interval	<p>Enter the maximum time in seconds the switch should wait before deleting multicast groups from VLAN ports when no hosts respond to IGMP query messages. The range is 1 to 25 seconds. The default is 1 second.</p>
Router Timeout	<p>Enter the maximum time the switch should wait for queries from multicast routers before removing them from its multicast tables. The range is 120 to 1200 seconds. The default is 250 seconds.</p>

3. Click the **Apply** button.
4. Select **Save Settings to Flash** to save your changes.

Adding or Deleting Static Multicast Router Ports

To support IGMP snooping, the switch needs to know which ports are connected to multicast routers. These are the ports on which the switch receives queries from multicast routers, transmits reports from host nodes to routers, and receives multicast traffic streams. The switch can learn the ports automatically by watching for multicast queries on its ports from the routers, or you can designate the ports manually. The latter ports are referred to as static multicast router ports. You might add static router ports for multicast routers that the switch might not learn automatically or for routers the switch should never unlearn.

Static multicast router ports have to belong to 802.1Q Tagged or Port-based VLANs. Thus, the first step to designating ports as static multicast router ports is to add the appropriate VLANs to the switch and add the ports to the VLANs. Refer to Chapter 16, “802.1Q Tagged Virtual LANs” on page 205 and Chapter 15, “Port-Based Virtual LANs” on page 183.

Perform the following procedure to add or delete static multicast router ports:

1. Select **Bridge > IGMP Snooping > IGMP Snooping Router Port** from the main menu. Refer to Figure 136.

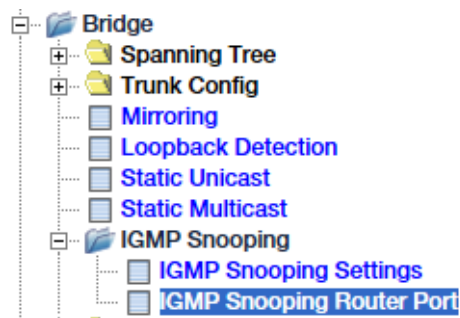


Figure 136. IGMP Snooping Router Port Menu Selection

The switch displays the IGMP Snooping Router Port window in Figure 137 on page 294. The window has two tables:

- ❑ 802.1Q VLAN: Lists static and dynamic router ports in 802.1Q Tagged VLANs. The VLANs are listed individually.
- ❑ Port-based VLAN: Lists static and dynamic router ports in Port-based VLANs. The table has only one entry for the Port-based VLANs. All the multicast router ports are grouped in the one entry.

IGMP Snooping Router Port

802.1Q VLAN ID

VLAN ID	Static Router Port	Dynamic Router Port	Action
1	N/A	N/A	Modify
21	N/A	7	Modify
22	11	N/A	Modify
30	N/A	N/A	Modify

Page 1 / 1 Page

Port-Base VLAN

VLAN Index	Static Router Port	Dynamic Router Port	Action
N/A	1,16-17	N/A	Modify

Page 1 / 1 Page

Figure 137. IGMP Snooping Router Port Window

The columns in the tables are described in Table 56.

Table 56. IGMP Snooping Router Port Window

Column	Description
VLAN ID	Displays the VID of a tagged VLAN with a static or dynamic router port. The VLANs are listed individually.
VLAN Index	Displays N/A. The Port-based VLANs are not listed individually. Rather, the static and dynamic router ports for all Port-based VLANs are combined into one entry.
Static Router Port	Displays ports that were manually designated as static multicast router ports in the VLAN.
Dynamic Router Port	Displays ports that the switch automatically learned as multicast router ports because it received queries from multicast routers on the ports.

Table 56. IGMP Snooping Router Port Window (Continued)

Column	Description
Action	Displays the Modify button for adding or removing static multicast router ports to the VLAN.

- To add or remove static multicast router ports from an 802.1Q Tagged or a Port-based VLAN, click the corresponding **Modify** button. The Modify IGS Static Router Port window is displayed. The example in Figure 138 is for an 802.1Q Tagged VLAN.

Modify IGS Static Router Port

802.1Q VLAN ID: 57

Static Router Port

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Figure 138. Modify IGS Static Router Port Window for an 802.1Q Tagged VLAN

The example in Figure 139 is for the Port-based VLANs. As mentioned previously, the static router ports on Port-based VLANs are combined into one entry.

Modify IGS Static Router Port

Port-Based VLAN Index: N/A

Static Router Port

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
All	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	15	16	17	18	19	20	21	22	23	24	25	26	27	28
	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Figure 139. Modify IGS Static Router Port Window for the Port-Based VLANs

3. Click the check boxes of ports to be designated or removed as static multicast router ports. A check mark designates a port as a static multicast router port. VLANs can have more than one static multicast router port.

The window includes these options:

- To select all ports, click the **All** button in the left margin.
 - To restore the original settings, click the **Restore** button in the right margin.
4. Click the **Apply** button.
 5. Select **Save Settings to Flash** from the main menu to save your changes.

Viewing Multicast Addresses and Host Ports

To view multicast addresses and ports with active host nodes on the switch, select **Bridge -> IGMP Snooping -> IGMP Snooping Settings** from the main menu. Refer to Figure 135 on page 290. The top section of window is explained in “Configuring IGMP Snooping” on page 290 for details. The table at the bottom of the window has two buttons:

- ❑ **Show result per 802.1Q VLAN ID** - Click this button to view the multicast addresses and ports with active host nodes of 802.1Q Tagged VLANs. This is the default. The VLANs are listed individually. Refer to Figure 140.
- ❑ **Show result per port-based VLAN Index** - Click this button to view the multicast addresses and ports with active host nodes of port-based VLANs. The ports are combined into one entry in the table. The VLANs are not listed individually.

The tables do not include ports with inactive hosts.

Show result per 802.1Q VLAN ID
Show result per port-based VLAN Index

802.1Q VLAN (Free entries: 254 Total entries: 2)

VLAN ID	Multicast Group Address	Member Ports
21	01:00:BE:00:00:01	3,5
22	01:00:BE:00:00:44	18-20

Page 0 / 0
First Page
Previous Page
Next Page
Last Page
Page
Go

Figure 140. Multicast MAC Address Tables in the IGMP Snooping Window

The columns in the table are described in Table 57.

Table 57. Multicast MAC Address Tables in the IGMP Snooping Settings Window

Columns	Description
VLAN ID	Displays the VID of a tagged VLAN.
VLAN Index	Displays N/A. The table has only one entry for Port-based VLANs.
Multicast Group Address	Displays active dynamic multicast addresses in the VLAN. Also displays active and inactive static multicast addresses.

Table 57. Multicast MAC Address Tables in the IGMP Snooping Settings Window (Continued)

Columns	Description
Member Ports	Displays active multicast router and host ports.

Chapter 25

Quality of Service and Class of Service

This chapter describes the Quality of Service (QoS) and Class of Service (CoS) features in the following sections:

- ❑ “QOS and COS Overview” on page 300
- ❑ “Mapping CoS Priorities to Egress Queues” on page 304
- ❑ “Mapping CoS Priorities to Ports” on page 306
- ❑ “Mapping DSCP Classes to Egress Queues” on page 307
- ❑ “Setting the Queue Scheduling Algorithm” on page 309
- ❑ “Mapping IPv6 Traffic Classes to Port Egress Queues” on page 311

QOS and COS Overview

Class of Service (CoS) is used to assign network packets to port egress queues in the switch based on their priority levels. The switch uses CoS to prioritize transmission of network packets, giving higher priority to designated packets, such as delay sensitive traffic, over other packets.

When a port on an Ethernet switch becomes oversubscribed, its egress queues contain more packets than the port can handle in a timely manner. When this happens, a port may be forced to delay the transmission of some packets, resulting in the delay of packets reaching their destinations. A port may be forced to delay transmission of packets while it handles other traffic, and, in some situations, some packets destined to be forwarded to an oversubscribed port from other switch ports may be discarded.

Minor delays are usually of no consequence to a network or its performance. But there are applications, referred to as delay or time-sensitive applications, that can be impacted by packet delays. Voice transmission and video conferences are two examples. The audio or video quality may suffer if packets carrying their data are delayed from reaching their destinations.

The topics of CoS are:

- ❑ “Packet Priority” next
- ❑ “Egress Queue vs Packet Priority Mapping” on page 301
- ❑ “Prioritizing Untagged Packets” on page 302
- ❑ “Scheduling Algorithm” on page 302

Packet Priority

CoS applies primarily to tagged packets. Tagged packets contain information that specifies the VLANs to which they belong. Tagged packets can also contain priority levels that are used by network switches and other networking devices to know how important (delay sensitive) packets are compared to other packets. Packets with high priorities are handled before packets with low priorities.

CoS, as defined in the IEEE 802.1p standard, has eight priority levels. The priority levels are 0 to 7, with 0 the lowest and 7 the highest priority.

When tagged packets are received on ports, the switch examines them for their priorities. It uses the priorities to determine which egress queues the packet should be directed to on the egress ports.

Egress Queue vs Packet Priority Mapping

Each port has four egress queues, labeled Low, Medium, High, Highest. Low is the lowest priority queue and Highest is the highest. Packets in a high-priority egress queue are usually transmitted sooner than packets in a low-priority queue. Table 58 lists the default mappings of the eight CoS priority levels and the four egress queues of the switch ports.

Table 58. Default Mappings Priority Levels to Priority Queues

IEEE 802.1p Priority Level	Port Priority Queue
0	Low
1	Low
2	Low
3	Low
4	Low
5	Low
6	Low
7	Low

You can change these mappings. For example, you might decide that packets with priority 4 or 5 should be handled in the High egress queues while packets with a priority of 6 or 7 should be handled in the Highest egress queues. The result is shown in Table 59.

Table 59. Example of Customized Mappings Priority Levels to Priority Queues

IEEE 802.1p Priority Level	Port Priority Queue
0	Low
1	Low
2	Low
3	Low
4	High
5	High
6	Highest
7	Highest

The procedure for changing the default mappings is found in “Mapping CoS Priorities to Ports” on page 306. These mappings are applied at the switch level. All ports use the same priority-to-egress queue mappings.

You can also map an IPv6 packet header’s 8 bit priority field, used by the switch to differentiate between classes or priorities of IPv6 ports, to one of the switch’s priority queues. The procedure is found in “Mapping IPv6 Traffic Classes to Port Egress Queues” on page 311.

The switch does not change the priority levels in tagged packets. It transmits them from the egress ports with the same priority levels they had when they were received. This is true even if you change the default priority-to-egress queue mappings.

Prioritizing Untagged Packets

Unlike tagged packets, untagged packets do not contain priority values, but they can still be prioritized by the switch. Priority values are assigned to untagged packets by their ingress ports. Each port has a priority value for ingress untagged packets. As untagged packets arrive on a port, the switch assigns them the port’s priority value. The priority determines the queue untagged packets are placed in on the egress ports. The default priority value for ingress untagged packets is “0”, which places them in the Low priority queue. You can change the setting as described in “Mapping CoS Priorities to Ports” on page 306.

Scheduling Algorithm

The switch has a process for controlling the order in which it transmits packets from the four egress queues of ports. The process is called the scheduling algorithm. The switch uses scheduling to determine the order in which ports handle the packets in their egress queues. For example, if all the queues of a port contain packets, the switch uses the process to determine whether to transmit all the packets from the highest (the highest priority queue) before moving on to the other queues, or transmit a few packets from each queue in a sequential fashion.

The switch has two types of scheduling:

- Strict priority
- Weighted round robin priority

To specify the scheduling, refer to “Mapping CoS Priorities to Ports” on page 306.

Note

Scheduling is set at the switch level and applies to all ports.

Strict Priority Scheduling

With strict priority scheduling, ports transmit all packets out of higher priority queues before transmitting any from the lower priority queues. For instance, as long as there are packets in the Highest queues, they do not handle any packets in the High queues. The value of this type of scheduling is that high-priority packets are always handled before low-priority packets which is required for voice or video data.

The problem with this method is that some low-priority packets might never be transmitted from the switch. This can happen if higher priority queues always contain packets because of high traffic volume.

Weighted Round Robin Priority Scheduling

The weighted round robin (WRR) scheduling method functions as its name implies. The port transmits a set number of packets from each queue, in a round robin fashion, so that each has a chance to transmit traffic. Normally, the higher a queue's priority, the more packets are transmitted from it as compared to lower priority queues. This method guarantees that every queue receives some attention from the port for transmitting packets.

Table 60 shows the WRR settings for the number of packets transmitted from each queue. You cannot change these settings.

Table 60. Weighted Round Robin Priority

Port Egress Queue	Maximum Number of Packets
Highest	8
High	4
Medium	2
Low	1

Mapping CoS Priorities to Egress Queues

Note

CoS is not compatible with Jumbo frames. Before enabling CoS, disable Jumbo frames on all ports. Refer to “Configuring Basic Port Settings” on page 51.

Perform the following procedure to map CoS priorities to port egress queues:

1. Select **Bridge > QoS > CoS** from the main menu in Figure 141.

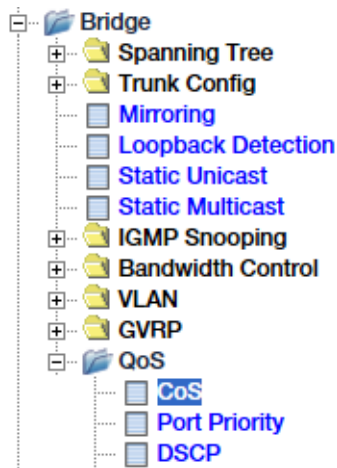


Figure 141. CoS Menu Selection

The CoS window is shown Figure 142 on page 305.

CoS

CoS Status:

Disabled ▾

Traffic Class	Queue			
0	<input checked="" type="radio"/> Low:	<input type="radio"/> Medium :	<input type="radio"/> High :	<input type="radio"/> Highest :
1	<input checked="" type="radio"/> Low:	<input type="radio"/> Medium :	<input type="radio"/> High :	<input type="radio"/> Highest :
2	<input checked="" type="radio"/> Low:	<input type="radio"/> Medium :	<input type="radio"/> High :	<input type="radio"/> Highest :
3	<input checked="" type="radio"/> Low:	<input type="radio"/> Medium :	<input type="radio"/> High :	<input type="radio"/> Highest :
4	<input checked="" type="radio"/> Low:	<input type="radio"/> Medium :	<input type="radio"/> High :	<input type="radio"/> Highest :
5	<input checked="" type="radio"/> Low:	<input type="radio"/> Medium :	<input type="radio"/> High :	<input type="radio"/> Highest :
6	<input checked="" type="radio"/> Low:	<input type="radio"/> Medium :	<input type="radio"/> High :	<input type="radio"/> Highest :
7	<input checked="" type="radio"/> Low:	<input type="radio"/> Medium :	<input type="radio"/> High :	<input type="radio"/> Highest :

Apply

Note: Disable will reset the settings to factory default and turn off the function.

Figure 142. CoS Window

- To enable or disable the mappings, select **Enabled** or **Disabled** from the CoS Status menu. The mappings have to be enabled for you to change the settings. The default is disabled.
- To change a mapping, click the **Low**, **Medium**, **High**, or **Highest** radio button for the appropriate traffic class. The default for all traffic classes is the Low queue.
- Click the **Apply** button.
- Select **Save Settings to Flash** from the main menu to save your changes.

Mapping CoS Priorities to Ports

This procedure explains how to set CoS port priority values on the individual ports on the switch. The switch assigns the priority values to ingress untagged packets. A port can have only one port priority value. The default setting is 0, which places untagged packets in the Low queue on egress ports. Refer to “Prioritizing Untagged Packets” on page 302.

Perform the following procedure to change the CoS port priority settings.

1. Select **Bridge > QoS > Port Priority** from the main menu. Refer to Figure 143.

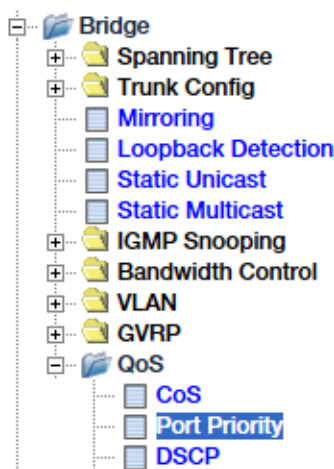


Figure 143. Port Priority Menu Selection

The Port Priority window is shown in Figure 144.

Port Priority		
Port	User Priority	Action
All	0 ▾	Apply
1	0 ▾	Apply
2	0 ▾	Apply
3	0 ▾	Apply
4	0 ▾	Apply

Figure 144. Port Priority Window

2. To change a port’s CoS priority, select the new value from the User Priority menu. The range is 0 to 7. The default is 0.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Mapping DSCP Classes to Egress Queues

This procedure explains how to map DSCP values (0-63) to port egress queues (Low to Highest). The default queue for all DSCP values is Low.

Perform the following procedure to map DSCP values to egress queues:

1. Select **Bridge > QoS > DSCP >** from the main menu. Refer to Figure 145.

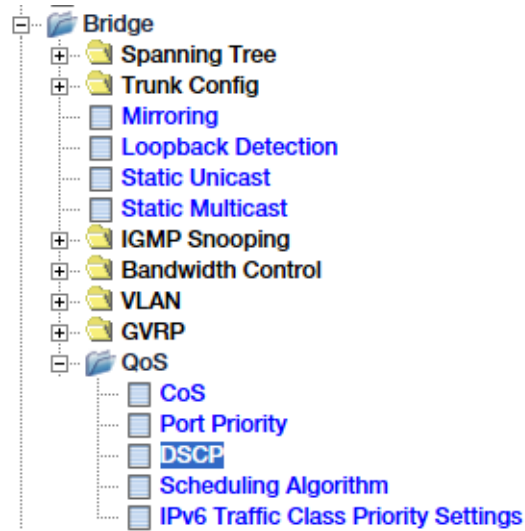


Figure 145. DSCP Menu Selection

The Port Priority window is shown in Figure 146 on page 308.

DSCP Class Mapping

DSCP Mapping Status: Disabled ▾

Apply

DSCP In	Queue	DSCP In	Queue	DSCP In	Queue	DSCP In	Queue
0-15	Ignore ▾	16-31	Ignore ▾	32-47	Ignore ▾	48-63	Ignore ▾
0	Low ▾	16	Low ▾	32	Low ▾	48	Low ▾
1	Low ▾	17	Low ▾	33	Low ▾	49	Low ▾
2	Low ▾	18	Low ▾	34	Low ▾	50	Low ▾
3	Low ▾	19	Low ▾	35	Low ▾	51	Low ▾
4	Low ▾	20	Low ▾	36	Low ▾	52	Low ▾
5	Low ▾	21	Low ▾	37	Low ▾	53	Low ▾
6	Low ▾	22	Low ▾	38	Low ▾	54	Low ▾
7	Low ▾	23	Low ▾	39	Low ▾	55	Low ▾
8	Low ▾	24	Low ▾	40	Low ▾	56	Low ▾
9	Low ▾	25	Low ▾	41	Low ▾	57	Low ▾
10	Low ▾	26	Low ▾	42	Low ▾	58	Low ▾
11	Low ▾	27	Low ▾	43	Low ▾	59	Low ▾
12	Low ▾	28	Low ▾	44	Low ▾	60	Low ▾
13	Low ▾	29	Low ▾	45	Low ▾	61	Low ▾
14	Low ▾	30	Low ▾	46	Low ▾	62	Low ▾
15	Low ▾	31	Low ▾	47	Low ▾	63	Low ▾

Apply
Reset to Default

Figure 146. DSCP Class Mapping Window

2. Select **Enabled** or **Disabled** from the DSCP Mapping Status menu to enable or disable the mappings. The default is disabled. You have to enable the DSCP mappings to set the values.
3. Click the **Apply** button under the DSCP Mapping Status menu.
4. Adjust the queue settings for the individual DSCP values, as needed. The queues are **Low**, **Medium**, **High**, and **Highest**. To set all the queues in a Queue column to the same value, use the Queue menu in the top 0-15 row.
5. Click the **Apply** button at the bottom of the window.

To return the DSCP class mapping to the default values, click the **Reset to Default** button.

6. Select **Save Settings to Flash** from the main menu to save your changes.

Setting the Queue Scheduling Algorithm

This procedure explains how to set the scheduling algorithm. The switch uses the scheduling algorithm to control the order in which it transmits packets from the four egress queues of ports. Refer to “Scheduling Algorithm” on page 302.

Perform the following procedure to change the scheduling algorithm for the port egress queues:

1. Select **Bridge > QoS > Scheduling Algorithm** from the main menu. Refer to Figure 147.

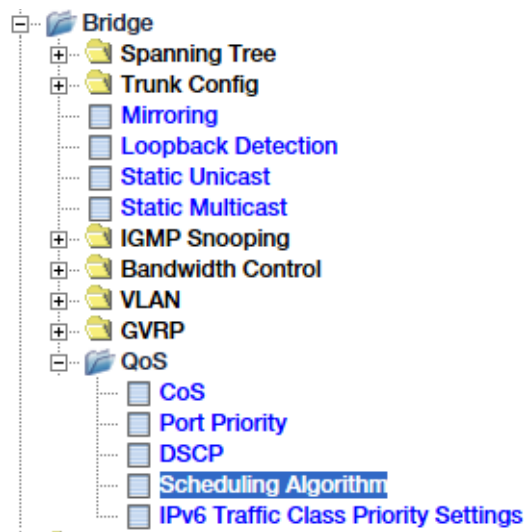


Figure 147. Scheduling Algorithm Menu Selection

The Scheduling Algorithm window is shown in Figure 148.



Figure 148. Scheduling Algorithm Window

2. Select one of the following algorithms from the **Scheduling Algorithm** menu:
 - Strict Priority** - The port transmits all packets out of higher priority queues before transmitting any from the lower priority queues. Refer to “Strict Priority Scheduling” on page 303.
 - Weighted Round Robin** - The port transmits a set number of packets from each queue, in a round robin, so that each has a

chance to transmit traffic. Refer to “Weighted Round Robin Priority Scheduling” on page 303.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Mapping IPv6 Traffic Classes to Port Egress Queues

Note

You cannot map IPv6 traffic class mapping to egress queues when Jumbo frames are enabled on ports. To disable Jumbo frames, refer to “Configuring Basic Port Settings” on page 51.

Perform the following procedure to map IPv6 traffic classes to port egress queues:

1. Select **Bridge > QoS > IPv6 Traffic Class Priority Settings** from the main menu. Refer to Figure 149.

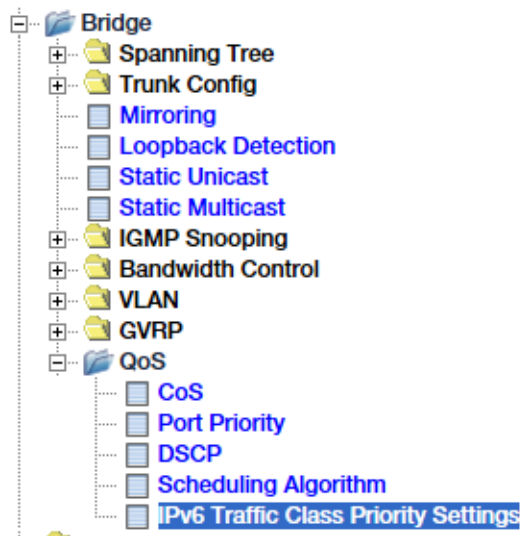


Figure 149. IPv6 Traffic Class Priority Settings Menu Selection

The IPv6 Traffic Class Priority Settings window is shown in Figure 150 on page 312.

IPv6 Traffic Class Priority Settings

State:

IPv6 Traffic Class:

Priority:

Free Entries: 188
Total Entries: 0

IPv6 Traffic Class	Priority	Action
<< Table is empty >>		

Page 0 / 0 Page

Figure 150. IPv6 Traffic Class Priority Settings Window

2. To enable or disable the mappings, do the following:
 - a. Select **Enabled** or **Disabled** from the State menu. The mappings have to be enabled for you to modify the settings. The default is disabled.
 - b. Click the **Apply** button.
3. To add new mappings, do the following:
 - a. Enter a value for the IPv6 packet header's 8 bit priority in the IPv6 Traffic Class field. The range is 0-255.
 - b. Select a **Low**, **Medium**, **High** or **Highest** queue from the Priority menu.
 - c. Click the **Apply** button. The new mapping is added to the table.
 - d. Repeat this step to add more mappings.
4. To delete IPv6 traffic class priority mappings, click the **Delete** button in the Action column to delete individual mappings or click the **Delete All** button to delete all the entries.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 26

Access Control Policies

This chapter describes access control lists in the following sections.

- ❑ “Access Control Policies Overview” on page 314
- ❑ “Adding Layer 2 IPv4 Policies” on page 317
- ❑ “Adding IPv6 Policies” on page 325
- ❑ “Enabling or Disabling Policies” on page 329
- ❑ “Modifying or Deleting Policies” on page 330
- ❑ “Displaying Policies” on page 331
- ❑ “Displaying Policies by Port Numbers” on page 332
- ❑ “Managing Rate Control Settings” on page 333

Access Control Policies Overview

Access control policies function as packet filters on switch ports. You can use them to control the flow of ingress packets on ports by designating the types of packets that the switch can accept or discard on ports. You might use the policies to increase port security or to add physical links that are dedicated to carrying specific types of traffic. For instance, you might add policies to ports so that they accept only those ingress packets that have a specific source or destination IPv4 or IPv6 address, and to reject all other packets.

Note

Before adding access control policies, be sure to configure the QoS parameters. QoS entries may have a direct effect on a policy's behavior. For more information, refer to Chapter 25, "Quality of Service and Class of Service" on page 299.

Policy Filters

You can add IPv4 or IPv6 port policies. The two types of policies support different sets of filters. Here are the available filters for IPv4 ingress packets:

- Source or destination MAC address
- Source or destination IPv4 address
- VLAN ID
- Ether type
- Ethernet protocol
- DSCP
- Source Transport Layer 4 Port

Here are the filters for IPv6 ingress packets:

- Source or destination IPv6 address
- VLAN ID
- 802.1p priority
- IPv6 traffic class
- Source Transport Layer 4 Port

Actions

The policy action defines a port's response to packets that match the filtering criterion of a policy. There are two actions:

- Permit - Instructs ports to accept ingress packets that match the policy.
- Deny - Instructs ports to discard ingress packets that match the policy.

Ports, by default, forward all ingress packets. Thus, a policy with a permit action is only required when you want a port to forward a subset of packets of a larger traffic flow that are otherwise to be blocked.

How Ingress Packets are Compared Against Policies

Ports that do not have any policies forward all ingress packets. Ports with one or more deny ACLs discard ingress packets that match the ACLs and forward all other traffic. A port that has one deny ACL that specifies a particular source IP address, for example, discards all ingress packets with that source address and forwards all other traffic. In situations where a port has more than one deny ACL, packets are discarded at the first match.

Since ports forward all ingress packets unless they have deny ACLs, permit ACLs are only necessary in situations where you want a port to forward packets that are a subset of a larger traffic flow that is to be blocked. An example is a port that is to forward only packets having a specified destination IP address. A permit ACL would specify the packets with the permitted destination IP address and a deny ACL would specify all traffic.

ACLs are assigned policy sequence numbers that control the order in which the switch uses them to filter packets when there are multiple ACLs on ports. The lower the number the higher the priority. The range is 1 to 65535. It is important to assign permit policies lower sequence numbers than deny policies so that packets are compared against them first when ports have both permit and deny ACLs. For example, you might reserve sequence numbers 1 to 500 for permit ACLs and numbers 501 to 1000 for deny ACLs.

Guidelines

Here are the policy guidelines:

- The policy action can be permit or deny. Permit actions allow ports to accept ingress packets of the designated traffic flows while deny actions cause ports to discard packets.
- A port can have more than one policy.
- A policy can be assigned to more than one port.
- You cannot assign a policy more than once to the same port.
- Policies filter ingress packets on ports, but not egress packets. Therefore, policies have to be applied to the ingress ports of traffic flows.
- Policies for static port or LACP trunks have to be assigned to all the member ports of the trunks.
- Policies are assigned sequence numbers that determine the order in which they are performed on ports that have more than one policy. The range is 1 to 65535.
- On ports that have more than one policy, packets are compared against the policies in the order of the policy sequence numbers, from

lowest to highest. Packets are forwarded or discarded at the first match.

- ❑ Ports can have policies with different filtering criteria. A port, for example, could have policies that filter on a source IP address and a VLAN ID.
- ❑ Because ports, by default, forward all ingress packets, permit policies are only required when ports are to forward packets that are subsets of larger packet flows that are blocked by deny policies.

Adding Layer 2 IPv4 Policies

Perform the following procedure to add Layer 2 IPv4 policies to switch ports:

1. Select **Access Control Config > Policy Settings** from the main menu. Refer to Figure 151.



Figure 151. Policy Settings Menu Selection

The Policy Settings window is shown in Figure 152.

Policy Settings

Policy type:

Free Policies : 188
Total Entries : 0

Index	Classifier	Sequence	Deny/Permit	Replaced-CoS	Replaced-DSCP	Rate Control	Port List	Status	Action
<< Table is empty >>									

Page 0 / 0 Page

Figure 152. Policy Settings Window

2. Click the **Add L2+IPv4** button. The IPv4 Policy Settings page is shown in Figure 153 on page 318.

Policy Settings

Policy type:

Policy Index: (1-65535)

Source MAC Address: : : : : : **Mask Length:** (1-48)

Destination MAC Address: : : : : : **Mask Length:** (1-48)

VLAN ID: (1-4093) **802.1p Priority:** (0-7)

Ether Type: 0x (0000-FFFF, ex: 0806; 0800)

Protocol: (1-255)

IPv4 Source IP Address: . . . **Mask Length:** (1-32)

IPv4 Destination IP Address: . . . **Mask Length:** (1-32)

DSCP: (0-63)

Source Layer 4 Port: (1-65535) **Destination Layer 4 Port:** (1-65535)

Policy Sequence: (1-65535)

Policy Action: ▼

Replaced-CoS: (0-7) **Rate Control Index:** (1-65535)

Replaced-DSCP: (0-63)

Port List: (e.g. 1,3,5-8)

Free Policies : 188

Total Entries : 0

Index	Classifier	Sequence	Deny/Permit	Replaced-CoS	Replaced-DSCP	Rate Control	Port List	Status	Action
<< Table is empty >>									

Page 0 / 0

Figure 153. IPv4 Policy Settings Window

3. Configure the IPv4 policy fields in Table 61.

Table 61. IPv4 Policy Settings Window

Field	Description
Policy Index	Enter a unique index number for the policy. Index numbers determine the order in which the policies are listed in the table at the bottom of the window. The range is 1 to 65535. Policies will be easier to manage if you make their index numbers and policy sequence numbers the same. This setting is required.
Source MAC Address	To filter ingress packets based on a source MAC address, enter the address in this field. You can enter only one MAC address.

Table 61. IPv4 Policy Settings Window (Continued)

Field	Description
Mask Length	Enter the mask for the source MAC address. The mask is a decimal number that represents the number of bits in the address, from left to right, on which to filter. The range is 1 to 48. For example, the mask for the MAC address 11:22:33:44:55:00 is 40. To filter on the whole address of a specific network device, enter 48 as the mask.
Destination MAC Address	To filter ingress packets based on a destination MAC address, enter the address in this field. You can enter only one MAC address.
Mask Length	Enter the mask for the destination MAC address. The mask is a decimal number that represents the number of bits in the address, from left to right, on which you want to filter. The range is 1 to 48. For example, the mask for the MAC address 11:22:33:44:55:00 would be 40. To filter on the whole address for a specific network device, enter 48 as the mask.
VLAN ID	To filter ingress tagged packets based on the VID of a 802.1Q tagged VLAN, enter the VID in this field. You can enter only one VID. The range is 1 to 4093. To view the VIDs of the 802,1 Tagged VLANs on the switch, refer to “Adding or Viewing 802.1Q Tagged VLANs” on page 212 or “Viewing Port-based and 802.1Q Tagged VLANs” on page 223. This filter does not apply to port-based or private VLANs.
802.1p Priority	To filter ingress packets based on the 802.1p priority level in packet headers, enter the level in this field. The range is 0 to 7.
Ether Type	To filter ingress packets based on an Ethernet frame protocol number, enter the protocol number here. You can enter only one number. The range is 0000 to FFFF.
Protocol	To filter ingress packets based on a IP protocol number, enter the number in this field. You can enter only one number. Refer to Table 62 on page 322.

Table 61. IPv4 Policy Settings Window (Continued)

Field	Description
IPv4 Source IP Address	To filter ingress packets based on a source IPv4 address, enter the address in this field. You can enter only one address.
Mask Length	Enter the mask for the source IPv4 address. The mask is a decimal number that represents the number of bits in the address, from left to right, on which you want to filter. The range is 1 to 32. For example, the mask is 24 for the IP address 111.222.333.0, To filter on a whole address for a specific network device, enter 32 as the mask.
IPv4 Destination IP Address	To filter ingress packets based on a destination IPv4 address, enter the address in this field. You can enter only one address.
Mask Length	Enter the mask for the destination IPv4 address. The mask is a decimal number that represents the number of bits in the address, from left to right, on which you want to filter. The range is 1 to 32. For example, the mask for the IP address 111.222.333.0 would be 24, To filter on a whole address for a specific network device, enter 32 as the mask.
DSCP	To filter ingress packets based on the Differentiated Services Code Point (DSCP) in the headers, enter the code in this field. You can enter only one code. The range is 0 to 63.
Source Layer 4 Port	To filter ingress source Transport Layer 4 packets, enter a port number in this field. The range is 1 to 65535.
Destination Layer 4 Port	To filter ingress destination Transport Layer 4 packets, enter a port number in this field. The range is 1 to 65535.

Table 61. IPv4 Policy Settings Window (Continued)

Field	Description
Policy Sequence	<p>Enter a sequence number for the policy. The range is 1 to 65535. Sequence numbers control the order in which policies are performed on ports that have more than one policy. Policies are performed from lowest to highest sequence numbers. Packets are forwarded or discarded at the first match. Consequently, for ports with both permit and deny policies, the permit policies should be assigned lower sequence numbers than the deny policies so that they are performed first. Otherwise, ports might discard packets you want them to forward.</p> <p>Policies will be easier to manage if you make their index numbers and policy sequence numbers the same. This value is required.</p>
Policy Action	<p>Select the action the switch performs on ingress packets that match the policy. The choices are listed here:</p> <ul style="list-style-type: none"> - Permit: Forward packets. - Deny: Discard packets.
Replaced CoS	<p>To replace the CoS value in ingress packets that match the policy, click the radio button and enter the new number. the range is 0 to 7. This option only applies to permit policies.</p>
Replaced DSCP	<p>To replace the DSCP value in ingress packets that match the policy, click the radio button and enter the new number. the range is 0 to 63. This option only applies to permit policies.</p>
Port List	<p>Enter the switch ports to which the policy is to be assigned. The port list can be specified as a consecutive list, a non-consecutive list, or a combination of the two. For example, you can specify ports 1-3,5,8. A policy has to have at least one port.</p> <p>You cannot mix individual ports and ports of a port trunk in a port list. For example, if ports 3 and 4 are members of a trunk, you may not assign ports 1-4 in the port list, but you may assign ports 3 and 4.</p>

Note

Policy Index, Policy Sequence, and Port List are required parameters.

4. Click the **Add** button to add the policy to the switch or the **Cancel** button to cancel the procedure. The new policy is added to the table at the bottom of the window. The initial status is disabled.
5. To activate the policy, click its **Enable** button in the Action column in the table.
6. Select **Save Settings to Flash** from the main menu to save your changes.

Table 62 lists the IP protocol numbers for the Protocol parameter.

Table 62. IP Protocol Numbers

Number	Description
1	Internet Control Message (RFC792)
2	Internet Group Management (RFC1112)
3	Gateway-to-Gateway (RFC823)
4	IP in IP (RFC2003)
5	Stream (RFC1190 and RFC1819))
6	TCP (Transmission Control Protocol) (RFC793)
8	EGP (Exterior Gateway Protocol) (RFC888)
9	IGP (Interior Gateway Protocol) (IANA)
11	Network Voice Protocol (RFC741)
17	UDP (User Datagram Protocol) (RFC768)
20	Host monitoring (RFC869)
27	RDP (Reliable Data Protocol) (RFC908)
28	IRTP (Internet Reliable Transaction Protocol) (RFC938)
29	ISO-TP4 (ISO Transport Protocol Class 4) (RFC905)

Table 62. IP Protocol Numbers (Continued)

Number	Description
30	Bulk Data Transfer Protocol [RFC969]
33	DCCP (Datagram Congestion Control Protocol) [RFC4340]
48	DSR (Dynamic Source Routing Protocol) [RFC4728]
50	ESP (Encap Security Payload) [RFC2406]
51	AH (Authentication Header) [RFC2402]
54	NARP (NBMA Address Resolution Protocol)[RFC1735]
58	ICMP for IPv6 [RFC1883]
59	No Next Header for IPv6 [RFC1883]
60	Destination Options for IPv6 [RFC1883]
88	EIGRP (Enhanced Interior Gateway Routing Protocol)
89	OSPFv2 [RFC1583]
97	Ethernet-within-IP Encapsulation / RFC3378
98	Encapsulation Header / RFC1241
108	IP Payload Compression Protocol / RFC2393
112	Virtual Router Redundancy Protocol / RFC3768
134	RSVP-E2E-IGNORE / RFC3175
135	Mobility Header / RFC3775
136	UDPLite / RFC3828
137	MPLS-in-IP / RFC4023
138	MANET Protocols / RFC-ietf-manet-iana-07.txt

Table 62. IP Protocol Numbers (Continued)

Number	Description
139 - 252	Unassigned / IANA
253 - 254	Use for experimentation and testing / RFC3692
255	Reserved / IANA

Adding IPv6 Policies

Perform the following procedure to add IPv6 policies to ports:

1. Select **Access Control Config > Policy Settings** from the main menu. Refer to Figure 151 on page 317.
2. Click the **Add IPv6** button in the Policy Type field at the top of the window. The IPv6 Policy Settings window is shown in Figure 154.

Policy Settings

Policy type: Add L2+IPv4 Add IPv6

Policy Index: (1-65535)

VLAN ID: (1-4093) 802.1p Priority: (0-7)

Next Header:

IPv6 Source IP Address: Prefix Length: (1-128)

IPv6 Destination IP Address: Prefix Length: (1-128)

IPv6 Traffic Class: (0-255)

Source Layer 4 Port: (1-65535) Destination Layer 4 Port: (1-65535)

Policy Sequence: (1-65535)

Policy Action: Permit ▼

Replaced-CoS: (0-7) Rate Control Index: (1-65535)

Port List: (e.g. 1,3,5-8)

Add Cancel

Free Policies : 188

Total Entries : 0 Delete All

Index	Classifier	Sequence	Deny/Permit	Replaced-CoS	Replaced-DSCP	Rate Control	Port List	Status	Action
<< Table is empty >>									

Page 0 / 0 First Page Previous Page Next Page Last Page Page Go

Figure 154. IPv6 Policy Settings Window

3. Configure the IPv6 policy fields in Table 63 on page 326.

Table 63. IPv6 Policy Settings Window

Field	Description
Policy Index	Enter a unique index number for the policy. A policy can have only one index number. Index numbers determine the order in which the switch displays the policies in the table at the bottom of the window. The range is 1 to 65535. Policies will be easier to manage if you make their index numbers and policy sequence numbers the same. This setting is required.
VLAN ID	To filter tagged packets based on a VID of a tagged VLAN, enter the VID in this field. You can enter only one VID. The range is 1 to 4093. To view the tagged VIDs on the switch, refer to “Adding or Viewing 802.1Q Tagged VLANs” on page 212 or “Viewing Port-based and 802.1Q Tagged VLANs” on page 223. This filter does not apply to Port-based VLANs.
802.1p Priority	To filter packets based on an 802.1p priority level in packet headers, enter the level in this field. The range is 0 to 7.
IPv6 Source IP Address	<p>To filter packets based on a source IPv6 address, enter the address in this field. You can enter only one address. The format for an IPv6 address is shown here:</p> <p>xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx</p> <p>Each x is a hexadecimal digit representing four bits. The eight groups of digits have to be separated by colons. Groups where all four digits are '0' can be omitted. Leading '0's in groups can also be omitted.</p> <p>As an example, the following IPv6 addresses are equivalent:</p> <p>3710:421e:09a8:0000:0000:0000:00a4:1c50</p> <p>3710:421e:9a8::a4:1c50</p>
Prefix Length	Enter the prefix length of the source address as a decimal number that represents the number of bits in the address, from left to right, on which you want to filter. The range is 1 to 128. To filter on a whole address for a specific network device, enter 128 as the prefix.

Table 63. IPv6 Policy Settings Window (Continued)

Field	Description
IPv6 Destination IP Address	To filter packets based on a destination IPv6 address, enter the address in this field. You can enter only one address.
Prefix Length	Enter the prefix length of the destination address as a decimal number that represents the number of bits in the address, from left to right, on which you want to filter. The range is 1 to 128. To filter on a whole address for a specific network device, enter 128 as the prefix.
IPv6 Traffic Class	To filter packets based on the IPv6 traffic class, enter a value from 0 to 255.
Source Layer 4 Port	To filter source Transport Layer 4 packets, enter a port number in this field. The range is 1 to 65535.
Destination Layer 4 Port	To filter destination Transport Layer 4 packets, enter a port number in this field. The range is 1 to 65535.
Policy Sequence	<p>Enter a sequence number for the policy. The range is 1 to 65535. Sequence numbers determine the order in which policies are performed when a port has more than one policy. Policies are performed from lowest to highest sequence number. Packets are forwarded or discarded at the first match. Consequently, for ports with both permit and deny policies, the permit policies should be assigned lower sequence numbers than the deny policies so that they are performed first. Otherwise, ports might discard packets you want them to forward.</p> <p>Policies will be easier to manage if you make their index numbers and policy sequence numbers the same. This value is required.</p>
Policy Action	<p>Select the action that the switch is to perform on ingress packets that match the policy. The choices are listed here:</p> <ul style="list-style-type: none"> - Permit: Forward the packets. - Deny: Discard the packets.

Table 63. IPv6 Policy Settings Window (Continued)

Field	Description
Replaced CoS	To replace the CoS value in packets that match the policy, click the radio button and enter the new number. The range is 0 to 7. This option only applies to permit policies.
Rate Control Index	To apply a Committed Information Rate (CIR) value to the ingress traffic that matches the policy, enter the corresponding CIR index number from the Rate Control Settings window. The range is 1 to 65535. The entry must already exist in the Rate Control Settings window, in “Managing Rate Control Settings” on page 333. This option only applies to permit policies.
Port List	<p>Enter the ports for the policy. The port list can be entered as a consecutive list, a non-consecutive list, or a combination of the two. For example, you can specify ports 1-3,5,8. A policy has to have at least one port.</p> <p>You cannot mix individual ports and ports of a port trunk in a port list. For example, if ports 3 and 4 are members of a trunk, you cannot assign ports 1-4 in a port list because ports 1 and 2 are individual ports and ports 3 and 4 part of a port trunk. However, you can assign ports 3 and 4 to a separate port list.</p>

Note

Policy Index, Policy Sequence, and Port List are required parameters.

4. Click the **Add** button to add the policy to the switch or the **Cancel** button to cancel the procedure. The new policy is added to the table at the bottom of the window. Its initial status is disabled.
5. To activate the policy on the ports, click the **Enable** radio button in the table Status column of the policy to be enabled.
6. Click the **OK** button.
7. Select **Save Settings to Flash** from the main menu to save your changes.

Enabling or Disabling Policies

Perform the following procedure to enable or disable policies:

1. Select **Access Control Config > Policy Settings** from the main menu. Refer to Figure 151 on page 317 and Figure 152 on page 317.
2. In the table at the bottom of the window, click the **Enable** or **Disable** radio button in the Status column of the policy you want to enable or disable.
3. Click the **OK** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying or Deleting Policies

Perform the following procedure to modify or delete policies:

1. Select **Access Control Config > Policy Settings** from the main menu. Refer to Figure 151 on page 317 and Figure 152 on page 317.
2. To modify a policy, do the following:
 - a. In the table at the bottom of the window, click the **Modify** button in the Action column of the policy you want to modify. You can modify only one policy at a time.
 - b. Change the fields as needed. Refer to Table 61 on page 318 for IPv4 policies or Table 63 on page 326 for IPv6 policies. You cannot change the policy index number.
3. To delete policies, do one of the following:
 - To delete a policy, click its **Delete** button in the Action column.
 - To delete all the policies, click the **Delete All** button.
4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Displaying Policies

The tables at the bottom of the Policy Settings windows lists all the policies on the switch. The table columns, however, do not include all possible IPv4 or IPv6 policy filters. To view policies with all their filters, perform “Modifying or Deleting Policies” on page 330, or this procedure:

1. Select **Access Control Config > Policy Settings** from the main menu. Refer to Figure 151 on page 317.
2. Click the **Add L2 IPv4** or **Add IPv6** button in the Policy Type field. The default is IPv4. Review to Figure 152 on page 317.
3. In the table at the bottom of the window. click the **Detail** button in the Classifier column of the policy you want to view. Refer to Table 61 on page 318 for descriptions of IPv4 policies or Table 63 on page 326 for IPv6 policies.
4. Click the **Back** button to return to the Policy Settings window.

Displaying Policies by Port Numbers

Perform the following procedure to view policies assigned to selected ports:

1. Select **Access Control Config > Policy Database** from the main menu. Refer to Figure 155.




Figure 155. Policy Database Menu Selection

The window is shown in Figure 156.

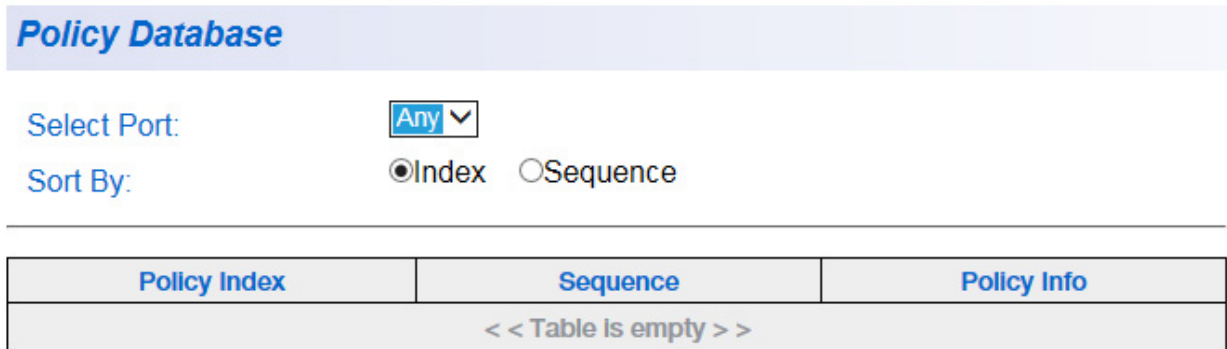


Figure 156. Policy Database Window

2. Select a switch port from the **Select Port** menu. Select **Any**, the default, to display all the policies on the switch.
3. Click either the **Index** or **Sequence** radio button to view the port policies by index numbers or policy sequence numbers: Index numbers control the order in which policies are listed in the tables. Policy sequence numbers determine the order in which the switch compares packets against policies on ports with more than one policy.
4. To view information about a policy, click its **Detail** button in the Policy Info column. For parameter descriptions, refer to Table 61 on page 318 for IPv4 policies or Table 63 on page 326 for IPv6 policies.
5. Click the **Back** button to return to the Policy Database page.

Managing Rate Control Settings

IPv6 policies can include a Rate Control Index parameter. It defines a Committed Information Rate (CIR) that sets a fixed bandwidth, in bits per second, for ingress traffic on ports. The parameter lets you define individual bandwidths for different virtual connections that share the same physical path, but that have different bandwidth requirements. For example, you can assign a high CIR to a connection with a high proportion of video signals to maintain signal quality.

Here are the procedures for managing rate control entries:

- “Adding Rate Control Entries,” next
- “Modifying or Deleting Rate Control Entries” on page 334

Adding Rate Control Entries

Perform the following procedure to add rate control entries:

1. Select **Access Control Config > Rate Control Settings** from the main menu. Refer to Figure 157.

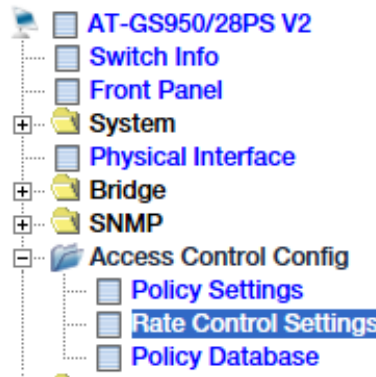


Figure 157. Rate Control Settings Menu Selection

The window is shown in Figure 158 on page 334.

Rate Control Settings

Index: (1-65535)

Committed Rate: 64kbps x (1-15625)

Free Entries : 92

Total Entries : 0

Index	Committed Rate	Action
<< Table is empty >>		

Page 0 / 0 Page

Figure 158. Rate Control Settings Window

2. Enter a unique number in the **Index** field in the range of 1 to 65535. The number determines the order in which rate control settings are displayed in the table and identified in the Rate Control Index entry in IPv6 policies.
3. Enter a number in the **Committed Rate** field in the range of 1 to 15625 in 64kbps increments to specify the bandwidth.
4. Click the **Add** button. The rate control entry is added to the table.
5. Select **Save Settings to Flash** from the main menu to save your changes.
6. To add the rate control setting to an IPv6 policy, refer to “Adding IPv6 Policies” on page 325 or “Modifying or Deleting Policies” on page 330.

Modifying or Deleting Rate Control Entries

Perform the following procedure to modify or delete rate control entries:

Note

You cannot delete rate control entries that are assigned to access control policies. You have to remove them from their policies first. Refer to “Modifying or Deleting Policies” on page 330.

1. Select **Access Control Config > Rate Control Settings** from the main menu. Refer to Figure 157 on page 333 and Figure 158 on page 334.

2. To modify a rate control entry, do the following:
 - a. Click the **Modify** button in the Action column of the entry you want to delete. You can modify only one entry at a time.
 - b. Enter the new rate in the **Committed Rate 64kbps x** field.
 - c. Click the **Apply** button in the Action column.
3. To delete entries, do one of the following:
 - To delete a single entry, click its **Delete** button in the Action column.
 - To delete all the entries, click the **Delete All** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Bandwidth Control

This chapter describes bandwidth control in the following sections:

- ❑ “Setting Threshold Limits for Ingress Unknown Unicast, Broadcast, and Multicast Packets” on page 338
- ❑ “Setting Ingress Traffic Rate Limits” on page 340
- ❑ “Setting Egress Traffic Rate Limits” on page 342

Setting Threshold Limits for Ingress Unknown Unicast, Broadcast, and Multicast Packets

You can set maximum threshold limits for the number of ingress packets that ports will accept per second. The switch discards Ingress packets that exceed the thresholds. You can set maximum thresholds for the following types of traffic:

- Unicast packets (destination lookup failures)
- Broadcast packets
- Multicast packets

The formula for calculating the bandwidth limit for the copper ports is as follows:

$$\text{Bandwidth} = 64\text{pps} \times \text{increment}$$

The increment is an integer ranging from 1 to 22194.

This feature supports packet sizes from 64 to 22194 bytes.

Perform the following procedure to set threshold limits for unknown unicast, broadcast, or multicast ingress packets:

1. Select **Bridge > Bandwidth Control > Storm Control** from the main menu. Refer to Figure 159.

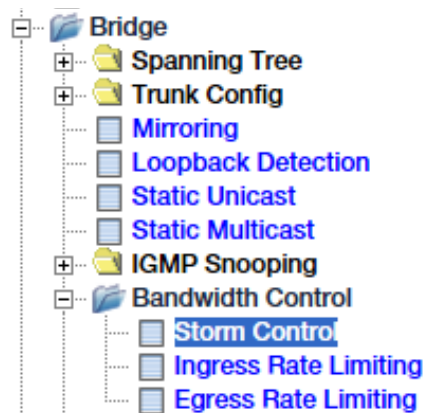


Figure 159. Storm Control Menu Selection

The Storm Control window is shown in Figure 160.

Storm Control					
Port	DLF	Broadcast	Multicast	Threshold	Action
All	Ignore ▾	Ignore ▾	Ignore ▾	64pps x <input type="text" value="1"/> (1-22194)	Apply
1	Disabled ▾	Disabled ▾	Disabled ▾	64pps x <input type="text" value="22194"/> (1-22194)	Apply
2	Disabled ▾	Disabled ▾	Disabled ▾	64pps x <input type="text" value="22194"/> (1-22194)	Apply
3	Disabled ▾	Disabled ▾	Disabled ▾	64pps x <input type="text" value="22194"/> (1-22194)	Apply
4	Disabled ▾	Disabled ▾	Disabled ▾	64pps x <input type="text" value="22194"/> (1-22194)	Apply

Figure 160. Storm Control Window

2. Select **Enabled** or **Disabled** from the **DLF** (unknown unicast packets), **Broadcast**, or **Multicast** menus to enable or disable the threshold limits for the individual ports. The default is disabled.
3. In the **Threshold** column, enter the increment used to calculate the threshold value. Here are guidelines:
 - Thresholds are calculated by multiplying 64 packets per second (pps) by the increment. For example, an increment of 1 results in a threshold of 64pps, an increment of 2 results in a threshold of 128pps, and so on.
 - The increment range is 1 to 22194.
 - A port has to have at least one enabled filter for you to set the threshold.
4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Setting Ingress Traffic Rate Limits

This section explains how to set threshold limits on the ingress traffic that ports will accept from their local or remote network device counterparts. Ingress traffic exceeding the threshold limits are discarded. The traffic is defined as kilobits per second (Kbps). The range is from 64Kbps to 1000Mbps, in increments of 64Kbps. You might use the feature to control the ingress bandwidths of the switch ports so as to protect their ingress queues from becoming overwhelmed or saturated by their network device counterparts during periods of heavy traffic. The formula for calculating the bandwidth limit for 10/100/1000Base-T ports is as follows:

$$\text{Bandwidth} = 64\text{Kbps} \times \text{increment}$$

The increment is an integer ranging from 1 to 15625.

Perform the following steps to configure port ingress traffic rate limits:

1. Select **Bridge > Bandwidth Control > Ingress Rate Limiting** from the main menu. Refer to Figure 161.

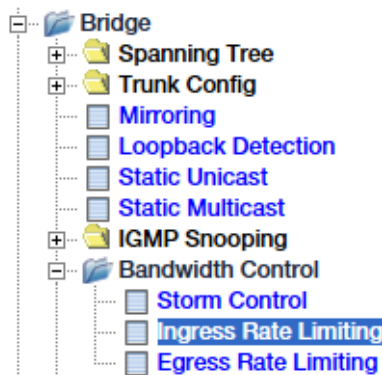


Figure 161. Ingress Rate Limiting Menu Selection

The Ingress Rate Limiting window is shown on Figure 162.

Ingress Rate Limiting

Bandwidth = 64kbps x rate limit

Port	Bandwidth	Status	Action
All	64kbps x <input type="text"/> (1-15625)	Ignore <input type="button" value="v"/>	<input type="button" value="Apply"/>
1	64kbps x <input type="text" value="15625"/> (1-15625)	Disabled <input type="button" value="v"/>	<input type="button" value="Apply"/>
2	64kbps x <input type="text" value="15625"/> (1-15625)	Disabled <input type="button" value="v"/>	<input type="button" value="Apply"/>
3	64kbps x <input type="text" value="15625"/> (1-15625)	Disabled <input type="button" value="v"/>	<input type="button" value="Apply"/>
4	64kbps x <input type="text" value="15625"/> (1-15625)	Disabled <input type="button" value="v"/>	<input type="button" value="Apply"/>

Figure 162. Ingress Rate Limiting Window

2. In the Status column, select **Enabled** or **Disabled** from the menus to enable or disable the thresholds on the ports. The default is disabled.
3. In the Threshold column, enter the increment for calculating the threshold value. Here are guidelines:
 - Thresholds are calculated by multiplying 64 kilobits per second (Kbps) by the increment. For example, an increment of 1 results in a threshold of 64Kbps, an increment of 2 results in a threshold of 128Kbps, and so on.
 - You have to enable the threshold on a port to set its threshold increment.
 - The increment range is 1 to 15625.
 - Ports can have different thresholds.
4. Click the **Apply** button in the Action column.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Setting Egress Traffic Rate Limits

This section explains how to set limits on the egress traffic that ports transmit to their local or remote network counterparts. The traffic is defined as kilobits per second (Kbps) and the range is from 64Kbps to 1000Mbps, in increments of 64Kbps. The feature provides a way to control the egress bandwidths of the switch ports so as to prevent them being overwhelmed or saturated by their network counterparts during periods of heavy traffic. The formula for calculating the bandwidth limit for 10/100/1000Base-T ports is as follows:

$$\text{Bandwidth} = 64\text{Kbps} \times \text{increment}$$

The increment is an integer ranging from 1 to 15625.

Perform the following steps to configure egress traffic port rate limits:

1. Select **Bridge > Bandwidth Control > Egress Rate Limiting** from the main menu. Refer to Figure 163.

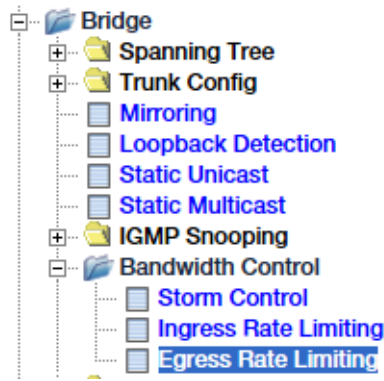


Figure 163. Egress Rate Limiting Menu Selection

The Egress Rate Limiting window is shown in Figure 164.

Egress Rate Limiting

Bandwidth = 64kbps x rate limit

Port	Bandwidth	Status	Action
All	64kbps x <input type="text" value=""/> (1-15625)	Ignore ▼	Apply
1	64kbps x <input type="text" value="15625"/> (1-15625)	Disabled ▼	Apply
2	64kbps x <input type="text" value="15625"/> (1-15625)	Disabled ▼	Apply
3	64kbps x <input type="text" value="15625"/> (1-15625)	Disabled ▼	Apply
4	64kbps x <input type="text" value="15625"/> (1-15625)	Disabled ▼	Apply

Figure 164. Egress Rate Limiting Window

2. In the Status column, select **Enabled** or **Disabled** from the pull-down menus to enable or disable the thresholds on the ports. The default is disabled.
3. In the Threshold column, enter the increment for calculating the threshold value. Here are guidelines:
 - Thresholds are calculated by multiplying 64 kilobits per second (Kbps) by the increment. For example, an increment of 1 results in a threshold of 64Kbps, an increment of 2 results in a threshold of 128Kbps, and so on.
 - The increment range is 1 to 15625.
 - You have to enable the threshold on a port to set the threshold increment.
 - Ports can have different thresholds.
4. Click the **Apply** button in the Action column.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 28

RMON

This chapter contains the following sections:

- “RMON Overview” on page 346
- “Enabling or Disabling RMON” on page 347
- “RMON Port Statistics Groups” on page 348
- “RMON Histories” on page 350
- “RMON Events” on page 353
- “RMON Alarms” on page 356

RMON Overview

The RMON (Remote MONitoring) MIB is used with SNMP applications to monitor the operations of network devices. The switch supports the four RMON MIB groups listed here:

- ❑ **Statistic groups** — Collects port statistics. Refer to “RMON Port Statistics Groups” on page 348.
- ❑ **History groups** — Collects histories of port statistics. Refer to “RMON Histories” on page 350.
- ❑ **Event groups** — Defines the actions of RMON alarms that switches perform when packet statistic thresholds are crossed. Refer to “RMON Events” on page 353.
- ❑ **Alarm groups** — Defines the statistics thresholds that trigger alarms and alarm actions. Refer to “RMON Alarms” on page 356.

For instructions on how to configure SNMP on your switch, refer to Chapter 31, “SNMPv1 and v2c” on page 385 or Chapter 32, “SNMPv3” on page 397.

Enabling or Disabling RMON

RMON allows you to monitor the statistics and histories of switch ports, with SNMP Network Management System (NMS) software and the RMON section of the MIB tree. You can also use RMON to add alarms in the forms of log entries or SNMP traps that alert you when port traffic cross defined thresholds.

Because RMON requires SNMP, the SNMP agent on the switch has to be enabled for the RMON feature to be active. Refer to Chapter 31, “SNMPv1 and v2c” on page 385 or Chapter 32, “SNMPv3” on page 397.

Perform the following procedure to enable or disable RMON:

1. Select **RMON > Global Settings** from the main menu. Refer to Figure 165.



Figure 165. RMON Global Settings Menu Selection

The RMON Basic Settings window is shown in Figure 166.




Figure 166. RMON Basic Settings window

2. Select one of the following from the RMON Status menu:
 - Enabled:** Enables the RMON feature.
 - Disabled:** Disable the RMON feature. This is the default setting.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

RMON Port Statistics Groups

Port statistics groups are used for the following functions:

- ❑ To view port statistics from the RMON portion of the MIB tree with SNMP NMS software.
- ❑ Combined with RMON alarms to alert you to defined traffic events.

Adding RMON Port Statistics Groups

Perform the following procedure to add RMON statistics groups to ports:

1. Select **RMON > Statistics** from the main menu. Refer to Figure 167.




Figure 167. RMON Statistics Menu Selection

The Ethernet Statistics Settings window is shown in Figure 168.

Ethernet Statistics Settings

Index (1~65535): *

Port: *

Owner: (32 characters limit)

Ethernet Statistics Table :

Index	Port	Drop Events	Octets	Packets	Broadcast Packets	Multicast Packets	Owner	Action
<< Table is empty >>								

Figure 168. Ethernet Statistics Settings Window

The table in the window lists the current Ethernet statistics groups, with the collected statistics.

2. To add a new RMON statistics group, configure the fields in Table 64 on page 349.

Table 64. Ethernet Statistics Settings Window

Parameter	Description
Index	Enter a unique ID number for the new statistics group. The range is 1 to 65535. Some SNMP programs identify statistics groups by their ID numbers and not by port numbers. Consequently, statistics groups will be easier to identify if you make their ID numbers the same as the port numbers. For instance, a group assigned to port 16 should be assigned the ID number 16.
Port	Enter the port to be monitored. You can enter only one port.
Owner	Enter the name of the person responsible for this entry. This optional parameter is for switches that are managed by more than one person.

3. Click the **Add** button to add the new entry to the switch. New statistics groups become active as soon as you add them.
4. To add RMON statistics groups for other ports, repeat steps 2 and 3.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying or Deleting RMON Port Statistics Groups

Perform the following procedure to modify or delete port statistics groups:

1. Select **RMON > Statistics** from the main menu. Refer to Figure 167 on page 348 and Figure 168 on page 348.
2. To modify a statistics group, do the following:
 - a. In the table at the bottom of the window, click the **Modify** button in the Action column of the group you want to modify. You can modify only one group at a time.
 - b. Modify the fields as needed. Refer to Table 64 on page 349. You cannot change the policy index number.
3. To delete groups, do one of the following:
 - To delete a group, click its **Delete** button in the Action column.
 - To delete all the groups, click the **Delete All** button.
4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

RMON Histories

RMON histories are snapshots of port statistics taken by the switch at predefined intervals. By comparing the snapshots you can identify trends or patterns in the numbers or types of ingress packets on the ports. The snapshots can be viewed with SNMP NMS software in the RMON portion of the MIB tree.

Adding RMON History Groups

Perform the following procedure to add RMON histories:

1. Select **RMON > History** from the main menu. Refer to Figure 169.




Figure 169. RMON History Menu Selection

The History Control Settings window is shown in Figure 170.

History Control Settings

Index (1~65535): *

Port: *

Buckets Requested (1~50):

Interval (1~3600): sec

Owner: (32 characters limit)

History Control Table:

Index	Port	Buckets Requested	Buckets Granted	Interval	Owner	Action
< < Table is empty > >						

Figure 170. History Control Settings Window

The table in the window lists the current RMON histories. The table columns are defined in Table 65 on page 351. The Buckets Granted column lists the number of buckets for a history group the switch can support in its free memory. Depending on the number of histories and requested buckets, the switch might not be able to support all of the requested buckets.

2. Configure the fields in Table 65.

Table 65. History Control Settings Window

Field	Description
Index	<p>Enter a unique ID number of the new group. The range is 1 to 65535.</p> <p>Some SNMP programs identify history groups by their ID numbers and not by the port numbers. Consequently, the groups will be easier to identify if you make their ID numbers the same as the port numbers. For instance, a history group assigned to port 16 should be assigned the ID number 16.</p>
Port	Enter the port to be monitored. You can enter only one port.
Buckets Requested	Enter the maximum number of buckets for storing snapshots of the port's statistics. Each bucket can store one snapshot of RMON statistics of one port. The more buckets in a group, the more snapshots it can store. Different ports can have different numbers of buckets. The range is 1 to 50 buckets.
Interval	Enter how frequently the switch is to take snapshots of the port's statistics. The range is 1 to 3600 seconds (1 hour). For example, enter 60 seconds to have the switch take a snapshot once a minute on a port.
Owner	Enter the name of the person responsible for the entry. This optional parameter is for switches that are managed by more than one person.

3. Click the **Add** button to add the new RMON history group.

After you add a new group, the switch determines whether it has sufficient free memory to add all the requested buckets. If it does not have enough memory, it reduces the number of buckets to a supportable amount. If it does not have any available free memory, it cancels the history group.

The switch takes the first snapshot at the end of the first interval. A history group that has an interval of 1800 seconds, for instance, does not take its first snapshot for 30 minutes. Once all the buckets of a group are full, the switch continues storing snapshots by deleting the oldest snapshots as it adds new snapshots. For instance, for a history group of three buckets, the switch deletes the first bucket when it adds the fourth bucket.

4. Repeat steps 2 and 3 to add additional RMON histories for other ports.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying or Deleting RMON History Groups

Perform the following procedure to modify or delete RMON history groups:

1. Select **RMON > History** from the main menu. Refer to Figure 169 on page 350. The History Control Settings window is shown in Figure 170 on page 350.
2. To modify a history group, do the following:
 - a. In the table at the bottom of the window, click the **Modify** button in the Action column of the group you want to modify. You can modify only one group at a time.
 - b. Change the fields as needed. Refer to Table 65 on page 351. You cannot change the policy index number.
3. To delete history groups, do one of the following:
 - To delete a group, click its **Delete** button in the Action column.
 - To delete all the groups, click the **Delete All** button.
4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

RMON Events

RMON events define the actions of RMON alarms. They specify the action the switch performs when port statistics cross defined thresholds, triggering an alarm. Events can log a message in the event log of the switch, send an SNMP trap to an SNMP network device, or both. Events can be used with more than one alarm. The switch supports up to 256 events.

Adding Events Perform the following procedure to add RMON events.

1. Select the **RMON > Event** from the main menu. Refer to Figure 171.




Figure 171. RMON Event Menu Selection

The RMON Event Settings window is shown in Figure 172.

RMON Event Settings

Index(1-65535): *

Description: * (32 characters limit)

Type: ▼

Community:

Owner: (32 characters limit)

Free Entries: 256
Total Entries: 0

Index	Description	Type	Community	Owner	Last Time Sent	Action
<< Table is empty >>						

Page 0 / 0 Page

Figure 172. RMON Event Settings Window

The table in the window lists the current RMON events. The columns are described in Table 66. The Last Time Sent column contains the date and time when an event's alarm was last triggered.

2. To add a new RMON event, configure the fields in Table 66.

Table 66. RMON Event Settings Window

Field	Description
Index	Enter a unique ID number of the new group. The range is 1 to 65535.
Description	Enter a description of the event, up to 32 characters.
Type	Select the action that the switch is to perform when it determines that a statistics threshold of an alarm has been crossed. The options are listed here: <ul style="list-style-type: none"> - None: Take no action. - Log: Log a message in the event log. - SNMP: Send an SNMP trap. - Both: Log a message in the event log and send an SNMP trap.
Community	Enter the SNMP trap community name to receive traps. Here are the guidelines: <ul style="list-style-type: none"> - This parameter is required when Type is SNMP or Both. - The community name has to already exist on the switch. Refer to "Adding SNMP Community Strings" on page 390. - The community name is case sensitive. - An event can have only one community name.
Owner	Enter the name of the person responsible for the entry. This optional parameter is for switches that are managed by more than one person.

3. Click the **Add** button to add the RMON event.
4. Repeat steps 2 and 3 to add more RMON events.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying or Deleting Events

Perform the following procedure to modify or delete events:

1. Select **RMON > Event** from the main menu. Refer to Figure 171 on page 353. The RMON Event Settings window is shown in Figure 172 on page 353.
2. To modify an event, do the following:
 - a. Click its **Modify** button in the Action column in the table. You can modify only one event at a time.
 - b. Change the fields as needed. Refer to Table 66 on page 354. You cannot change the event index number.
3. To delete events, do one of the following:
 - To delete a event, click its **Delete** button in the Action column.
 - To delete all the events, click the **Delete All** button.

Note

You cannot delete events when they are assigned to alarms. You have to remove the events from their alarms or delete the alarms. Refer to “Modifying or Deleting Alarms” on page 359.

4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

RMON Alarms

RMON alarms are used to generate alert messages when packet activity on designated ports rises above or falls below defined threshold values. The alert messages can take the form of messages that are entered in the event log on the switch, traps that are sent to your SNMP network devices, or both. You can use alarms to alert you when traffic flows on ports raise above or fall below defined thresholds. The switch supports up to eight alarms.

RMON alarms have rising threshold and falling thresholds. An alarm is triggered if the value of the monitored RMON statistic of the designated port raises above the rising threshold or falls below the falling threshold. The response of the switch is determined by an RMON event. The response can be to enter a message in the event log, send an SNMP trap, or both.

The frequency with which the switch samples the thresholds of an alarm against the actual RMON statistic is controlled by a time interval parameter. You can adjust this interval for each alarm.

RMON alarms have the following three components:

- ❑ RMON statistics group: A port has to have an RMON statistics group to have an alarm. When you add an alarm, you specify the port to which it is to be assigned, not by the port number, but rather by the ID number of the port's statistics group. (As explained in "RMON Port Statistics Groups" on page 348, statistics groups are also used to remotely view port statistics in the RMON portion of the MIB tree.)
- ❑ RMON event: An event specifies the action of the switch when the ingress packet activity on a port crosses a statistical threshold defined in an alarm. The choices are to log a message in the event log of the switch, send an SNMP trap to an SNMP network device, or both.
- ❑ Alarm: The last component is the alarm itself. It defines the port statistic to be monitored and the rising and falling thresholds that trigger the switch to perform the event. The thresholds of an alarm can have the same event or different events.

Adding RMON Alarms

Perform the following procedure to add RMON alarms.

1. Select **RMON > Alarm** from the main menu. Refer to Figure 173.




Figure 173. RMON Alarms Menu Selection

The RMON Alarm Settings page is shown in Figure 174.

RMON Alarm Settings

Index (1~65535): *

Interval (1~2³¹-1): sec

Variable: *

Sample type: ▾

Rising Threshold (0~2³¹-1): *

Falling Threshold (0~2³¹-1): *

Rising Event Index (1~65535):

Falling Event Index (1~65535):

Owner: (32 characters limit)

Free Entries: 256
Total Entries: 0

Index	Interval	Variable	Sample Type	Rising Threshold	Falling Threshold	Rising Event Index	Falling Event Index	Owner	Action
<< Table is empty >>									

Page 0 / 0 Page

Figure 174. RMON Alarm Settings Window

The table in the window lists the current RMON alarms. The table columns are defined in Table 67.

2. To add a new alarm, configure the fields in Table 67.

Table 67. RMON Alarm Settings Window

Field	Description
Index	Enter a unique ID number for the new alarm. The range is 1 to 65535.
Interval	Enter the time period (in seconds) the switch waits to poll the statistics group to determine whether a threshold has been crossed. The range is 1 to 2 ³¹ -1 (2147483647) seconds.
Variable	Enter the RMON MIB object that the event is to monitor.

Table 67. RMON Alarm Settings Window (Continued)

Field	Description
Sample Type	Select the change to the port statistic that triggers the alarm: Here are the choices: <ul style="list-style-type: none"> - Absolute Value: Compares the current statistic value against the threshold. - Delta Value: Compares the difference between the previous and current statistic values against the threshold.
Rising Threshold	Enter the rising threshold for the monitored statistics that, when crossed, causes the switch to perform the specified rising event. The range is 0 to $2^{31}-1$ (2147483647).
Falling Threshold	Enter the falling threshold for the monitored statistics that, when crossed, causes the switch to perform the specified falling event. The range is 0 to $2^{31}-1$ (2147483647).
Rising Event Index	Enter the event index for the rising threshold. The event determines the action of the switch when the monitored statistic exceeds the rising threshold. The range is 1 to 65535. The specified event should already exist on the switch. Refer to "RMON Events" on page 353. Rising and falling thresholds can use the same event. This field is required.
Falling Event Index	Enter the event index for the falling threshold. The event determines the action of the switch when the monitored statistic falls below the falling threshold. The range is 1 to 65535. The specified event should already exist on the switch. Refer to "RMON Events" on page 353. The rising and falling thresholds can use the same event. This field is required.
Owner	Enter the name of the person responsible for this entry. This optional parameter is for switches that are managed by more than one person.

3. Click the **Apply** button to add the new alarm to the table.
4. Repeat steps 2 and 3 to add more RMON alarms.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying or Deleting Alarms

Perform the following procedure to modify or delete alarms:

1. Select **RMON > Alarms** from the main menu. Refer to Figure 173 on page 357. The RMON Event Settings window is shown in Figure 174 on page 357.
2. To modify an alarm, do the following:
 - a. Click its **Modify** button in the Action column in the table. You can modify only one alarm at a time.
 - b. Change the fields as needed. Refer to Table 67 on page 357. You cannot change the event index number.
3. To delete alarms, do one of the following:
 - To delete an alarm, click its **Delete** button in the Action column.
 - To delete all the alarms, click the **Delete All** button.
4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Section V

Port Trunks

This section contains the following chapters:

- ❑ Chapter 29, “Static Port Trunks” on page 363
- ❑ Chapter 30, “LACP Trunks” on page 371

Chapter 29

Static Port Trunks

This chapter describes static port trunks in the following sections:

- ❑ “Static Port Trunk Overview” on page 364
- ❑ “Adding Static Port Trunks” on page 366
- ❑ “Modifying Static Port Trunks” on page 368
- ❑ “Deleting Static Port Trunks” on page 370

Static Port Trunk Overview

Static port trunks are an economical way to increase the bandwidth between the Ethernet switch and another networking device, such as a network server, router, workstation, or another Ethernet switch. They consist of two or more ports that function as a single virtual link between the switch and another device. Port trunks improve performance by distributing network traffic across multiple ports between the devices and enhance reliability by reducing the reliance on a single physical link. They are commonly used in situations where the bandwidth of a single physical link between devices is not sufficient to efficiently handle the traffic load.

The example in Figure 175 illustrates a port trunk of three links between two GS950/28PS V2 Switches.




Figure 175. Static Port Trunk Example

Network equipment vendors tend to employ different techniques to implement static port trunks. Consequently, a static port trunk on one device may not be compatible with the same feature on a device from a different manufacturer. For this reason, static port trunks are typically employed only between devices from the same manufacturer.

Also, note that a static port trunk does not provide for redundancy or link backup. If a port in a static trunk loses its link, the trunk's total bandwidth is reduced. Although the traffic carried by the lost link is shifted to one of the remaining ports in the trunk, the bandwidth remains reduced until the lost link is re-established or you add another port to the trunk.

Guidelines

Here are the guidelines to static port trunks:

- ❑ Allied Telesis recommends employing static port trunks between the same Allied Telesis networking devices to ensure compatibility.
- ❑ A static trunk can have up to ten ports.
- ❑ A port can belong to only one static trunk at a time.

- ❑ The ports of a static trunk must be of the same medium type. They can be all twisted-pair ports or fiber optic ports, but not both.
- ❑ Trunk ports can be consecutive (for example, ports 2 through 4) or nonconsecutive (for example, ports 3, 5, and 7).
- ❑ The port settings have to be the same for all trunk ports. This includes speed, duplex mode, flow control, back pressure settings, and VLAN membership.
- ❑ A change to the speed, duplex mode, flow control, or back pressure of a port in a trunk automatically implements the same change on all the other trunk ports.
- ❑ Trunk ports can be tagged or untagged.
- ❑ The switch selects a port in the trunk to handle broadcast packets and packets of unknown destination. The switch makes this choice based on a hash algorithm, depending on the source and destination MAC addresses.

Adding Static Port Trunks

This procedure explains how to add static port trunks.



Caution

Do not connect the cables of a port trunk until after you have configured the ports on both switches. Connecting the cables prior to configuring the ports can form loops in your network topology, which can result in broadcast storms that can reduce network performance.

To add static port trunks, perform the following procedure:

1. Select **Bridge -> Trunk Config. -> Trunking** from the main menu. Refer to Figure 176.




Figure 176. Trunking Menu Selection

The Trunking window is shown in Figure 177.

Trunking																															
Trunk ID 1:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply	
Trunk ID 2:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 3:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 4:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 5:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 6:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 7:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 8:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply

Note: Disabling will turn off the function and return all values to default.

Figure 177. Trunking Window

The switch supports up to eight LACP and static port trunks, Each row represents a separate trunk. The numbers and check boxes represent the switch ports. Ports that have check marks in the check boxes are members of a trunk.

2. Click the check boxes of the ports for the trunk. Here are the guidelines:
 - A check in a box indicates the port is a member of the trunk. No check means the port is not a member.
 - A static port trunk can have up to ten ports.
 - A port can be a member of only one trunk at a time. To add a port that is already a member of a trunk to another trunk, you first have to remove it from its current trunk.
3. Set the trunk mode status with the pull-down menu in the right column of the corresponding Trunk ID row. The choices are listed here:
 - Manual** - Activates the static port trunk.
 - Disable** - Disables the static port trunk. The ports function as regular networking ports. This is the default setting.

Note

The Active and Passive menu settings are for LACP trunks. Refer to Chapter 30, "LACP Trunks" on page 371.

4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.
6. Configure the port trunk on the other switch.
7. Connect the Ethernet cables to the trunk ports on the two switches.

Modifying Static Port Trunks

This procedure explains how to enable or disable port trunks and add or remove ports.



Caution

Before disabling or modifying a port trunk, disconnect all cables from its ports. Leaving the cables connected can create loops in your network topology, which can result in broadcast storms that can reduce network performance.

To add or remove ports from a trunk, perform the following procedure:

1. Disconnect all Ethernet cables from the ports of the trunk.
2. Select **Bridge -> Trunk Config. -> Trunking** from the main menu. Refer to Figure 176 on page 366.

The Trunking window is shown in Figure 177 on page 366. The window lists the eight possible trunks on the switch. The rows of numbers and check boxes represent the switch ports.

3. To add or remove ports, click the check boxes. Here are the guidelines:
 - A check in a box indicates the port is a member of the trunk. No check means the port is not a member.
 - A port trunk can contain up to ten ports.
 - A port can be a member of only one trunk at a time. To add a port that is already a member of a trunk to another trunk, you first have to remove it from its current trunk.
4. To enable or disable a trunk, change the status with the pull-down menu to the right column of the corresponding Trunk ID row. The choices are listed here:
 - Manual** - Activates the static port trunk.
 - Disable** - Disables the static port trunk. The ports function as regular networking ports. This is the default setting.

Note

The Active and Passive menu settings are for LACP trunks. Refer to Chapter 30, "LACP Trunks" on page 371.

5. Click the **Apply** button.

6. Select **Save Settings to Flash** from the main menu to save your changes.
7. If necessary, modify the port trunk on the other switch.
8. Connect the Ethernet cables to the trunk ports on the two switches.

Deleting Static Port Trunks



Caution

Before deleting a port trunk, disconnect all cables from its ports. Leaving the cables connected can create loops in your network topology, which can result in broadcast storms that can reduce network performance.

Perform the following procedure to delete static port trunks:

1. Disconnect all Ethernet cables from the trunk ports.
2. Select **Bridge > Trunk Config. > Trunking** from the main menu. Refer to Figure 176 on page 366. The Trunking window is shown in Figure 177 on page 366.
3. In the appropriate Trunk ID row, remove the ports from the trunk by clicking their check boxes to remove the check marks.
4. Set the status of the trunk to **Disabled** with the menu in the right column of the trunk row.
5. Click the **Apply** button.
6. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 30

LACP Trunks

This chapter describes LACP port trunks in the following sections:

- ❑ “LACP Port Trunk Overview” on page 372
- ❑ “Displaying LACP Group Status” on page 375
- ❑ “Adding LACP Trunks” on page 377
- ❑ “Setting LACP Port Priorities” on page 379
- ❑ “Modifying LACP Trunks” on page 380
- ❑ “Deleting LACP Trunks” on page 382

LACP Port Trunk Overview

LACP (Link Aggregation Control Protocol) port trunks are used to increase the bandwidth between network devices by distributing the traffic load over multiple physical links. They are useful in situations where the bandwidth of a single physical link is not sufficient to efficiently handle the traffic load between network devices.

LACP on the GS950 PS V2 Switch is compliant with the IEEE 802.3ad standard. This makes it inter-operable with equipment from other vendors that also comply with the standard, allowing for LACP trunks between GS950 PS V2 Switches and network devices from other manufacturers.

To add an LACP trunk to the switch, you designate the ports of the trunk. A trunk can have any number of ports, but only eight ports can be active at any one time in a trunk. The other ports operate in a standby mode. Ports in the standby mode do not pass network traffic, but they do transmit and accept LACP Data Unit (LACPDU) packets, which the switch uses to search for LACP compliant devices. If a link on an active LACP trunk port goes down, the switch automatically activates a standby port so that the maximum possible bandwidth of the trunk is maintained. This adds redundancy and resiliency to the trunk. For example, an LACP trunk of ten ports will have eight active ports and two standby ports. If the link on one of the active ports fails, the switch activates one of the standby ports.

The main component of an LACP trunk is the *aggregator*. It manages a group of ports on the switch. On the GS950 PS V2 Switch, the ports assigned to a trunk group are automatically assigned to an aggregator. Only one aggregator can be assigned to each trunk group, which is referred to as an *aggregate trunk*.

Only ports on the switch that are part of an aggregator transmit LACPDU packets. If a port that is part of an aggregator does not receive LACPDU packets from its corresponding port on another device, it assumes that the other port is not part of an LACP trunk and functions as a normal Ethernet port by forwarding network traffic. However, it does continue to send LACPDU packets. If it begins to receive LACPDU packets, it automatically transitions to an active or standby mode as part of an aggregate trunk.

System Priority and ID Numbers

When two network devices form an aggregate trunk, a conflict may occur if they have different LACP implementations. For example, the two devices might not support the same number of active ports in an aggregate trunk or might not agree on which ports are to be active and which are to be in the standby mode.

If a conflict does occur, the two devices need a mechanism for resolving the conflict and deciding whose LACP settings take precedence. This is accomplished with the following values:

- ❑ System priority: This value is fixed at 32768 on the GS950 PS V2 Switch. It cannot be changed.
- ❑ System ID numbers: This is the MAC address of the GS950 PS V2 Switch. It cannot be changed.

If the two switches of an LACP trunk have different system priorities, the LACP settings on the switch with the lower value takes precedence. If the two switches have the same system priorities, as will be the case with LACP trunks between GS950 PS V2 Switches, they compare system IDs. System IDs for GS950 PS V2 Switches, are their MAC addresses. The LACP settings on the switch with the lower MAC address takes precedence.

Port Priority Value

The switch uses ports' LACP priorities to determine which ports are active and which are in the standby mode in situations where the number of ports in the aggregate trunk exceeds the maximum number of allowed active ports. This parameter has a range of 1 to 65535. The lower the number, the higher the priority. Ports with the highest priorities are designated as the active ports in an aggregate trunk. The default value is a port's number. For example, the default priority values for ports 1 and 2 are 1 and 2, respectively.

As an example, if both 802.3ad-compliant devices support up to eight active ports, and there are a total of nine or more ports in the trunk, the eight ports with the highest priorities (lowest priority values) are designated as the active ports, and the others are placed in the standby mode. If an active link goes down on a active port, the standby port with the highest priority is automatically activated to take its place.

The selections of active links in an aggregate trunk are dynamic and change as links are added, removed, lost, or re-established. For example, if an active port loses its link and is replaced by a port in the standby mode, the re-establishment of the link on the originally active port causes the port to return to the active state because it has a higher priority value than the replacement port, which returns to the standby mode.

Two conditions must be met for a port in an aggregate trunk to function in the standby mode. First, the number of ports in the trunk has to exceed the highest allowed number of active ports, and second, the port must be receiving LACPDU packets from the other device. A port functioning in the standby mode does not forward network traffic. However, it continues to send LACPDU packets. If a port that is part of an aggregator does not receive LACPDU packets, it functions as a normal Ethernet port and forwards network packets along with LACPDU packets.

Note

You can adjust a port's priority value.

General Guidelines

Here are LACP guidelines:

- ❑ LACP has to be activated on both the GS950 PS V2 switch and its partner device.
- ❑ The other device has to be 802.3ad-compliant.
- ❑ GS950 PS V2 Switches support up to eight active ports in an aggregate trunk at a time.
- ❑ The ports of an aggregate trunk have to be the same medium type: all twisted pair ports or all fiber optic ports.
- ❑ The ports of a trunk can be consecutive (for example ports 1-5) or nonconsecutive (for example, ports 2, 4, 6, 8).
- ❑ A port can belong to only one aggregator at a time.
- ❑ The ports of an aggregate trunk have to be untagged members of the same VLAN.
- ❑ Copper ports have to be set to Auto-Negotiation or 1000 Mbps, full-duplex mode. LACP trunking is not supported in half-duplex mode.
- ❑ 1000Base-X fiber optic ports have to be set to full-duplex mode.
- ❑ You can create an aggregate trunk of transceivers with 1000Base-X fiber optic ports.
- ❑ Only ports in an aggregator transmit LACPDU packets.
- ❑ Member ports in an aggregator function as part of an aggregate trunk only when receiving LACPDU packets from a remote device. If they do not receive LACPDU packets, they function as regular Ethernet ports, forwarding network traffic, while also continuing to transmit LACPDU packets.
- ❑ The port with the highest priority in an aggregate trunk carries broadcast packets and packets with an unknown destination.
- ❑ Prior to adding an aggregate trunk between an Allied Telesis device and another vendor's device, refer to the vendor's documentation to determine the maximum number of active ports the device can support in a trunk. If the number is less than eight, the maximum number for GS950 PS V2 Switches, you should assign the other vendor's device a higher system LACP priority than your GS950 PS V2 Switch. This can help avoid a conflict between the devices if some ports are placed in the standby mode when the devices form the trunk. For background information, refer to "System Priority and ID Numbers" on page 372.
- ❑ LACPDU packets are transmitted as untagged packets.

Displaying LACP Group Status

To display the LACP Group Status, select **Bridge** -> **Trunk Config.** -> **LACP Group Status** from the main menu. Refer to Figure 178.




Figure 178. LACP Group Status Menu Selection

The LACP Group Status window is shown in Figure 179.

LACP Group Status

System Priority: 32768
System ID: E0:1A:EA:56:3F:54

Group	Aggregator	Active Port List	Standby Port List
1	This group does not exist		
2	This group does not exist		
3	This group does not exist		
4	This group does not exist		
5	This group does not exist		
6	This group does not exist		
7	This group does not exist		
8	This group does not exist		

Figure 179. LACP Group Status Window

The fields and table columns are described in Table 68.

Table 68. LACP Group Status Window

Field or Column	Description
System Priority	Displays the switch's LACP System Priority value, 32768. You cannot change this value. Refer to "System Priority and ID Numbers" on page 372.
System ID	Displays the MAC address of the switch. You cannot change this value.

Table 68. LACP Group Status Window (Continued)

Field or Column	Description
Group	Displays the group ID number of an LACP trunk. This value is not adjustable.
Aggregator	Displays the trunk status: - This group does not exist: The trunk has not be defined. - 1: This group is created.
Active Port List	Displays the ports that are active members of the LACP trunk. The GS950 PS V2 Switch supports up to eight active trunk members per trunk.
Standby Port List	Displays the ports that are standby members of the LACP trunk. The ports transmit and accept LACPDU packets, but do not forward network traffic.

Adding LACP Trunks



Caution

Do not connect the cables of a port trunk until after you have configured the ports on both switches. Connecting the cables prior to configuring the ports can form loops in your network topology, which can result in broadcast storms that can reduce network performance.

Perform the following procedure to add LACP trunks to the switch:

1. Select **Bridge** -> **Trunk Config.** -> **Trunking** from the main menu. Refer to Figure 180.




Figure 180. Trunking Menu Selection

The Trunking window is shown in Figure 181.

Trunking																															
Trunk ID 1:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply	
Trunk ID 2:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 3:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 4:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 5:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 6:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 7:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply
Trunk ID 8:	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Disabled ▾	Apply

Note: Disabling will turn off the function and return all values to default.

Figure 181. Trunking Window

The switch can have up to eight LACP and static port trunks, Each row represents a separate trunk. The numbers and check boxes represent the switch ports. Ports with check marks are members of a trunk.

2. Click the check boxes of the ports for the trunk. Here are the guidelines:
 - A check in a box indicates the port is a member of the trunk. No check means the port is not a member.
 - Ports can be members of only one trunk at a time. Ports that are already members of a trunk have to be removed from their current trunk assignments before you can assign them to a different trunk.
 - An LACP trunk can have any number of ports. Up to eight ports can be active at one time.
3. Select one of the following from the menu in the right column of the Trunk ID row:
 - Active:** Activates the ports for an LACP trunk. The ports send and respond to LACPDU packets by forming an LACP trunk.
 - Passive:** Activates the ports for an LACP trunk. The trunk ports respond to LACPDU packets by forming an LACP trunk, but do not send LACPDU packets. You can use this setting if the other network device does not need to receive LACPDU packets to form LACP trunks.
 - Disabled:** Disables an LACP trunk.

Note

The Manual setting in the menu is for static port trunks. Refer to Chapter 29, "Static Port Trunks" on page 363.

4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.
6. Configure the ports on the other network device for LACP.
7. Connect the Ethernet cables between the trunk ports on the two devices.

Setting LACP Port Priorities

The switch uses LACP port priorities to select the active ports of a trunk when the number of ports exceeds the maximum number of eight active ports. Refer to “Port Priority Value” on page 373. Perform the following procedure to set LACP port priorities:

1. Select the **Bridge -> Trunk Config -> Port Priority** from the main menu. Refer to Figure 182.




Figure 182. Port Priority Menu Selection

The Port Priority window is show in Figure 183.

Port Priority

System Priority 32768
System ID E0:1A:EA:56:3F:54

Port	Priority (0-65535)
1	0
2	0
3	0
4	0

Figure 183. Port Priority Window

2. To change a port's priority, enter a value from 0 to 65535 in the Priority field for the port. The lower the number the higher the priority. The default is 0.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying LACP Trunks



Caution

Before modifying a port trunk, disconnect all cables from its ports. Leaving the cables connected can create loops in your network topology, which can result in broadcast storms that can reduce network performance.

Perform this procedure to add or remove ports from LACP trunks:

1. Disconnect all Ethernet cables from the trunk ports.
2. Select **Bridge > Trunk Config. > Trunking** from the main menu. Refer to Figure 178 on page 375. The Trunking window is shown in Figure 179 on page 375. The eight Trunk ID rows are the static and LACP trunks.
3. To add or remove ports from a trunk, click the port check boxes. Here are the guidelines:
 - A check in a box indicates the port is a member of the trunk. No check means the port is not a member.
 - An LACP trunk can have any number of ports. Up to eight ports can be active at one time.
 - Ports can be members of only one trunk at a time. Ports that are already members of a trunk have to be removed from their current trunk before you can add them to a different trunk.
4. To enable or disable a trunk, change its status with the menu in the right column of the Trunk ID row. The choices are listed here:
 - Active:** Activates the ports for an LACP trunk. The ports send and respond to LACPDU packets by forming an LACP trunk.
 - Passive:** Activates the ports for an LACP trunk. The trunk ports respond to LACPDU packets by forming an LACP trunk, but do not send LACPDU packets. You can use this setting if the other network device does not need to receive LACPDU packets to form LACP trunks.
 - Disabled:** Disables a static port or LACP trunk.

Note

The Manual setting in the menu is for static port trunks. Refer to Chapter 29, “Static Port Trunks” on page 363.

5. Click the **Apply** button.

6. Select **Save Settings to Flash** from the main menu to save your changes.
7. Configure the port trunk on the other switch.
8. Connect the Ethernet cables between the trunk ports on the two devices.

Deleting LACP Trunks



Caution

Before deleting a port trunk, disconnect all cables from its ports. Leaving the cables connected can create loops in your network topology, which can result in broadcast storms that can reduce network performance.

Perform the following procedure:

1. Disconnect all Ethernet cables from the trunk ports.
2. Select **Bridge > Trunk Config. > Trunking** from the main menu. Refer to Figure 180 on page 377. The Trunking window is shown in Figure 181 on page 377.
3. Remove the ports from the trunk by clicking their check boxes to remove the check marks.
4. Set the status of the trunk to **Disabled** with the pull-down menu to the right of the trunk row.
5. Click the **Apply** button.
6. Select **Save Settings to Flash** from the main menu to save your changes.

Section VI

Simple Network Management Protocol

This section contains the following chapters:

- ❑ Chapter 31, “SNMPv1 and v2c” on page 385
- ❑ Chapter 32, “SNMPv3” on page 397

Chapter 31

SNMPv1 and v2c

This chapter described SNMPv1 and SNMPv2c in the following sections:

- ❑ “Enabling or Disabling SNMP Management” on page 386
- ❑ “SNMPv1 and SNMPv2c User and Group Names” on page 387
- ❑ “SNMP Community Strings” on page 390
- ❑ “SNMP Traps” on page 392

Enabling or Disabling SNMP Management

Perform the following procedure to enable or disable SNMP management on the switch:

1. Select **System** -> **User Interface** from the main menu. Refer to Figure 37 on page 87. The User Interface window is shown in Figure 38 on page 87.
2. Choose one of the following from the **SNMP Agent** menu:
 - Enabled** - Enables the SNMP agent, allowing you to manage the switch with SNMP.
 - Disabled** - Disables the SNMP agent, blocking you from managing the switch with SNMP.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

SNMPv1 and SNMPv2c User and Group Names

SNMPv1 and SNMPv2c User Name and Group Name definitions are the basis for adding SNMP communities. Use the following sections to create or delete User and Group Names:

- ❑ “Adding User and Group Names,” next
- ❑ “Modifying User and Group Names” on page 388
- ❑ “Deleting User and Group Names” on page 389

A community string has attributes for controlling who can use the string and what the string allows a network management station to do on the switch.

The GS950 PS V2 Switch does not have any default community strings. You have to define an SNMP User and Group Name on the SNMP User/Group window and then define a community name on the SNMP Community Table window.

Adding User and Group Names

Perform the following procedure to add SNMP Users and Group Names:

1. Select **SNMP > SNMP User/Group** from the main menu. Refer to Figure 178.

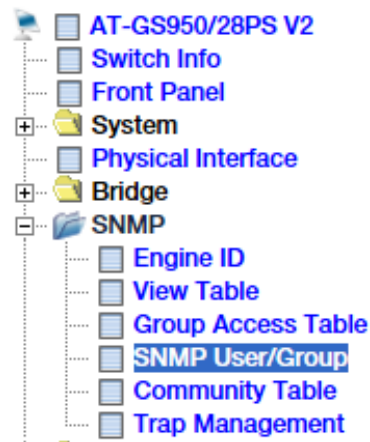


Figure 178. SNMP User/Group Menu Selection

The SNMP User/Group window is shown in Figure 179.

SNMP User/Group

User Name: * (32 characters limit)

Group Name: * (32 characters limit)

SNMP Version: v1 ▼ encrypted

Auth-Protocol: MD5 ▼ Password:

Priv-Protocol: DES ▼ Password:

User Name	Group Name	SNMP Version	Auth-Protocol	Priv-Protocol	Action
ReadOnly	ReadOnly	v1	None	None	<input type="button" value="Delete"/>
ReadOnly	ReadOnly	v2c	None	None	<input type="button" value="Delete"/>
ReadWrite	ReadWrite	v1	None	None	<input type="button" value="Delete"/>
ReadWrite	ReadWrite	v2c	None	None	<input type="button" value="Delete"/>

Figure 179. SNMP User/Group Window

- Configure the settings in Table 68.

Table 68. SNMP User/Group Window for v1/v2c

Parameter	Description
User Name	Enter a user name up to 32 characters.
Group Name	Enter the group name ReadOnly or ReadWrite.
SNMP Version	Select either v1 or v2c .

Note

The Auth-Protocol, Priv-Protocol, and Password fields are for SNMPv3 user and group names.

- Click the **Add** button.
- Select **Save Settings to Flash** from the main menu to save your changes.

Modifying User and Group Names

You cannot modify entries in the SNMP User/Group window. To change an entry, you have to delete it and then enter it again with the changes.

Deleting User and Group Names

Perform this procedure to delete entries from the SNMP User/Group window:

1. Select **SNMP > SNMP User/Group** from the main menu. Refer to Figure 178 on page 387 and Figure 179 on page 388.
2. In the **Action** column of the table, click the **Delete** button of the User Name and Group Name to be deleted.

Note

You cannot delete the default User and Group names.

3. Select **Save Settings to Flash** from the main menu to save your changes.

SNMP Community Strings

A community string has attributes for controlling who can use the string and what the string will allow a network management station to do on the switch. The GS950 PS V2 Switch does not have default community strings. You have to first define an SNMP User and Group Name on the SNMP User/Group page and then define a community name on the SNMP Community Table page.

Here are the SNMP community strings procedures:

- ❑ “Adding SNMP Community Strings,” next
- ❑ “Modifying SNMP Community Strings” on page 391
- ❑ “Deleting SNMP Community Strings” on page 391

Adding SNMP Community Strings

Perform the following procedure to add SNMPv1 or SNMPv2c community strings:

1. Select **SNMP > Community Table** from the main menu. Refer to Figure 180.

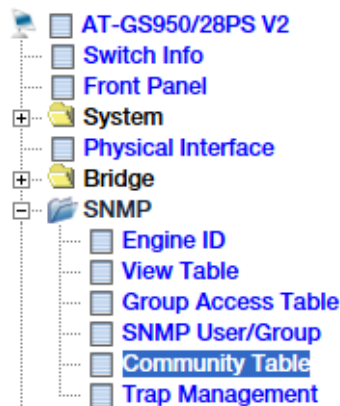


Figure 180. Community Table Menu Selection

The Community Table window is shown in Figure 181 on page 391.

SNMP Community Table

Community Name: * (32 characters limit)

User Name (View Policy): * (32 characters limit)

Community Name	User Name (View Policy)	Action
<< Table is empty >>		

Figure 181. SNMP Community Table Window

- Configure the fields in Table 69.

Table 69. SNMP Community Table Window

Parameter	Description
Community Name	Enter a new community name up to 32 characters.
User Name (View Policy)	Enter a user name up to 32 characters. This name has to match one of the user names in the SNMP User/Group window, explained in "Adding User and Group Names" on page 387.

- Click the **Add** button to add the new community name to the table.
- Select **Save Settings to Flash** from the main menu to save your changes.

Modifying SNMP Community Strings

To modify a Community Table entry, you have to delete it and enter it again with the desired changes.

Deleting SNMP Community Strings

Perform the following procedure to delete SNMP community names.

- Select **SNMP > Community Table** from the main menu. Refer to Figure 180 on page 390.
- Click the **Delete** button in the Action column of the entry you want to remove. You can delete only one community at a time.
- Select **Save Settings to Flash** from the main menu to save your changes.

SNMP Traps

The switch can send SNMP traps to designated host IP addresses of management devices to alert you to network operational events. Here are the procedures:

- ❑ “Adding Trap Host Table Entries,” next
- ❑ “Modifying Trap Host Table Entries” on page 394
- ❑ “Deleting Trap Host Table Entries” on page 395

Adding Trap Host Table Entries

Use the following procedure to add trap Host Table entries:

1. Select **SNMP > Trap Management** from the main menu. Refer to Figure 182.

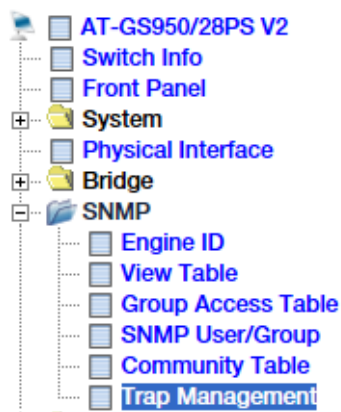


Figure 182. Trap Management Menu Selection

The Trap Management window is shown in Figure 183 on page 393.

Trap Management

Trap Status:

Add Host Table

Host IP Address: . . . IPv4
 IPv6

SNMP Version:

Community Name/User Name: * (32 characters limit)

Host IP Address	SNMP Version	Community Name/User Name	Action
<< Table is empty >>			

Figure 183. Trap Management Window

To enable or disable trap management, perform steps 2 and 3. To add new trap destinations, go to step 4.

2. Select one of the following from Trap Status:
 - Enabled:** This activates trap management. You have to enable trap management to configure its settings.
 - Disabled:** This deactivates trap management. The switch does not transmit traps when trap management is disabled.
3. Click **Apply**.
4. Configure the settings in Table 70 on page 394.

Table 70. Trap Management Window

Parameter	Description
Host IP Address	<p>Enter the IP address of a host node to receive traps from the switch. The address can be either IPv4 or IPv6. Here are the guidelines for entering an IPv4 address:</p> <ul style="list-style-type: none"> - The IPv4 address has to be entered in this format. nnn.nnn.nnn.nnn - Each “nnn” is a decimal number from 0 to 255 - The numbers have to be separated by periods. <p>Here are the guidelines for entering an IPv6 address:</p> <ul style="list-style-type: none"> - The IPv6 address has to be entered in this format. nnn:nnn:nnn:nnn:nnn:nnn:nnn:nnn - N is a hexadecimal digit from 0 to F. - The eight groups are separated by colons. - You can omit groups where all four digits are “0”. - You can also omit leading “0”s in groups. As an example, the following two addresses are equivalent: <ul style="list-style-type: none"> - 12c4:421e:09a8:0000:0000:0000:00a4:1c50 - 12c4:421e:09a8:a4:1c50
SNMP Version	Select v1 or v2c .
Community Name/ User Name	Enter an existing community name from the SNMP Community Table. Refer to “SNMP Community Strings” on page 390.

5. Click the **Add** button.
6. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying Trap Host Table Entries

To modify an trap host entry, you have to delete it and enter it again with the necessary changes.

Deleting Trap Host Table Entries

Perform the following procedure to delete trap host entries:

1. Select **SNMP > Trap Management** from the main menu. Refer to Figure 182 on page 392 and Figure 183 on page 393.
2. Click the **Delete** button in the Action column of the entry to be deleted. No confirmation message is displayed.
3. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 32

SNMPv3

This chapter describes SNMPv3 in the following sections:

- ❑ “SNMPv3 Overview” on page 398
- ❑ “SNMPv3 User and Group Names” on page 402
- ❑ “SNMPv3 View Names” on page 405
- ❑ “SNMPv3 View Table” on page 409
- ❑ “SNMPv3 Traps” on page 412
- ❑ “SNMP Engine ID” on page 413

SNMPv3 Overview

The SNMPv3 protocol builds on the existing SNMPv1 and SNMPv2c protocol implementation which is described in Chapter 31 on page 385. In SNMPv3, User-based Security Model (USM) authentication is implemented along with encryption, allowing you to configure a secure SNMP environment.

The SNMPv3 protocol uses different terminology than the SNMPv1 and SNMPv2c protocols. In the SNMPv1 and SNMPv2c protocols, the terms agent and manager are used. An agent is the software within an SNMP user, while a manager is an SNMP host. In the SNMPv3 protocol, agents and managers are called entities. In any SNMPv3 communication, there is an authoritative entity and a non-authoritative entity. The authoritative entity checks the authenticity of the non-authoritative entity. The non-authoritative entity checks the authenticity of the authoritative entity.

With the SNMPv3 protocol, you create users, determine the protocol used for message authentication and determine if data transmitted between two SNMP entities are encrypted. In addition, you can restrict user privileges by defining which portions of the Management Information Bases (MIBs) can be viewed by specific users. In this way, you restrict which MIBs a user can display and modify. In addition, you can restrict the types of messages, or traps, the user can send. (A trap is a type of SNMP message.) After you have created a user, you define SNMPv3 message notification. This consists of determining where messages are sent and which types of messages can be sent. This configuration is similar to the SNMPv1 and SNMPv2c configurations because you configure IP addresses of trap receivers, or hosts.

This section describes the features of the SNMPv3 protocol. The following subsections are included:

- ❑ "SNMPv3 Authentication Protocols"
- ❑ "SNMPv3 Privacy Protocol" on page 399
- ❑ "SNMPv3 MIB Views" on page 399
- ❑ "SNMPv3 Configuration Process" on page 400

SNMPv3 Authentication Protocols

The SNMPv3 protocol supports two authentication protocols— HMAC-MD5-96 (MD5) and HMAC-SHA-96 (SHA). Both MD5 and SHA use an algorithm to generate a message digest. Each authentication protocol authenticates a user by checking the message digest. In addition, both protocols use keys to perform authentication. The keys for both protocols are generated locally using the Engine ID and the user password. You can only modify a key by modifying the user password.

In addition, you have the option of assigning no user authentication. In this case, no authentication is performed for this user. You may want to make this configuration for someone with super-user capabilities.

SNMPv3 Privacy Protocol

After configuring an authentication protocol, you have the option of assigning a privacy protocol. In SNMPv3 protocol terminology, privacy is equivalent to encryption. Currently, the DES protocol is the only encryption protocol supported. The DES privacy protocol requires the authentication protocol to be configured as either MD5 or SHA.

If you assign a DES privacy protocol to a user, then you are also required to assign a privacy password. If you choose to not assign a privacy value, then SNMPv3 messages are sent in plain text format.

SNMPv3 MIB Views

The SNMPv3 protocol allows you to configure MIB views for users and groups. The MIB tree is defined by RFC 1155 (Structure of Management Information). See Figure 184.

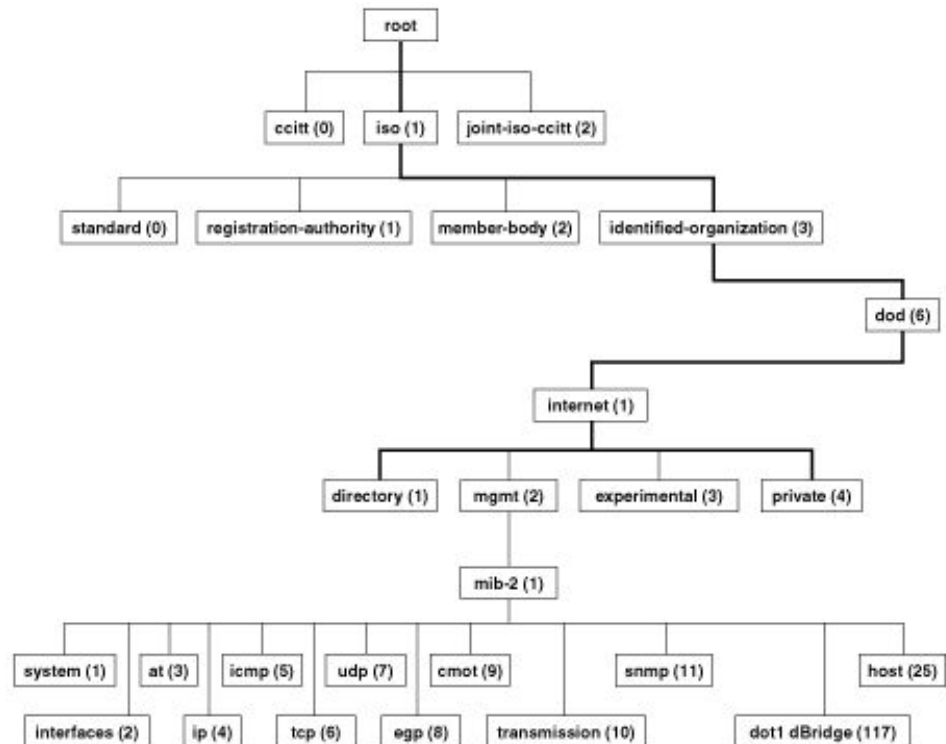


Figure 184. MIB Tree

The GS950 PS V2 Switches support the MIB tree, starting with the Internet MIBs, as defined by 1.3.6.1. There are two ways to specify an MIB view. You can enter the OID number of the MIB view or its equivalent text name. For example, to specify MIBs in the Internet view, you can enter the OID format “1.3.6.1” or the text name, “internet.”

In addition, you can define an MIB view that the user can access or an MIB view that the user cannot access. When you want to permit a user to access an MIB view, you include a particular view. When you want to deny a user access to an MIB view, you exclude a particular view.

After you specify an MIB subtree view you have the option of further restricting a view by defining a subtree mask. The relationship between an MIB subtree view and a subtree mask is analogous to the relationship between an IP address and a subnet mask. The switch uses the subnet mask to determine which portion of an IP address represents the network address and which portion represents the node address. In a similar way, the subtree mask further refines the subtree view and enables you to restrict an MIB view to a specific row of the OID MIB table. You need a thorough understanding of the OID MIB table to define a subtree mask.

SNMPv3 Configuration Process

The SNMPv3 parameters are contained in the following tables:

- SNMPv3 User/Group table
- SNMPv3 Access table
- SNMPv3 View table
- SNMPv3 Community table
- Trap Management

The SNMPv3 configuration information must be entered in a specific sequence:

Note

The SNMP Interface has to be activated first. Refer to “Enabling or Disabling SNMP Management” on page 386.

1. Create a user name and associated group name in the SNMPv3 User/Group table.
2. Define view names in the Access table for each group name.
3. Define the MIB view in the SNMPv3 View table for each view name.
4. Enter information in the Community table based on a pre-defined user name.

Note

The community strings do not have a default value defined and are initially blank.

5. Define the traps on the Trap Management page based on the community or user name.

See Figure 185 for an illustration of how the user configuration tables are linked.

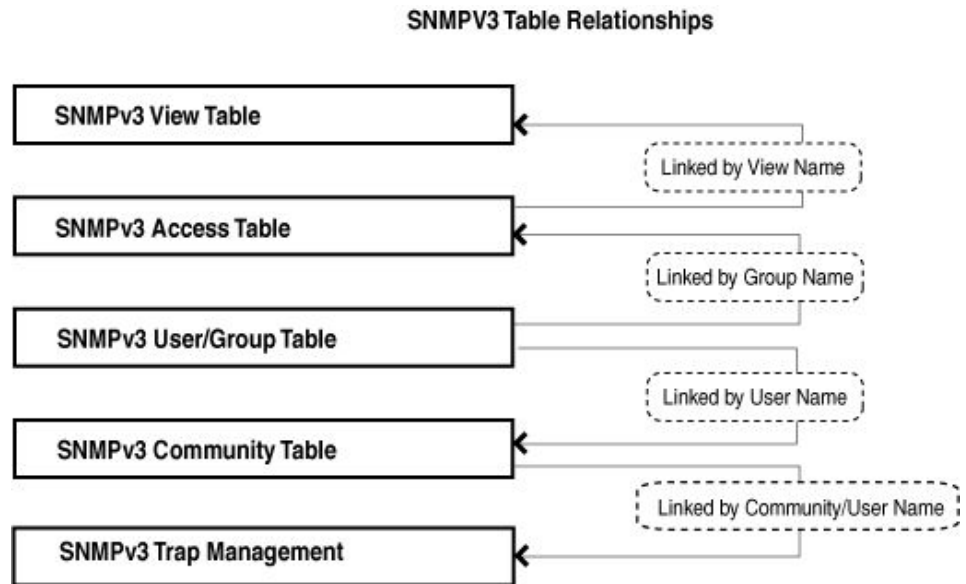


Figure 185. SNMPv3 Table Relationships

SNMPv3 User and Group Names

SNMPv3 user and group names are the basis for all SNMPv3 tables. Here are the procedures:

- "Adding SNMPv3 User and Group Names"
- "Modifying SNMPv3 User and Group Names" on page 403
- "Deleting SNMPv3 User and Group Names" on page 404

Adding SNMPv3 User and Group Names

Perform this procedure to add SNMPv3 user and group names:

1. Select **SNMP > SNMP User/Group** from the main menu. Refer to Figure 178 on page 387. The SNMP User/Group window is shown in Figure 179 on page 388.
2. Configure the settings in Table 71.

Table 71. SNMP User/Group Window for v3

Parameter	Description
User Name	Enter a user name up to 32 characters.
Group Name	Enter a group name up to 32 characters.
SNMP Version	Select v3 .
encrypted	To add encryption to the user name, click the encrypted check box to add a check mark.
Auth-Protocol	Select one of the following: <ul style="list-style-type: none"> - MD5: The switch uses MD5 authentication to authenticate the SNMPv3 user. This is the default value. - SHA: The switch uses SHA authentication to authenticate the SNMPv3 user.
Password	Enter a password for Auth-Protocol.
Priv-Protocol	Select one of the following: <ul style="list-style-type: none"> - DES: The switch uses DES encryption to encrypt the SNMP packets. This is the default value. - none: The switch does not encrypt the SNMP packets.

Table 71. SNMP User/Group Window for v3 (Continued)

Parameter	Description
Password	Enter a password for Priv-Protocol. Auth-Protocol has to have a password if DES for Priv-Protocol has a password.

- Click the **Add** button.

The new user name and group name are displayed on the SNMP User/Group page. See Figure 186 for an example.

SNMP User/Group

User Name: * (32 characters limit)

Group Name: * (32 characters limit)

SNMP Version: v1 encrypted

Auth-Protocol: MD5 Password:

Priv-Protocol: DES Password:

User Name	Group Name	SNMP Version	Auth-Protocol	Priv-Protocol	Action
Jarad	Managers	v3	MD5	DES	Delete
Jordan	Managers	v3	MD5	DES	Delete
ReadOnly	ReadOnly	v1	None	None	Delete
ReadOnly	ReadOnly	v2c	None	None	Delete
ReadWrite	ReadWrite	v1	None	None	Delete
ReadWrite	ReadWrite	v2c	None	None	Delete

Figure 186. SNMP User Group, SNMPv3 Window Example

- Select **Save Settings to Flash** from the main menu to save your changes.

Modifying SNMPv3 User and Group Names

You cannot modify user or group name entries. To change an entry, you have to delete it and enter it again with the new changes. Refer to “Deleting SNMPv3 User and Group Names” on page 404 and “Adding SNMPv3 User and Group Names” on page 402.

Deleting SNMPv3 User and Group Names

Perform this procedure to delete entries from the SNMP User/Group window.

1. Select **SNMP > SNMP User/Group** from the main menu. Refer to Figure 178 on page 387. The SNMP User/Group window is shown in Figure 179 on page 388.
2. Click the **Delete** button in the **Action** column of the user name to be removed. You cannot delete the default SNMPv1 and v2c entries.
3. Select **Save Settings to Flash** from the main menu to save your changes.

SNMPv3 View Names

The SNMPv3 view names are defined in the SNMP Group Access table and are based on the user and group names. You can add and delete view names with the following procedures:

- "Adding SNMPv3 View Names"
- "Modifying SNMPv3 View Names" on page 408
- "Deleting SNMPv3 View Names" on page 408

Adding SNMPv3 View Names

To add an SNMPv3 view name, you first have to define its group name with the SNMP User/Group window. Refer to "Adding SNMPv3 User and Group Names" on page 402.

Perform this procedure to add SNMPv3 view names.

1. Select **SNMP > Group Access Table** from the main menu. Refer to Figure 187.

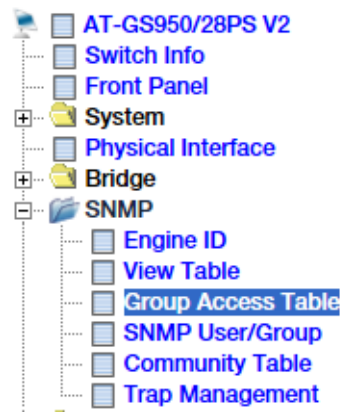


Figure 187. Group Access Table Menu Selection

The SNMP Group Access Table window is shown in Figure 188 on page 406.

SNMP Group Access Table

Group Name: * (32 characters limit)

Read View Name: (32 characters limit)

Write View Name: (32 characters limit)

Notify View Name: (32 characters limit)

Security Model:

Security Level:

Group Name	Read View	Write View	Notify View	Security Model	Security Level	Action
ReadOnly	ReadWrite	---	ReadWrite	v1	NoAuthNoPriv	Delete
ReadOnly	ReadWrite	---	ReadWrite	v2c	NoAuthNoPriv	Delete
ReadWrite	ReadWrite	ReadWrite	ReadWrite	v1	NoAuthNoPriv	Delete
ReadWrite	ReadWrite	ReadWrite	ReadWrite	v2c	NoAuthNoPriv	Delete

Figure 188. SNMP Group Access Table Window

- Configure the settings in Table 72.

Table 72. SNMP Group Access Table Window for v3

Parameter	Description
Group Name	Enter the name of an existing group in the Group Name field. To add group names, refer to “Adding SNMPv3 User and Group Names” on page 402.
Read View Name	Enter a read view name of up to 32 characters. The name is optional.
Write View Name	Enter a write view name of up to 32 characters. The name is optional.
Notify View Name	Enter a notify view name of up to 32 characters. The name is optional.
Security Model	Select v3 .

Table 72. SNMP Group Access Table Window for v3 (Continued)

Parameter	Description
Security Level	<p>Choose a Security Level from the menu. The selections are listed here:</p> <ul style="list-style-type: none"> - NoAuthNoPriv: This selection is appropriate when no Auth-Protocol or Priv-Protocol (no encryption) are selected on the SNMP User/Group page. - AuthNoPriv: Choose this selection when encryption has been enabled, but only the Auth-Protocol has a password assigned, and the Priv-Protocol has been selected as none on the SNMP User/Group page. - AuthPriv: When both the Auth-Protocol and Priv-Protocol have been enabled, choose this selection.

- Click the **Add** button. See Figure 189 for an example.

SNMP Group Access Table

Group Name: * (32 characters limit)

Read View Name: (32 characters limit)

Write View Name: (32 characters limit)

Notify View Name: (32 characters limit)

Security Model: v1 ▾

Security Level: NoAuthNoPriv ▾

Group Name	Read View	Write View	Notify View	Security Model	Security Level	Action
Managers	---	---	---	v3	NoAuthNoPriv	Delete
ReadOnly	ReadWrite	---	ReadWrite	v1	NoAuthNoPriv	Delete
ReadOnly	ReadWrite	---	ReadWrite	v2c	NoAuthNoPriv	Delete
ReadWrite	ReadWrite	ReadWrite	ReadWrite	v1	NoAuthNoPriv	Delete
ReadWrite	ReadWrite	ReadWrite	ReadWrite	v2c	NoAuthNoPriv	Delete

Figure 189. SNMP Group Access Table Example for SNMPv3

- Select **Save Settings to Flash** from the main menu to save your changes.

Modifying SNMPv3 View Names

You cannot modify entries in the SNMP Group Access window. To change an entry, you have to delete it and enter it again with the changes. Refer to “Deleting SNMPv3 View Names” on page 408 and “Adding SNMPv3 View Names” on page 405.

Deleting SNMPv3 View Names

Perform this procedure to delete entries from the SNMP Group Access Table window:

1. Select **SNMP > Group Access Table** from the main menu. Refer to Figure 187 on page 405. The SNMP Group Access Table window is shown in Figure 188 on page 406.
2. Click the **Delete** button in the **Action** column of the group name to be removed.

Note

The views corresponding to the ReadOnly and ReadWrite Group Names are default values and cannot be deleted.

3. Select **Save Settings to Flash** from the main menu to save your changes.

SNMPv3 View Table

The SNMPv3 View table specifies the MIB object access criteria for each view name. If the view name is not specified on this page, then it has access to all MIB objects. You can specify specific areas of the MIB that can be accessed or denied based on the entries in this table. You can add and delete entries in the View table by following the procedures in these sections:

- "Adding SNMPv3 View Table Entries"
- "Modifying SNMPv3 View Table Entries" on page 410
- "Deleting SNMPv3 View Table Entries" on page 411

Adding SNMPv3 View Table Entries

Perform this procedure to add entries to the SNMPv3 View table.

1. Select **SNMP > View Table** from the main menu. Refer to Figure 190.




Figure 190. View Table Menu Selection

The SNMP View Table window is shown in Figure 191.

SNMP View Table

View Name: * (32 characters limit)

Subtree OID: *

OID Mask: *

View Type: included ▼

View Name	Subtree OID	OID Mask	View Type	Action
ReadWrite	1	1	Included	Delete

Figure 191. SNMP View Table Window

2. Configure the settings in Table 73.

Table 73. SNMP View Table Window for v3

Parameter	Description
View Name	Enter the name of an existing view. Refer to "Adding SNMPv3 View Names" on page 405.
Subtree OID	Enter the subtree OID.
OID Mask	Enter "1" for the mask.
View Type	Select one of the following: - Included: This selection allows the view to include the specified MIB object. - Excluded: This selection blocks the view of the specified MIB object.

3. Click the **Add** button.

The updated view is displayed in the View table. See Figure 192.

SNMP View Table

View Name: * (32 characters limit)

Subtree OID: *

OID Mask: *

View Type: included ▼

Add
Reset

View Name	Subtree OID	OID Mask	View Type	Action
Read	1	1	Included	Delete
ReadWrite	1	1	Included	Delete

Figure 192. SNMP View Table Window Example

4. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying SNMPv3 View Table Entries

You cannot modify an entry in the View Table page. To change an entry, you have to delete it and enter it again with the changes. Refer to "Deleting SNMPv3 View Table Entries" and "Adding SNMPv3 View Table Entries" on page 409.

Deleting SNMPv3 View Table Entries

Perform this procedure to delete entries from the SNMPv3 View table.

1. Select **SNMP > View Table** from the main menu. Refer to Figure 190 on page 409 and Figure 191 on page 409.
2. Click the **Delete** button in the Action column of the entry you want to remove.

Note

The views corresponding to the ReadOnly and ReadWrite Group Names are default values and cannot be removed.

3. Select **Save Settings to Flash** from the main menu to save your changes.

SNMPv3 Traps

The procedures for adding, modifying, or deleting SNMPv3 traps are the same as for SNMPv1/v2 traps. Refer to “SNMP Traps” on page 392.

SNMP Engine ID

An SNMP agent has an engine ID that uniquely identifies the agent in a device and the MIB objects in a domain. The engine ID of the switch follows the RFC 3411 standard and consists of the enterprise ID and the MAC address of the switch. You can modify or reset the SNMP engine ID by following these procedures:

- ❑ “Modifying SNMP Engine ID” next
- ❑ “Resetting SNMP Engine ID” on page 413

Modifying SNMP Engine ID

Perform this procedure to modify the engine ID.

1. Select **SNMP > Engine ID** from the main menu. Refer to Figure 193.




Figure 193. SNMP Engine ID Settings Menu Selection

The SNMP Engine ID Settings window is shown in Figure 194.




Figure 194. SNMP Engine ID Settings Window

2. Enter the new engine ID in the Engine ID field.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Resetting SNMP Engine ID

Perform this procedure to reset the engine ID to the previous value or default setting.

1. Select **SNMP > Engine ID** from the main menu. Refer to Figure 193 on page 413. The SNMP Engine ID Settings page is shown in Figure 194 on page 413.
2. Do one of the following:

- To reset the engine ID to the previous setting, click the **Reset** button.
 - To reset the engine ID to the default setting, click the **Reset to Default** button.
3. Select **Save Settings to Flash** from the main menu to save your changes.

Section VII

Network Security

This section contains the following chapters:

- ❑ Chapter 33, “Port Authentication” on page 417
- ❑ Chapter 34, “Local Dial-in User Accounts” on page 439
- ❑ Chapter 35, “RADIUS Client” on page 445
- ❑ Chapter 36, “TACACS+ Client” on page 451
- ❑ Chapter 37, “DHCP Snooping” on page 457
- ❑ Chapter 38, “Destination MAC Address Filters” on page 475
- ❑ Chapter 39, “Port Mirroring” on page 479
- ❑ Chapter 40, “Denial of Service” on page 485

Chapter 33

Port Authentication

This chapter contains background information on the port authentication feature of the switch. The chapter contains the following sections:

- ❑ “Overview” on page 418
- ❑ “Authentication Methods” on page 420
- ❑ “Authenticator Port Operational Settings” on page 421
- ❑ “Authenticator Port Operating Modes” on page 422
- ❑ “Supplicant and VLAN Associations” on page 426
- ❑ “Guest VLAN” on page 428
- ❑ “RADIUS Accounting” on page 429
- ❑ “General Steps” on page 430
- ❑ “Guidelines” on page 431
- ❑ “Configuring Port Access Control” on page 433

Overview

Port authentication is a network security feature. Network users have to log with log-on credentials before the switch forwards their traffic. Depending on the authentication method, network users may be required to manually provide user names and passwords when they log on or their workstations may automatically transmit their MAC addresses as their log-on user names and passwords.

Here are several feature terms:

- ❑ Supplicant - A supplicant is an end user or node that wants to access the network through a switch port. Supplicants are also referred to as clients.
- ❑ Authenticator - The authenticator is a port on the switch that prohibits network access until a supplicant has logged on and been validated by an authentication device.
- ❑ Authentication device - The device that authenticates the user names and passwords from the supplicants on the switch ports. This can be RADIUS or TACACS+ servers on your network, or the switch's local dial-in user accounts.

Authentication Devices

Port-based network access control is a security feature for network clients. It requires that clients log on by providing user names and passwords when they initially send traffic through the switch ports. The feature is used to prevent unauthorized individuals from connecting computers to switch ports or accessing unattended workstations. The switch permits only those users who have been assigned log on credentials to forward traffic through its ports.

The device that performs the verification of user credentials is referred to as the authentication device. The GS950 PS V2 Switch supports the following three authentication devices:

- ❑ RADIUS server with Extensible Authentication Protocol (EAP) extensions: The switch has a RADIUS client so that it can communicate with RADIUS servers on your network. The client on the switch acts as an intermediary between the network users and the RADIUS servers. When network users provide their credentials to log on the network, the client on the switch forwards the information to the RADIUS server, which validates the credentials and notifies the switch as to whether the credentials are valid. Refer to Chapter 35, “RADIUS Client” on page 445.
- ❑ TACACS+ server: The switch also has a TACACS+ client for validating log on credentials with TACACS+ servers. Refer to Chapter 36, “TACACS+ Client” on page 451.
- ❑ Local dial-in user accounts: If your network does not have RADIUS or TACACS+ servers, you can instead use the switch’s own internal authentication system. When network users log on, the switch checks their log on credentials itself, using its internal database. The switch can store up to 64 local dial-in user accounts. Refer to Chapter 34, “Local Dial-in User Accounts” on page 439.

Authentication Methods

The switch supports two authentication methods:

- ❑ 802.1x port-based network access control
- ❑ MAC address-based authentication

802.1x Port-based Network Access Control

Supplicants of this type of port authentication use usernames and passwords as their logon credentials. They have to provide their unique credentials when they initially begin to forward traffic through the ports on the switch. Supplicants may provide their usernames and passwords manually when prompted by their workstations or their network devices can provide the information automatically. The switch forwards the user credentials for verification to an authentication device, which can be a RADIUS or TACACS+ server on your network, or internally to its local dial-in user accounts.

Supplicants that manually enter their logon credentials are not tied to any specific computer or node. They can log on from any system and still be verified as valid users of the switch and network.

The supplicants must have 802.1x client software to support this port authentication method.

MAC Address- based Authentication

The log-on credentials for supplicants of this type of port authentication consist of the MAC addresses of the network nodes. The MAC addresses of the devices are used as the usernames and passwords of the supplicants. Supplicants are not prompted for this information. Rather, the switch transmits the initial frames with the source MAC address of the node to the authentication device for authentication.

The advantage to this approach is that supplicants do not need 802.1x client software. The disadvantage is that because clients are not prompted for usernames and passwords, unauthorized individuals can access your network through unattended network nodes or by counterfeiting valid network MAC addresses.

Authenticator Port Operational Settings

An authenticator port on the switch can have one of three possible operational settings:

- ❑ Auto - Activates port authentication on a port. Supplicants must provide logon credentials for verification before a port begins to forward their network traffic. This is the default setting for an authenticator port.
- ❑ Force-authorized - Disables port authentication and automatically places the port in the authorized state without any authentication exchange required. The port transmits and receives normal traffic without authenticating the client.

Note

A supplicant connected to an authenticator port set to force-authorized must have 802.1x client software if the port's authenticator mode is 802.1x. Though the force-authorized setting prevents an authentication exchange, the supplicant must still have the client software to forward traffic through the port.

- ❑ Force-unauthorized - Causes the port to remain in the unauthorized state, ignoring all attempts by the supplicant to authenticate. This setting is analogous to disabling a port.

Authenticator Port Operating Modes

Authenticator ports support three modes:

- ❑ Single host mode
- ❑ Single host mode with Piggy-backing
- ❑ Multiple Host mode

Single Host Mode

An authenticator port set to the single host mode permits only one supplicant to log on and forwards only the traffic of that supplicant. After one supplicant has logged on, the port discards packets from any other supplicant that might try to log on.

In Figure 195, port 10 is an authenticator port set to the single host mode. It permits only one supplicant to log on and forwards the traffic from only that supplicant.

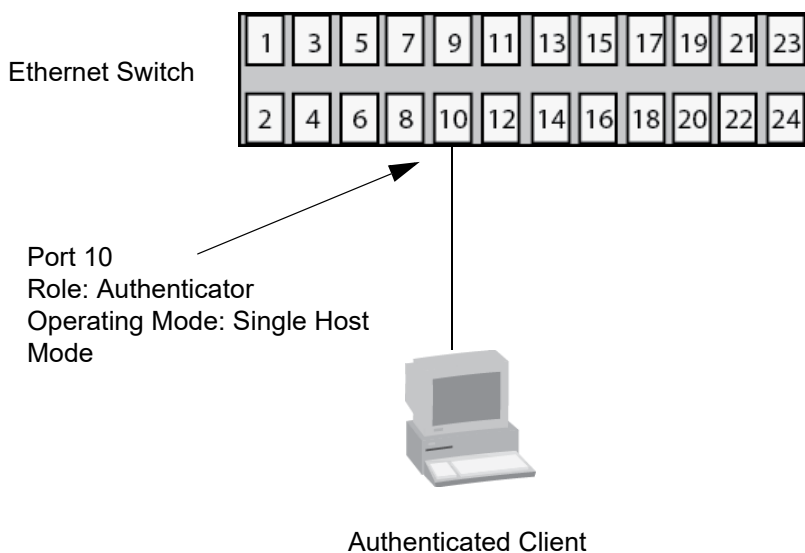


Figure 195. Single Host Mode

Single Host Mode with Piggy Backing

This mode permits multiple clients on an authenticator port, but only one of the clients is authenticated. An authenticator mode forwards packets from all the clients after one client has successfully logged on. This mode is typically used in situations where you want to add authentication to a switch port that is supporting multiple clients, but do not want to create individual accounts for all the clients on the authentication device.

This is referred to as “piggy-backing.” After one client has successfully logged, the port permits the other clients to piggy-back onto the initial client’s log on, so that they can forward packets through the port without

being authentication.

Note, however, that should the client who performed the initial log on fail to periodically reauthenticate or log out, the authenticator port reverts to the unauthenticated state. It bars all further traffic to and from all of the clients until the initial client or another client logs on.

Figure 196 is an example of this mode. Port 10 is connected to an Ethernet hub or non-authentication compliant switch, which in turn is connected to several supplicants. The switch blocks all client traffic until one client logs on. Afterwards, it forwards traffic from all the clients.

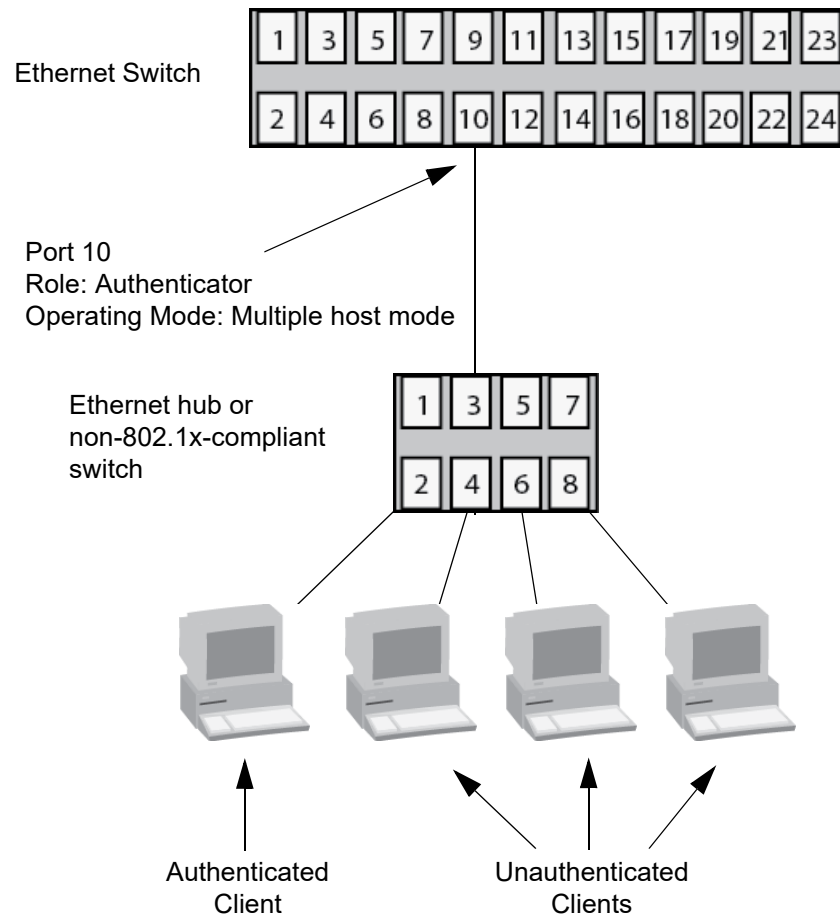


Figure 196. Multiple Host Operating Mode

If the port is set to the 802.1x authentication method, one client has to have 802.1x client firmware and provide a username and password during authentication. (The other clients do not need 802.1x client firmware to forward traffic through the port after one client is authenticated.)

If the port is using MAC address-based authentication, 802.1 client firmware is not required. The first client to forward traffic through the port is

used for authentication. When that client is authenticated, all supplicants have access to the port.

As mentioned earlier, should the client who performed the initial log on fail to reauthenticate when necessary or log out, the port reverts to the unauthenticated state, blocking all traffic to and from all clients. Another client has to be authenticated in order for the remaining clients to continue to forward traffic through the port.

Multiple Host Mode

This mode requires the authentication of all the clients on an authenticator port. This mode is appropriate in situations when you want all the clients to be authenticated on authenticator ports that are supporting more than one client.

If you are using 802.1x authentication, you have to provide each client with a separate username and password combination and the clients have to provide their credentials to forward traffic through a switch port.

An example of this authenticator operating mode is illustrated in Figure 197 on page 425. The clients are connected to a hub or non-authentication switch which is connected to an authenticator port on the switch. If the authenticator port is set to 802.1x authentication, the clients have to provide their username and password credentials before they can forward traffic through the switch.

If the authentication method is MAC address-based, the authenticator port uses the MAC addresses of the clients as the username and password combinations. The port accepts and forwards traffic only from those clients whose MAC addresses have been entered on the authentication devices and denies access to all other users.

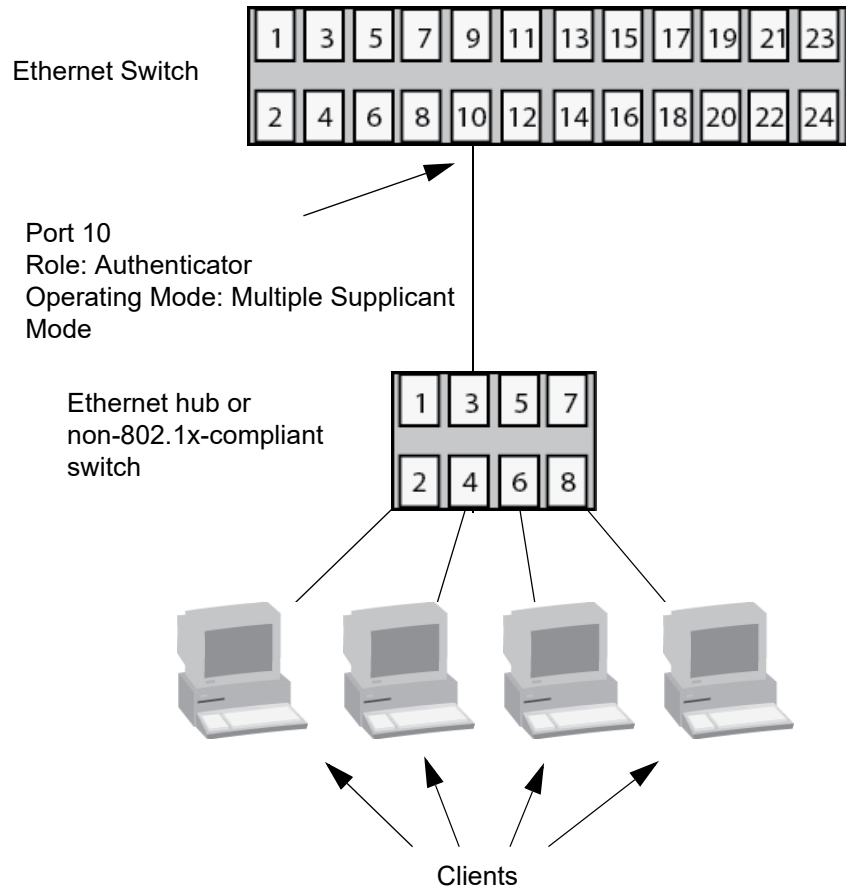


Figure 197. Multiple Supplicant Mode

Supplicant and VLAN Associations

One of the challenges to managing a network is accommodating end users who roam. These are individuals whose work requires that they access the network resources from different points at different times. The difficulty arises in providing them with access to the same network resources and, conversely, restricting them from unauthorized areas, regardless of the workstation from where they access the network. A closely related issue is where a workstation is employed at various times by different individuals with unique requirements in terms of network resources and security levels.

Providing network users with access to their network resources while also maintaining network security is often achieved with VLANs. As explained in Chapter 15, “Port-Based Virtual LANs” on page 183 and Chapter 16, “802.1Q Tagged Virtual LANs” on page 205, VLANs are n independent traffic domains where the traffic generated by the nodes within a VLAN are restricted to nodes of the same VLAN, unless there is a router or Layer 3 device. Typically, network users are assigned to different VLANs depending on their resource requirements and security levels.

The problem with a port-based VLAN is that VLAN membership is determined by the port on the switch to which the device is connected. If a different device that needs to belong to a different VLAN is connected to the port, the port must be manually moved to the new VLAN using the management software.

With port authentication, you can link username and password credentials or MAC addresses to specific VLANs so that the switch automatically moves ports to the appropriate VLANs when clients log on. This frees you from having to reconfigure VLANs as end users access the network from different points or where the same workstation is used by different individuals at different times.

To use this feature, you have to enter a VLAN identifier, along with other information, when you add supplicant accounts on the authentication device. The server passes the identifier to the switch when a user logs on with a valid username and password credential or MAC address, depending on the authentication method.

How the switch responds when it receives VLAN information during the authentication process can differ depending on the operating mode of an authenticator port.

Single Host Mode

Here are the operating characteristics for the switch when an authenticator port is set to the single host mode:

- ❑ If the switch receives a valid VLAN ID or VLAN name from the RADIUS server, it moves the authenticator port to the designated

guest VLAN and changes the port to the authorized state. Only the authenticated supplicant is allowed to use the port. All other supplicants are denied entry.

- ❑ If the switch receives an invalid VLAN ID or VLAN name from the authentication device (e.g., the VID of a nonexistent VLAN), it leaves the port in the unauthorized state to deny access to the port.

Multiple Host Mode

Here are the operating characteristics for the switch when an authenticator port is set to the multiple host mode:

- ❑ If the switch receives a valid VLAN ID or VLAN name from the authentication device, it moves the authenticator port to the designated VLAN and changes the port to the authorized state. All clients are allowed access to the port and the same VLAN after the initial authentication.
- ❑ If the switch receives an invalid VLAN ID or VLAN name from the authentication device (e.g., the VID of a nonexistent VLAN), it leaves the port in the unauthorized state to deny access to the port.

Multiple Supplicant Mode

The initial authentication on an authenticator port running in the multiple supplicant mode is handled in the same fashion as with the Single operating mode. If the switch receives a valid VLAN ID or name from the authentication device, it moves the authenticator port to the designated VLAN and changes the port to the authorized state. All subsequent supplicants who log on to the same port are also authenticated. However, the port remains in the VLAN specified in the initial authentication.

Supplicant VLAN Attributes on RADIUS Servers

If you are using RADIUS servers as the authentication devices, you have to include the following information as part of supplicant accounts when associating supplicants to VLANs.

- ❑ Tunnel-Type: This is the protocol to be used by the tunnel specified by Tunnel-Private-Group-Id. The only supported value is VLAN (13).
- ❑ Tunnel-Medium-Type: This is the transport medium to be used for the tunnel specified by Tunnel-Private-Group-Id. The only supported value is 802 (6).
- ❑ Tunnel-Private-Group-ID This is the ID of the tunnel the authenticated user should use. This must be the name of VID of the VLAN of the switch.

Guest VLAN

An authenticator port in the unauthorized state typically accepts and transmits only 802.1x packets while waiting to authenticate a supplicant. However, you can configure authenticator ports to be members of a Guest VLAN when no supplicants are logged on. Client using the ports are not required to log on and have full access to the resources of the Guest VLAN.

If the switch receives 802.1x packets on the port, signaling that a supplicant is logging on, it moves the port to its predefined VLAN and places it in the unauthorized state. The port remains in the unauthorized state until the log on process between the supplicant and the authentication device is completed. When the supplicant logs off, the port automatically returns to the Guest VLAN.

Note

The Guest VLAN feature is only supported on authenticator ports in the Single operating mode.

RADIUS Accounting

The switch supports RADIUS accounting for switch ports in the authenticator role. This feature sends information to the RADIUS server about the status of the supplicants so that you can monitor network activity and use.

The switch sends accounting information to the RADIUS server when the following events occur:

- Supplicants log on
- Supplicants logs off
- Authenticator ports change states during active supplicant sessions (for example, a port is reset or is changed from the Authenticator role to None role while a supplicant is logged on)

The event information sent to the RADIUS server includes:

- The port number where an event occurred.
- The date and time when an event occurred.
- The number of packets transmitted and received by a switch port during a supplicant's session. (This information is sent only when a client logs off.)

You can also configure the accounting feature to send interim updates so that you can monitor which clients are still active.

Here are the guidelines to using the accounting feature:

- The management software supports the Network level of accounting, but not the System or Exec.
- This feature is only available on Authenticator ports.
- You have to configure 802.1x Port-based Network Access Control as explained in this chapter and designate the Authenticator ports.
- You have to configure the RADIUS client.

General Steps

Here are the general steps to implementing port-based authentication on the switch:

1. To use RADIUS or TACACS+ servers as the authentication devices, install server software on one or more of your network devices.
2. Add supplicant accounts to the authentication device:
 - ❑ For 802.1x authentication, the accounts should have user names and passwords. The maximum length for user names and passwords is 23 alphanumeric characters. Spaces and special characters are not supported.
 - ❑ For MAC address-based authentication, the accounts use the MAC addresses of the supplicant devices as the user names and passwords.
3. Verify that clients connected to authenticator ports set to 802.1x authentication have 802.1x client software. (MAC address authentication does not require 802.1x client software.)
4. If the authentication devices are RADIUS or TACACS+ servers, configure the RADIUS or TACACS+ client on the switch. This involves entering the IP addresses and encryption keys of the servers on your network. Refer to Chapter 35, “RADIUS Client” on page 445 or Chapter 36, “TACACS+ Client” on page 451.
5. Configure the port access control settings on the switch. Refer to “Configuring Port Access Control” on page 433.

Guidelines

Here are the general feature guidelines:

- ❑ Ports configured for authentication do not support dynamic MAC address learning.
- ❑ A port connected to a RADIUS or TACACS+ authentication server cannot be set to the authenticator role because an authentication server cannot authenticate itself.
- ❑ The authentication method can be 802.1x or MAC address authentication.
- ❑ Supplicants connected to authenticator ports set to 802.1x authentication need to have 802.1x client software.
- ❑ Supplicants do not need 802.1x client software for MAC address authentication.
- ❑ The log-on credentials for 802.1x supplicants are not tied to the MAC addresses of an end node. This allows end users to use the same log-on credentials when working from different workstations.
- ❑ The MAC addresses of authenticated clients are added to the MAC address table as authenticated addresses. They remain in the table until clients log off the network or fail to reauthenticate, at which point they are removed. The addresses are not timed out, even if the nodes are inactive.

Note

End users of port authentication should be instructed to always log off at the conclusion of every work session. This can prevent unauthorized individuals from accessing the network through unattended network workstations.

- ❑ Authenticator and supplicant ports need to be untagged ports. They cannot be tagged ports.
- ❑ Authenticator ports cannot use destination MAC address filters. Refer to Chapter 38, “Destination MAC Address Filters” on page 475.
- ❑ Authenticator ports cannot be members of LACP or static port trunks or the port mirror.
- ❑ The Guest VLAN feature requires that the designated VLAN already exist on the switch.
- ❑ The Guest VLAN can be a port-based or tagged VLAN.
- ❑ The switch needs to have a IP address to communicate with RADIUS or TACACS+ servers. Refer to Chapter 3, “IPv4 Address” on page 59 or Chapter 4, “IPv6 Address” on page 65.

Here are the guidelines to adding VLAN assignments to supplicant accounts:

- The VLAN can be either a port-based or tagged VLAN.
- The VLAN has to already exist on the switch.
- A supplicant account can have only one VLAN.
- When a supplicant logs on, the switch moves the port to the designated VLAN as an untagged port.

Configuring Port Access Control

Perform the following procedure to configure port-based access control:

1. Select **Security > Port Access Control** from the main menu. Refer to Figure 198.

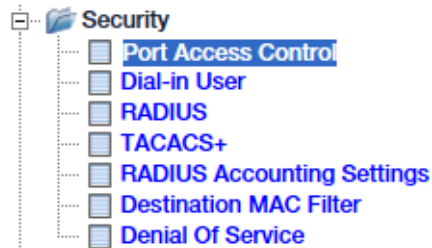


Figure 198. Port Access Control Menu Selection

The Port Access Control Settings window is shown in Figure 199.

Port Access Control Settings

NAS ID:

Port Access Control Status:

Authentication Method:

Figure 199. Port Access Control Settings Window

2. Configure the parameters in Table 74.

Table 74. Port Access Control Settings Window

State	Description
NAS ID	Enter an 802.1x identifier that the switch applies to all ports. The NAS ID can be up to 16 characters. Valid characters are 0 to 9, a to z, and A to Z. Spaces are allowed. The default is fsNAS1. Specifying an NAS ID is optional.
Port Access Control Status	Select the status of port access control from the pull-down menu: <ul style="list-style-type: none"> - Enabled: Enables Port Access Control. - Disabled: Disables the feature.

Table 74. Port Access Control Settings Window (Continued)

State	Description
Authentication Method	Select the authentication method from the menu: <ul style="list-style-type: none"> - RADIUS: Selects remote authentication with RADIUS servers. Refer to the Chapter 35, "RADIUS Client" on page 445 to define the servers. - TACACS+: Selects remote authentication with TACACS+ servers. Refer to Chapter 36, "TACACS+ Client" on page 451 to define the servers. - Local: Selects local authentication. Refer to Chapter 34, "Local Dial-in User Accounts" on page 439 to define the supplicants.

3. Click the **Apply** button.

When you enable the feature, the window adds the port settings. You can also displays the port settings by clicking the **Settings** button. The window can display the settings of only one port at a time. Refer to Figure 200.

Port Access Control Settings

NAS ID: (Max. length: 16 characters)

Port Access Control:

Authentication Method:

Port:

Authentication Mode:

Port Control:

Re-authentication Status:

Control Direction:

Supplicant Mode:

Piggyback Mode:

VLAN Assignment:

Secure VLAN:

Guest VLAN ID: (1-4093)

Transmission Period	30	Sec. (1-65535)	Maximum Request	2	(1-10)
Quiet Period	60	Sec. (1-65535)	Re-authentication Period	3600	Sec. (1-65535)
Supplicant Timeout	30	Sec. (1-65535)	Server Timeout	30	Sec. (1-65535)

Note: In MAC based-authentication mode, re-authentication status is always "Enabled", and default period is 600 sec.

Figure 200. Port Access Control Settings Window - Port Settings

4. Configure the fields in Table 75.

Table 75. Port Access Control Settings Window - Advanced Settings

Field	Description
Port	Select a port to configure from the pull-down menu. You can configure only one port at a time.
Authentication Mode	<p>Select an authentication mode for the port from the pull-down menu:</p> <ul style="list-style-type: none"> - 802.1x: Designates 802.1x as the authentication mode for the port. Clients on ports in this mode have to provide user names and passwords to log in to access the network. - MAC Based: Designates MAC address-based authentication on the port. Clients on ports in this mode are logged on using the MAC addresses of their network devices. The switch sends the initial frames from the clients to the authentication server. The server uses the source MAC addresses in the frames as the user names and passwords for the clients. Supplicant nodes do not need 802.1x client software for this authentication method.
Port Control	<p>Select a control method for the port from the pull-down menu:</p> <ul style="list-style-type: none"> - Forced Authorized: Sets the port to Forced-Authorized port control. Ports transition to the authorized state without any authentication exchanges required. The port transmits and receives traffic normally without 802.1x based authentication of a client. This is the default setting. - Forced Unauthorized: Sets the port to the 802.1x authenticator role, in the unauthorized state. The port blocks all authentications, preventing any client from logging on and forwarding packets. - Auto: Sets the port to the 802.1x port-based authenticator role. The port begins in the unauthorized state, forwarding only EAPOL frames, until a client successfully logs on. <p>Refer to “Authenticator Port Operational Settings” on page 421.</p>

Table 75. Port Access Control Settings Window - Advanced Settings

Field	Description
Re-authentication Status	<p>Select a re-authentication state from the menu:</p> <ul style="list-style-type: none"> - Enabled: Activates re-authentication on the port. The client has to periodically reauthenticate according to the time interval set with the Re-authentication Period. - Disabled: Deactivates re-authentication. The client does not have to reauthenticate after the initial authentication. Reauthentication is still required if there is a change to the status of the link between a client and the switch, or the switch is reset or power cycled. This is the default setting.
Control Direction	<p>Displays the control direction. The default is Both. This parameter cannot be changed.</p>
Supplicant Mode	<p>Select a supplicant mode from the menu:</p> <ul style="list-style-type: none"> - Single: Enables the single supplicant mode. Ports in this mode support only one supplicant at a time. They blocks all other clients once one client has successfully logged in. This is the default setting. Refer to “Single Host Mode” on page 422. - Multiple: Enables the multiple supplicant mode. Ports in this mode permit more than one supplicant at a time to access them. Use the piggyback mode to control whether only one client or all clients have to log in to use ports in this mode. Refer to “Multiple Host Mode” on page 424.

Table 75. Port Access Control Settings Window - Advanced Settings

Field	Description
Piggyback Mode	<p>Select a piggyback mode for the port from the menu:</p> <ul style="list-style-type: none"> - Enabled: Activates the piggyback mode. This mode is used with the multiple supplicant mode. It is commonly used when you want to add 802.1x port-based network access control to a port that is supporting multiple clients, but do not want to add individual accounts for all the clients on the authentication server. After one client has successfully logged on, the port permits other clients to piggy-back onto the initial client's logon, so that they can use the port without being authenticated. Refer to "Single Host Mode with Piggy Backing" on page 422. - Disabled: Deactivates the piggyback mode. This is the default mode. All clients on ports set to the multiple supplicant mode have to have authentication credentials and log in.
VLAN Assignment	<p>Select the VLAN assignment mode from the pull-down menu:</p> <ul style="list-style-type: none"> - Enabled: The switch moves the port to the VLAN specified in the clients login credentials on the authentication device when the clients log on. - Disabled: The port remains in its current VLAN when the client logs in on.
Secure VLAN	Reserved for future use.
Guest VLAN ID	Enter the VID for a Guest VLAN. Clients do not have to log in to access the Guest VLAN. The default is no Guest VLAN. The range is 0 to 4093. Refer to "Guest VLAN" on page 428.
Transmission Period	Enter the switch-to-client retransmission time for EAP request frames. The range is 1 to 65535 seconds. The default is 30 seconds.
Quiet Period	Enter the number of seconds that authenticator ports wait after a failed authentication before accepting authentication requests again. The range is 1 to 65535 seconds.

Table 75. Port Access Control Settings Window - Advanced Settings

Field	Description
Supplicant Timeout	Enter the switch-to-client retransmission time interval for EAP request frames. The range is 1 to 65535 seconds. The default is 30 seconds.
Maximum Request	Enter the maximum number of times authenticator ports transmit EAP Request packets to clients before timing out authentication sessions. The range is 1 to 10. The default is 2.
Re-authentication Period	Enter the time interval in seconds when an authenticator port requires a client to reauthenticate. The range is 1 to 65,535 seconds. The default is 3600 seconds.
Server Timeout	Enter the amount of time in seconds the switch waits for a response from an authentication server. The range is 1 to 65535 seconds. The default is 30 seconds.

5. Click the **Apply** button.
6. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 34

Local Dial-in User Accounts

This chapter describes the local dial-in user accounts feature in the following sections:

- ❑ “Local Dial-in User Authentication Overview” on page 440
- ❑ “Adding Local Dial-in User Accounts” on page 441
- ❑ “Modifying or Deleting Local Dial-in User Accounts” on page 443

Local Dial-in User Authentication Overview

Local dial-in user authentication is part of the 802.1x port-based network access control feature. You can use it instead of RADIUS or TACACS+ servers to authenticate clients on authenticator ports on the switch. The feature consists of an internal database that the switch maintains of client credentials of user names, passwords, and VLAN assignments. As clients log on, the switch refers to its internal database to validate the log on credentials.

Here are the guidelines:

- ❑ The switch supports up to 64 user accounts.
- ❑ You have to configure port-based network access control before activating the local dial-in user authentication. Refer to Chapter 33, “Port Authentication” on page 417.
- ❑ The switch uses standard EAP over LAN (EAPOL) transactions to authenticate clients.

Adding Local Dial-in User Accounts

Perform the following procedure to add local user accounts to the switch:

1. Select **Security > Dial-in User** from the main menu. Refer to Figure 201.




Figure 201. Dial-in User Menu Selection

The Dial-in User window is shown in Figure 202.

Dial-In User

User Name: (Maximum length is 23)

Password: (Maximum length is 23)

Dynamic VLAN: (1-4093)

Free Entries: 64
Total Entries: 0

Username	Password	Dynamic VLAN	Action
<< Table is empty >>			

Page 0 / 0 Page

Figure 202. Dial-In User Window

2. Configure the fields in Table 76.

Table 76. Dial-In User Window

Field	Description
User Name	Enter a unique user name for the account. The maximum length is 23 alphanumeric characters. The user name is not case-sensitive. Spaces are allowed and special characters are not allowed. This field is required.
Password	Enter a password for the account. The maximum length is 23 alphanumeric characters. The password is not case-sensitive. Spaces are allowed and special characters not allowed. This field is required.
Dynamic VLAN	<p>Enter a VID of the 802.1Q tagged VLAN where the switch moves the port when a client logs on. Here are the guidelines:</p> <ul style="list-style-type: none"> - The VID range is 1 to 4093. - An account can have only one VID. - The 802.1Q tagged VLAN has to already exist on the switch. - You cannot specify a Port-based or private VLAN. - Leave this field empty if you want ports to remain in their pre-defined VLANs when clients log on. - If you enter a VID, you have to enable the VLAN Assignment parameter on the switch ports where clients will log on. Refer to “Configuring Port Access Control” on page 433. - The default is no VID.

3. Click the **Add** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying or Deleting Local Dial-in User Accounts

Perform the following procedure to modify or delete local dial-in user accounts:

1. Select **Security > Dial-in User** from the main menu. Refer to Figure 201 on page 441. The Dial-in User window is shown in Figure 202 on page 441.
2. To modify an account, do the following:
 - a. Click its **Modify** button in the Action column of the table. You can modify only one account at a time. The account's details are displayed in the window.
 - b. Modify the account by referring to Table 76 on page 442.
3. To delete accounts, do one of the following:
 - To delete a single account, click its **Delete** button in the Action column in the table.
 - To delete all the accounts on the switch, click the **Delete All** button.
4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 35

RADIUS Client

This chapter describes the RADIUS client in the following sections:

- ❑ “RADIUS Client Overview” on page 446
- ❑ “Adding Server IP Addresses to the RADIUS Client” on page 447
- ❑ “Modifying or Deleting Server IP Addresses from the RADIUS Client” on page 449
- ❑ “Configuring RADIUS Accounting” on page 450

RADIUS Client Overview

Remote Authentication Dial In User Services (RADIUS) is one of the ways that the switch can authenticate network clients of 802.1x port-based access control. The switch has a RADIUS client that allows it to communicate with RADIUS servers on your network. When a network user logs in to the switch, the switch uses its RADIUS client to forward the client's credentials to the RADIUS server for authentication.

The following guidelines apply to the RADIUS client.

- ❑ You have to designate RADIUS as the authentication method for 802.1x port-based access control on the switch. Refer to “Configuring Port Access Control” on page 433.
- ❑ You have to install RADIUS server software on a network server or management station.
- ❑ You can designate up to five RADIUS servers on the switch.
- ❑ You can specify RADIUS servers by their IPv4 or IPv6 addresses.
- ❑ The RADIUS server has to communicate with the switch through a port that is an untagged member of the Default VLAN and is configured for Forced-Authorized (802.1x) port control. Refer to “Configuring Port Access Control” on page 433.
- ❑ If the RADIUS server is on a different subnet from the switch, be sure to specify a System Default Gateway in the IP Setup page, so that the switch and server can communicate with each other via the gateway. Refer to Chapter 3, “Manually Setting the IPv4 Address” on page 61, Chapter 3, “Setting the IPv4 Address from a DHCP or BOOTP Server” on page 64, or “Setting the IPv6 Address” on page 67.
- ❑ You have to specify the user name and password credentials when configuring the RADIUS server software on the authentication server.
- ❑ RADIUS servers have to support Extensible Authentication Protocol (EAP) extensions.

Note

This guide does not explain how to configure RADIUS server software. Refer to the documentation that comes with the server software for instructions.

For more information on the RADIUS authentication protocol, refer to the RFC 2865 standard.

Adding Server IP Addresses to the RADIUS Client

Perform the following procedure to add the IPv4 or IPv6 addresses of up to five RADIUS servers to the client on the switch:

1. Select **Security > RADIUS** from the main menu. Refer to Figure 203

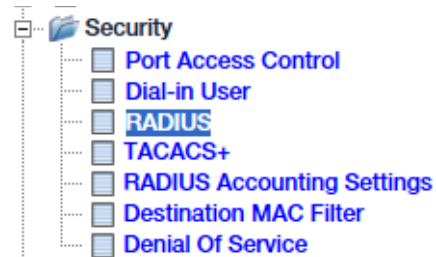


Figure 203. RADIUS Menu Selection

The RADIUS window is shown in Figure 204.

RADIUS

Server Priority: (Highest :1, Lowest :5)

Server IP Address: . . . IPv4 IPv6

Server Port: (1-65535)

Accounting Port: (1-65535)

Shared Secret: (Maximum length is 32)

Server Priority	Server IP Address	Server Port	Accounting Port	Shared Secret	Action
<< Table is empty >>					

Figure 204. RADIUS Window

The table in the window lists the current RADIUS server IP addresses. The columns are defined in Table 77.

2. Configure the fields in Table 77.

Table 77. RADIUS Window

Field	Description
Server Priority	Enter a priority number for the server address. The range is 1 to 5. The switch queries the servers in ascending order, starting with 1. A RADIUS IP address can have only one priority number.

Table 77. RADIUS Window (Continued)

Field	Description
Server IP Address	<p>Enter the IPv4 or IPv6 address of the TACACS+ server. The format for an IPv4 address is shown here:</p> <p>nnn nnn nnn nnn</p> <p>Each N is a decimal number from 0 to 255.</p> <p>The format for an IPv6 address is shown here:</p> <p>xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx</p> <p>Each x is a hexadecimal digit representing four bits. The eight groups of digits have to be separated by colons. Groups where all four digits are '0' can be omitted. Leading '0's in groups can also be omitted.</p> <p>As an example, the following IPv6 addresses are equivalent:</p> <p>3710:421e:09a8:0000:0000:0000:00a4:1c50</p> <p>3710:421e:9a8::a4:1c50</p>
Server Port	<p>Enter the UDP destination port for RADIUS authentication requests. The range is 1 to 65535. The default port is 1812.</p>
Accounting Port	<p>Enter the UDP destination port for RADIUS accounting requests. The range is 1 to 65535. The default port is 1813.</p>
Shared Secret	<p>Enter the encryption key used by the RADIUS server. The maximum length is 32 characters.</p>

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Modifying or Deleting Server IP Addresses from the RADIUS Client

Perform the following procedure to modify or delete IP addresses of RADIUS servers from the client on the switch:

1. Select **Security > RADIUS** from the main menu. Refer to Figure 203 on page 447. The RADIUS window is shown in Figure 204 on page 447.
2. To modify an IP address in the table, perform the following steps:
 - a. Click the **Modify** button in the **Action** column of the IP address to be modified. You can modify only one address at a time.
 - b. Modify the fields in Table 77 on page 447.
3. To delete an IP address, click the **Delete** button of the IP address in the table to be deleted.
4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Configuring RADIUS Accounting

The switch supports RADIUS accounting on ports in the authenticator role. The accounting information sent by the switch to a RADIUS server include the date and time when clients log on or off, as well as the number of packets sent and received by a switch port during a client session. The switch sends the information in packets to the server with the Account-Status Type attribute. For instructions on how to specify the IP addresses of the servers, refer to “Adding Server IP Addresses to the RADIUS Client” on page 447.

Perform the following procedure to enable or disable RADIUS accounting:

1. Select **Security > RADIUS Accounting Settings** from the main menu. Refer to Figure 205.

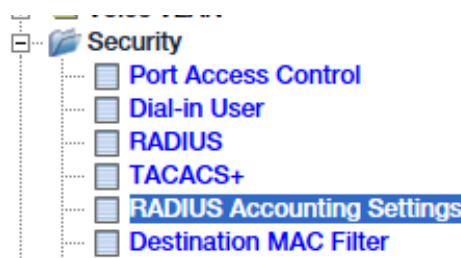


Figure 205. RADIUS Accounting Settings Menu Selection

The RADIUS Accounting Global Settings window is shown in Figure 206.




Figure 206. RADIUS Accounting Global Settings Window

2. Select one of the following from the menu:
 - Enabled:** Enables RADIUS accounting.
 - Disabled:** Disables RADIUS accounting.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 36

TACACS+ Client

The TACACS+ client is described in the following sections:

- “TACACS+ Client Overview” on page 452
- “Adding Server IP Addresses to the TACACS+ Client” on page 453
- “Modifying or Deleting Server IP Addresses from the TACACS+ Client” on page 455

TACACS+ Client Overview

The GS950 PS V2 Switch has a Terminal Access Controller Access-Control System Plus (TACACS+) client. The client is one of several ways that the switch can authenticate supplicants of 802.1x port-based access control. When supplicants log in to authenticator ports, the switch uses the client to transmit the user names and passwords to TACACS+ servers on a network for verification.

Here are the TACACS+ client guidelines.

- ❑ You have to designate TACACS+ as the authentication method for 802.1x port-based access control on the switch. Refer to “Configuring Port Access Control” on page 433.
- ❑ You have to install TACACS+ server software on a network server or management station.
- ❑ You can specify the IP addresses of up to five TACACS+ servers.
- ❑ You can identify the servers by IPv4 or IPv6 addresses.
- ❑ The TACACS+ server has to communicate with the switch through a port that is an untagged member of the Default VLAN and is configured for Forced-Authorized (802.1x) port control.
- ❑ If the TACACS+ server is on a different subnet from switch, be sure to specify a System Default Gateway in the IP Setup window, so that the switch and server can communicate with each other via the gateway. Refer to “Manually Setting the IPv4 Address” on page 61, “Setting the IPv4 Address from a DHCP or BOOTP Server” on page 64, or “Setting the IPv6 Address” on page 67.
- ❑ You have to configure the TACACS+ servers with the user name and password credentials of the network clients.

Note

This guide does not explain how to configure TACACS+ server software. Refer to the software documentation for instructions.

You can also authenticate supplicants of 802.1x port-based access control with RADIUS. The differences between TACACS+ and RADIUS are as follows:

- ❑ TACACS+ separates authentication and authorization in a user profile, whereas, RADIUS combines authentication and authorization.
- ❑ TACACS uses TCP instead of UDP.

Adding Server IP Addresses to the TACACS+ Client

Perform the following procedure to add the IP addresses of up to five TACACS+ servers to the TACACS+ client on the switch:

1. Select **Security > TACACS+** from the main menu. Refer to Figure 207.




Figure 207. TACACS+ Menu Selection

The TACACS+ window is shown in Figure 208.

TACACS+

Server Priority: (Highest :1, Lowest :5)

Server IP Address: . . . IPv4
 IPv6

Server Port: (1-65535)

Timeout: (1-255) sec

Shared Secret: (Maximum length is 32)

Server Priority	Server IP Address	Server Port	Timeout	Shared Secret	Action
<< Table is empty >>					

Figure 208. TACACS+ Window

The table in the window lists the current TACACS+ server IP addresses. The columns are defined in Table 78.

2. Configure the fields in Table 78.

Table 78. TACACS+ Window

Field	Description
Server Priority	Select a unique priority number for the server address from the menu. The range is 1 to 5. The switch queries the server addresses in ascending order, starting with 1. An IP address can have only one priority number.
Server IP Address	<p>Enter the IPv4 or IPv6 address of the TACACS+ server. The format for an IPv4 address is shown here:</p> <p>nnn nnn nnn nnn</p> <p>Each N is a decimal number from 0 to 255.</p> <p>The format for an IPv6 address is shown here:</p> <p>xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx</p> <p>Each x is a hexadecimal digit representing four bits. The eight groups of digits have to be separated by colons. Groups where all four digits are '0' can be omitted. Leading '0's in groups can also be omitted.</p> <p>As an example, the following IPv6 addresses are equivalent:</p> <p>3710:421e:09a8:0000:0000:0000:00a4:1c50</p> <p>3710:421e:9a8::a4:1c50</p>
Server Port	Enter the TCP destination port for TACACS+ authentication requests. The range is 1 to 65535. The default port is 49.
Timeout	Enter the length in time (seconds) the switch waits for a response from a TACACS+ server to an authentication request before querying the next server in the list.
Shared Secret	Enter the encryption key used by the TACACS+ server. The maximum length is 32 characters.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** in the main menu to save your changes.

Modifying or Deleting Server IP Addresses from the TACACS+ Client

Perform this procedure to modify or delete the IP addresses of TACACS+ servers in the client on the switch:

1. Select **Security > TACACS+** from the main menu. Refer to Figure 207 on page 453. The TACACS+ window is shown in Figure 208.
2. To modify an IP address in the table, perform the following steps:
 - a. Click the **Modify** button in the **Action** column of the IP address to be modified.
 - b. Modify the fields in Table 78 on page 454.
3. To delete an IP address, click the **Delete** button in the **Action** column of the IP address to be deleted.
4. Click the **Apply** button.
5. Select **Save Settings to Flash** in the main menu to save your changes.

DHCP Snooping

This chapter describes DHCP snooping in the following sections:

- “DHCP Snooping Overview” on page 458
- “General Guidelines” on page 463
- “Configuring DHCP Snooping” on page 464
- “Adding DHCP Snooping to VLANs” on page 467
- “Designating Trusted and Untrusted Ports” on page 469
- “Managing the Binding Database” on page 471

DHCP Snooping Overview

DHCP snooping is a network security feature. It allows the switch to monitor and block packets from unauthorized DHCP servers.

The switch monitors ARP packets from nodes and discards invalid packets that contain fake or counterfeit IP addresses. The switch maintains a table called the binding database that contains the MAC and IP addresses of the nodes on its ports. When it receives an ARP packet it compares the packet's source MAC and IP addresses against the addresses in the database. If there is no match, it discards the packet.

security by inspecting ingress packets for the correct IP and MAC address information. The DHCP snooping feature defines the GS950 PS V2 ports as either trusted or untrusted. With DHCP snooping enabled, two network security issues are addressed:

- ❑ All ingress DHCP packets are examined on the untrusted ports, and only authorized packets are passed through the switch. Unwanted ingress DHCP packets are discarded. Refer to "Unauthorized DHCP Servers," next below.
- ❑ DHCP ingress packets on an untrusted port are inspected to insure that the source IP Address and MAC Address combination in each packet is valid when compared to the DHCP snooping Binding table. If a match is not found, the packet is discarded.

Trusted Ports

By definition, trusted ports inherently trust all ingress Ethernet traffic. There is no checking or testing on ingress packets for this type of port. A trusted port connects to a DHCP server in one of the following ways:

- ❑ Directly to the legitimate trusted DHCP Server
- ❑ A network device relaying DHCP messages to and from a trusted server
- ❑ Another trusted source such as a switch with DHCP snooping enabled

Untrusted Ports

Ethernet traffic on untrusted ports are inherently not trusted. The switch tests ingress packets against specified criteria to determine if they should be forwarded or discarded. Untrusted ports should be connected to DHCP clients and to traffic that originates outside of the LAN.

Unauthorized DHCP Servers

A local area network typically has a single DHCP server that supplies the network configurations, including as IP addresses, to the network devices. A switch port connected to a trusted DHCP server should be designated as a trusted port.

It is possible that an unauthorized DHCP server could be connected to the

network. This situation might occur if a network client activates a DHCP server application or if someone outside the network sends DHCP packets to your network. These situations pose a security risk.

A network device initially sends out a DHCPDISCOVER packet so that a DHCP server will respond. It waits for, and then accepts, the first DHCPOFFER packet from the server that it receives. This packet contains the DHCP server's IP address and mask. If the unauthorized DHCP server responds first, then the network device will use the information from the unauthorized DHCP server for the default gateway or DNS server.

Untrusted ports are connected to the DHCP clients and to traffic that originated outside the LAN. By definition, untrusted ports do not accept DHCP packets originating from a DHCP server and immediately drop them when they are detected. The DHCP packets types that are not accepted are DHCPOFFER and DHCPACK.

However, untrusted ports do accept both DHCP DISCOVER and DHCPREQUEST packets sent from DHCP clients. This behavior allows DHCP clients to respond to a trusted DHCP server and not respond to a DHCP server that is untrusted.

DHCP with Option 82

When the switch receives DHCP request packets from end nodes, it can add information before forwarding them on to the DHCP server. Referred to as Option 82, the information can be used to identify the physical locations of the nodes. Here is the Option 82 information that the switch can add to the packets.

- Remote ID: This is the MAC address of the switch.
- Circuit ID: This is the port number on the switch where the node is connected and the VLAN ID where the port is a member.
- Subscriber ID: This is not applicable to the GS950 PS V2 switch.

The process is transparent to the end nodes because they never see the Option 82 information. The switch adds the information to the ingress DHCP request packets from end nodes before forwarding the packets to the DHCP server, and removes the information from the ingress reply packets from the DHCP server. The process is illustrated in Figure 209.

Option 82 can be enabled or disabled on the switch. The default setting is disabled. If the switch is at the network edge and there are end nodes directly connected to untrusted ports, enabling the option adds the information to the DHCP request packets from the nodes.

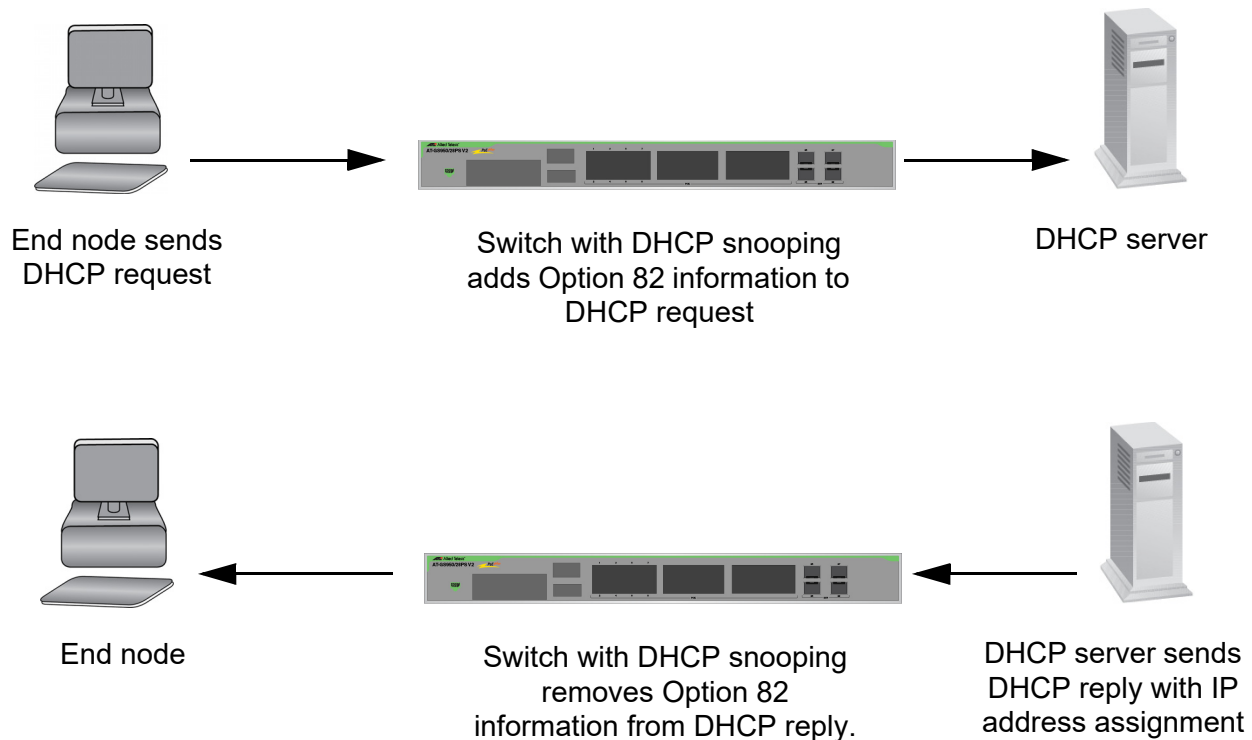


Figure 209. DHCP with Option 82

Option 82 Pass Through

As explained in the previous section, Option 82 is information the switch adds to DHCP request packets from end nodes to identify the node locations. At its default settings, the switch discards ingress DHCP request packets that already contain Option 82 information. This can occur when a switch port is connected to another switch that is adding Option 82 information to DHCP requests.

You can control this with the Option 82 pass through feature. When you enable the feature, the switch forwards ingress DHCP request packets that already contain Option 82 information. Figure 210 on page 461 illustrates what happens when Option 82 pass through is disabled and Figure 211 on page 462 when the option is enabled.

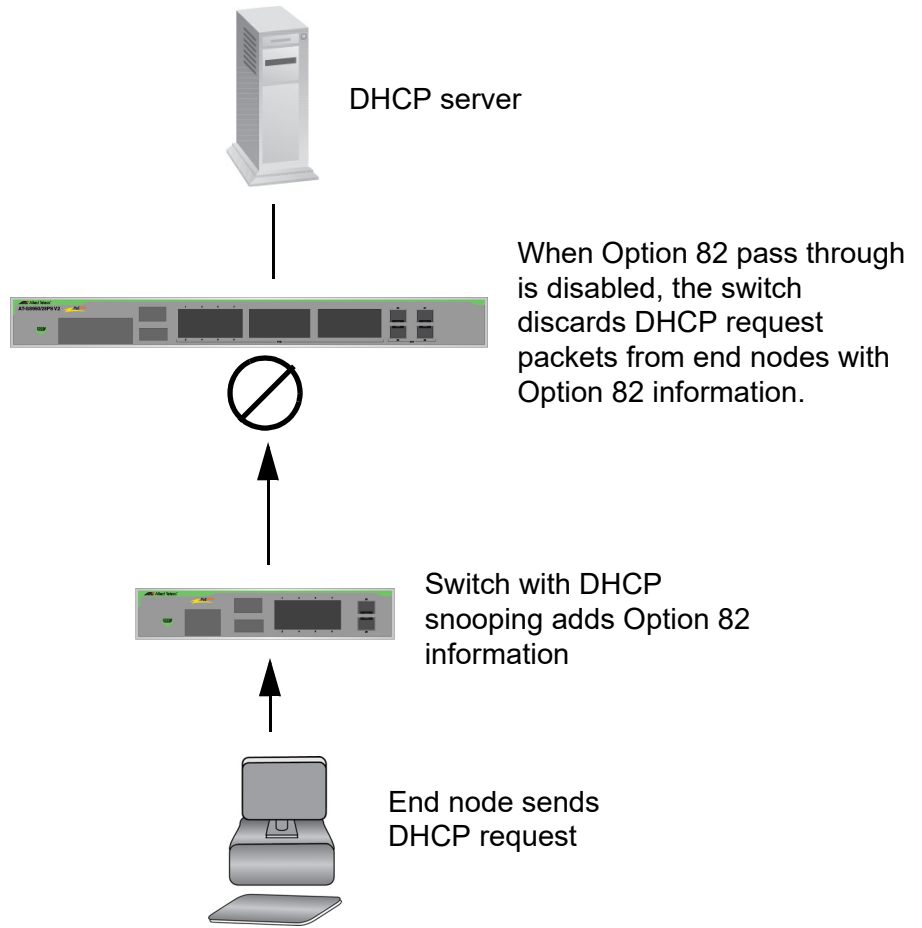


Figure 210. DHCP with Option 82 Pass Through Disabled

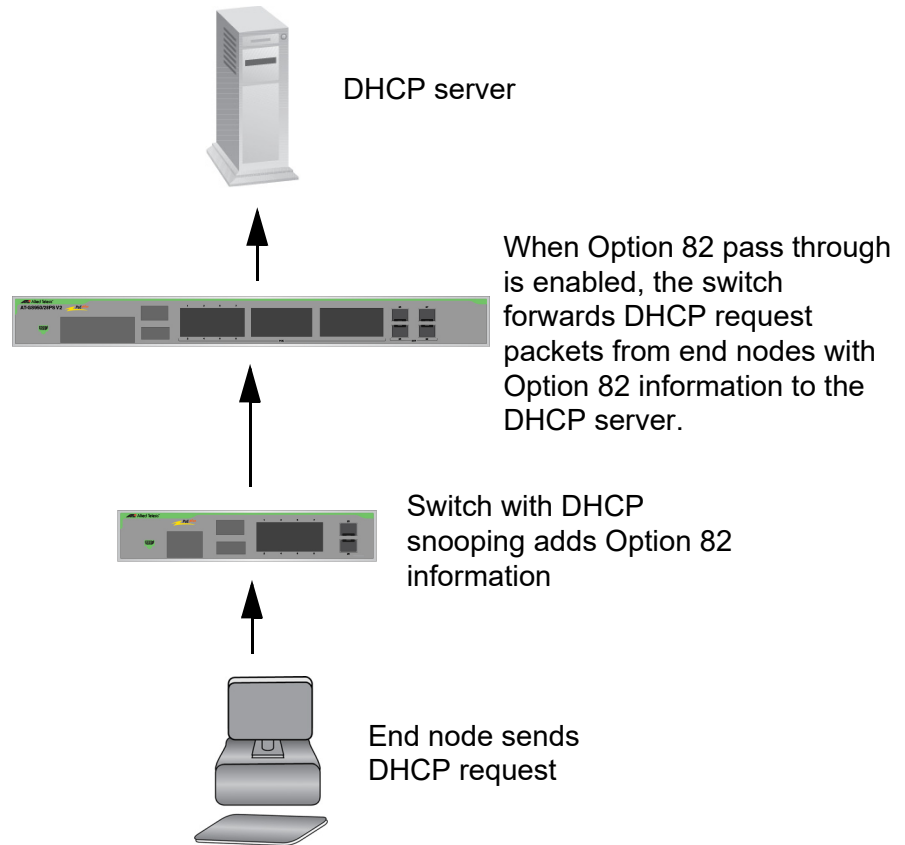


Figure 211. DHCP with Option 82 Pass Through Enabled

General Guidelines

Here is a summary of the rules to observe when configuring DHCP snooping:

- ❑ A trusted port is connected to one of the following:
 - Directly to a trusted DHCP Server.
 - A network device relaying DHCP messages to and from a trusted server.
 - Another trusted source, such as a switch with DHCP snooping enabled.
- ❑ Untrusted ports are connected to DHCP clients and to traffic that originates outside of the local area network.
- ❑ The VLANs to which the DHCP snooping feature applies must be specified in the DHCP snooping VLAN Setting configuration.
- ❑ Static IP addresses on the network have to be manually added to the Binding Database.

Configuring DHCP Snooping

Perform the following procedure to enable or disable DHCP snooping and configure its general settings:

1. Select **DHCP Snooping > General Settings** from the main menu. Refer to Figure 212.




Figure 212. DHCP Snooping General Settings

The General Settings window is shown in Figure 213.

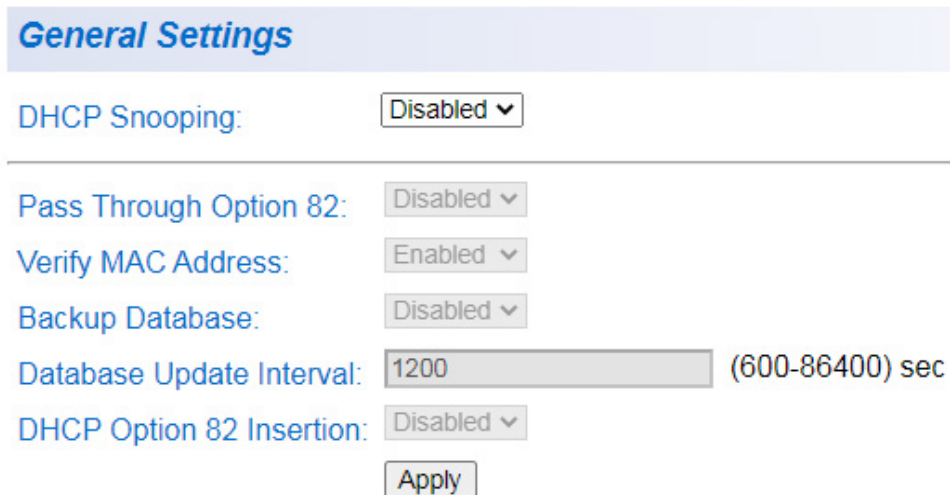


Figure 213. DHCP Snooping General Settings Window

2. Configure the fields in Table 79.

Table 79. DHCP Snooping General Settings Window

Field	Description
DHCP Snooping	<p>Select one of the following from the menu:</p> <ul style="list-style-type: none"> - Enabled: Enables DHCP snooping on the designated VLANs on the switch. Refer to “Adding DHCP Snooping to VLANs” on page 467. You have to enable DHCP snooping before you can configure the settings. - Disabled: Disables DHCP snooping. This is the default setting.

Table 79. DHCP Snooping General Settings Window (Continued)

Field	Description
Pass Through Option 82	<p>Select one of the following choices from the menu:</p> <ul style="list-style-type: none"> - Enabled - Configures the switch to forward DHCP packets with Option 82 information from nodes on untrusted ports to DHCP servers. Refer to “Option 82 Pass Through” on page 460. - Disabled - Configures the switch to discard DHCP packets with Option 82 information from nodes on untrusted ports. This is the default setting.
Verify MAC Address	<p>Select one of the following choices from the menu:</p> <ul style="list-style-type: none"> - Enabled - Configures the switch to examine the MAC and IP addresses of ingress ARP packets on untrusted ports and to forward only those packets that have a match in the binding database. Invalid ARP packets are discarded. This is the default setting. - Disabled - Configures the switch to not verify the MAC and IP addresses of ingress ARP packet against the binding table. The switch forwards all ARP packets without regard to their IP and MAC addresses.
Backup Database	<p>Select one of the following choices from the menu:</p> <ul style="list-style-type: none"> - Enabled - Configures the switch to save a backup copy of the binding database to flash memory at the specified time interval. - Disabled - Prevents the switch from saving a backup copy of the binding database to flash memory. This is the default setting.
Database Update Interval	<p>Enter the time interval for backing up the binding database. The range is 600 to 86400 seconds. The default is every 1200 seconds (20 minutes).</p>

Table 79. DHCP Snooping General Settings Window (Continued)

Field	Description
DHCP Option 82 Insertion	Select one of the following options: - Enabled - Inserts DHCP Option 82 information into DHCP request packets from hosts on untrusted ports and removes them from reply packets from DHCP servers. Refer to “DHCP with Option 82” on page 459. - Disabled - Disables the addition of DHCP Option 82 information into DHCP packets. This is the default setting.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.
5. Go to “Designating Trusted and Untrusted Ports” on page 469.
6. Go to “Adding DHCP Snooping to VLANs” on page 467.

Adding DHCP Snooping to VLANs

This section contains the procedure for adding DHCP snooping to 802.1Q tagged VLANs. Here are the guidelines:

- ❑ DHCP snooping is supported on 802.1Q tagged VLANs. Refer to Chapter 16, “802.1Q Tagged Virtual LANs” on page 205.
- ❑ DHCP snooping is not supported on port-based or private VLANs

Perform the following procedure to add DHCP snooping to 802.1Q tagged VLANs:

1. Select **DHCP Snooping > VLAN Settings** from the main menu. Refer to Figure 214.




Figure 214. DHCP Snooping VLAN Settings Menu Selection

The DHCP snooping VLAN Settings window is shown in Figure 215.

VLAN Settings

VLAN ID: (1~4093)

VLAN Table :

VLAN ID	Action
<< Table is empty >>	

Page 0 / 0 Page

Figure 215. DHCP Snooping VLAN Settings Window

2. To add DHCP snooping to an 802.1Q tagged VLAN, do the following:
 - a. Enter the VID of VLAN in the **VLAN ID** field. You can add only one VLAN at a time.

- b. Click the **Add** button. The VLAN is added to the table. (The switch displays an error message if you enter the VID of a nonexistent 802.1Q tagged VLAN or the VLAN index number of a port-based VLAN.)
 - c. Repeat this step to add DHCP snooping to more 802.1Q tagged VLANs.
 3. To remove DHCP snooping from VLANs, do one of the following:
 - To remove DHCP snooping from a single VLAN, click the **Delete** button in the Action column of the entry.
 - To remove DHCP snooping from all the VLANs, click the **Delete All** button.
 4. Select **Save Settings to Flash** from the main menu to save your changes.
 5. Go to “Designating Trusted and Untrusted Ports” on page 469 to designate ports in the VLAN as trusted or untrusted. Trusted ports are connected to authorized DHCP servers and untrusted ports to clients of the DHCP servers.

Designating Trusted and Untrusted Ports

For background information, refer to “Trusted Ports” on page 458 and “Untrusted Ports” on page 458. DHCP snooping has to be enabled for you to configure ports. Refer to “Configuring DHCP Snooping” on page 464.

Perform the following procedure to designate ports as trusted or untrusted for DHCP snooping:

1. Select **DHCP Snooping > Trusted Interfaces** from the main menu. Refer to Figure 216.




Figure 216. DHCP Snooping Trusted Interfaces Menu Selection

The DHCP Snooping Trusted Interfaces window is shown in Figure 217.

<i>Trusted Interfaces</i>		
Port	Trust	Action
All	Enabled ▼	Apply
1	Enabled ▼	Apply
2	Enabled ▼	Apply
3	Enabled ▼	Apply
4	Enabled ▼	Apply

Figure 217. DHCP Snooping Trusted Interfaces Window

2. In the Trust column for a port, select one of the following options from the menu:
 - Disabled** - Defines ports as untrusted. This is the correct setting for ports connected to hosts or to devices outside the network. The switch will forward DHCP packets to and from hosts on the port, but will discard ingress packets from DHCP servers. This protects against unauthorized DHCP servers connecting to the network.
 - Enabled** - Defines ports as trusted. This is the correct setting for posts connected to authorized DHCP servers. The switch will forward packets to and from the DHCP server on the port. This is the default setting.

3. Click the **Apply** button in the **Action** column for the port.
4. Repeat steps 2 and 3 to designate other ports as trusted or untrusted.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Managing the Binding Database

The DHCP snooping binding database contains the MAC addresses and IP addresses of the hosts on untrusted switch ports. (As explained earlier, untrusted ports are typically connected to end nodes or to devices outside the network.) The IP addresses can either be those assigned by DHCP servers or those you assigned manually. Along with the addresses the table lists the switch ports to which the hosts are connected and their VLAN assignments.

The switch builds the database by examining DHCP packets from DHCP servers on trusted ports and hosts on untrusted ports. You can also add to the database by entering static entries.

The switch uses the binding database to prevent unauthorized individuals from accessing the network on untrusted ports by transmitting fake ARP messages containing valid IP addresses. As ARP packets arrive on the ports, the switch compares their source MAC and IP addresses against the host entries in the table and discards packets that do not match the table entries.

Perform the following procedure to view or edit the entries in the binding database:

1. Select **DHCP Snooping > Binding Database** from the main menu. Refer to Figure 218.




Figure 218. DHCP Snooping Binding Database Menu Selection

The Binding Database window is shown in Figure 219 on page 472.

Binding Database

MAC Address : : : : : : (e.g. 00:11:ab:cd:ef:22)

IP Address: . . . IPv4
 : : : : : IPv6

VLAN : (1~4093)

Port: ▼

Type ▼

Lease Time : (10 - 4294967295) sec

Free Policies :188 Total Entries: 0

MAC Address	VLAN ID	IP Address	Port	Type	Lease Time	Action
<< Table is empty >>						

Page 0 / 0 Page

Figure 219. DHCP Snooping Binding Database Window

The table in the window displays the current dynamic and static entries of the IP address assignments to hosts. Refer to Table 80.

Table 80. DHCP Snooping Binding Database Window

Column	Description
MAC Address	Displays the host’s MAC address.
VLAN	Displays the VLAN ID of the port.
IP Address	Displays the host’s static IPv4 or IPv6 address.
Port	Displays the port number where the host is connected.
Type	Displays one of the following: - Static: The entry was entered manually into the database. - Learned: The switch and DHCP snooping learned the entry from a DHCP sever.
Lease Time	Displays the amount of time the IP address from the DHCP server is valid, before it has to be reauthorized. Does not apply to static entries.

2. To add a static entry, do the following:

- a. Enter the fields in Table 81.

Table 81. DHCP Snooping Binding Database Window

Field	Description
MAC Address	Enter the host's MAC address.
IP Address	Enter the host's static IPv4 or IPv6 address.
VLAN	Enter the VLAN ID.
Port	Select the port number where the host is connected, from the pull-down menu.
Type	Select Static from the pull-down menu.
Lease Time	Not applicable to static IP addresses.

- b. Click the **Add** button.
3. To delete entries, do one of the following:
- To delete a single entry, click its **Delete** button in the Action column.
 - To delete all the entries, click the **Delete All** button above the table.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 38

Destination MAC Address Filters

This chapter describes the destination MAC address filters in the following sections:

- ❑ “Destination MAC Address Filters Overview” on page 476
- ❑ “Adding Destination MAC Address Filters” on page 477
- ❑ “Deleting Destination MAC Address Filters” on page 478

Destination MAC Address Filters Overview

Destination MAC address filters are a network security feature that block network users from accessing specified devices on your network. The restricted devices are identified by their MAC addresses. The switch discards ingress packets whose destination MAC addresses match the filters.

You can use the feature to block switch clients from gaining unauthorized access to secured devices on your network. For instance, if the users connected to the switch are part of the Sales department, you could add filters to block them from accessing the accounting and engineering servers.

Here are the guidelines to destination MAC address filters:

- ❑ The switch supports up to 188 filters.
- ❑ Each filter can have one destination MAC address.
- ❑ Filters have to contain the MAC addresses of specific devices. Filters cannot contain MAC address ranges.
- ❑ The filters are set at the switch level and apply to all switch ports.

Adding Destination MAC Address Filters

Perform the following procedure to add destination MAC address filters:

1. Select **Security** > **Destination MAC Filter** from the main menu. Refer to Figure 220.

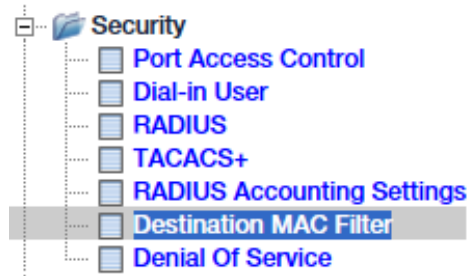


Figure 220. Destination MAC Filter Menu Selection

The Destination MAC Filter window is shown in Figure 221.

Destination MAC Filter

MAC Address : : : : : : (e.g. 00:11:ab:cd:ef:22)

Free Policies : 188
Total Entries : 0

MAC Address	Action
<< Table is empty >>	

Page 0 / 0 Page

Figure 221. Destination MAC Filter Window

The table in the window lists the current MAC address filters.

2. Enter a MAC address of a destination device in the **MAC Address** field to block network clients from accessing the device.
3. Click the **Add** button. The switch immediately activates the new filter.
4. Repeat steps 2 and 3 to add more MAC address filters.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Deleting Destination MAC Address Filters

Perform the following procedure to delete destination MAC address filters:

1. Select **Security > Destination MAC Filter** from the main menu. Refer to Figure 220 on page 477. The Destination MAC Filter window is shown in Figure 221 on page 477.
2. Perform one of the following:
 - To delete a single filter, click its **Delete** button in the **Action** column.
 - To delete all the filters, click the **Delete All** button above the table.
3. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 39

Port Mirroring

This chapter describes port mirroring in the following sections:

- “Port Mirroring Overview” on page 480
- “Enabling Port Mirroring” on page 481
- “Disabling Port Mirroring” on page 483

Port Mirroring Overview

Port mirroring lets you unobtrusively monitor the traffic received or transmitted on one or more ports by copying the traffic to another switch port. By connecting a data analyzer to the port where the traffic is copied, you can monitor the traffic on the other ports without impacting network performance or speed. The feature can be used to troubleshoot network problems or investigate possible network attacks.

Port mirroring has two components. The ports whose traffic you want to mirror are called *source port(s)*. The port where the traffic will be copied is called the *mirroring port*.

Observe the following guidelines when using the port mirror:

- ❑ You can designate only one mirroring port.
- ❑ You can designate more than one source port. However, the more ports you mirror, the less likely the mirroring port will be able to handle all the traffic. For example, if you mirror the traffic of six heavily active ports, the mirroring port will likely drop packets, meaning that it will not provide an accurate mirror of the traffic of the source ports.
- ❑ The source and mirroring ports have to be located on the same switch.
- ❑ You can mirror the ingress traffic, egress traffic or both of the source ports.
- ❑ The mirroring port cannot be a member of a static or LACP trunk.
- ❑ The mirroring port is dedicated to monitoring the traffic from the source ports. It cannot be used for regular network operations.

Enabling Port Mirroring

This section explains how to activate port mirroring. You need to know the following to use the feature:

- ❑ Which port will be the mirroring port? The switch copies the traffic from the source port to this port. There can be only one mirroring port. The default is port 1.
- ❑ Which ports will be the source ports? These are the ports whose traffic you want to monitor. There can be more than one source port.
- ❑ Do you want to mirror the ingress traffic, egress traffic, or both on the source ports?

Perform the following procedure to enable port mirroring:

1. Select the **Bridge > Mirroring** from the main menu. Refer to Figure 222.



Figure 222. Mirroring Menu Selection

The Mirroring window is shown in Figure 223.

Mirroring

Status: Disabled ▾

Mirror Target Port: 1 ▾

Ingress Port:

All	1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28
	<input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>

Egress Port:

All	1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28
	<input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>

Apply

Figure 223. Mirroring Window

2. Select **Enabled** from the Status menu. You have to enable the feature to configure it.
3. Select the mirroring port from the Mirror Target Port menu. You can select only one mirroring port.
4. Click the check boxes of the ports that are to be source ports. Use the Ingress Port row to mirror ingress traffic on source ports and the Egress Port row to mirror egress traffic on source ports. To select all ports for ingress or egress traffic monitoring, click the **All** button.
5. Click the **Apply** button. The switch immediately activates port mirroring.
6. Select **Save Settings to Flash** from the main menu to save your changes.
7. Connect a data analyzer to the mirroring port to monitor the Ethernet traffic on the source ports.

Disabling Port Mirroring

Perform the following procedure to disable port mirroring:

8. Select **Bridge > Mirroring** from the main menu. Refer to Figure 222 on page 481. The Mirroring window is shown in Figure 223 on page 481.
9. Select **Disabled** from the Status field.
10. Click the **Apply** button.

Port mirroring is now disabled on the switch. Traffic on the source ports are no longer copied to the mirroring port. The parameters in the window become inactive. You can now use the mirroring port for regular network operations.

11. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 40

Denial of Service

This chapter contains the following sections:

- “Configuring Denial of Service Protection” on page 486

Configuring Denial of Service Protection

The switch has filters for blocking selected types of TCP/IP packets that might pose network security risks. The filters are set at the switch level and apply to all the ports. Perform the following procedure to configure the Denial or Service protection filters:

1. Select **Security > Denial or Service** from the main menu. Refer to Figure 224.




Figure 224. Denial of Service Menu Selection

The Denial of Service window is shown in Figure 225.

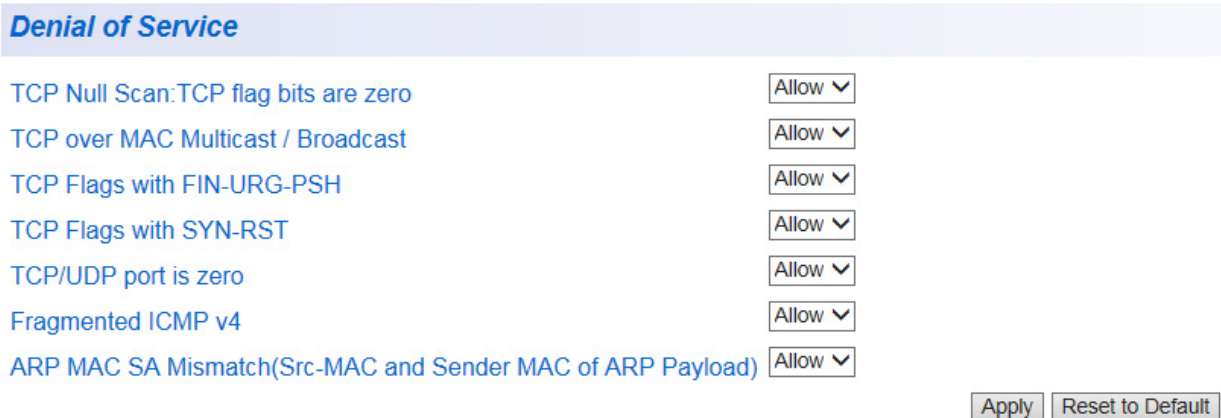


Figure 225. Front Panel Window

2. Adjust the filters. The options are listed here:

- Allow: The switch forwards the packets.
- Deny: The packets are discarded.

The filters are described in Table 82 on page 487.

Table 82. Denial of Service Protection

Blocked Packets	Description
TCP Null Scan: TCP flag bits are zero	Blocks TCP packets that do not have any flags set.
TCP over MAC Multicast/Broadcast	Blocks TCP packets with multicast or broadcast IP addresses.
TCP Flags with FIN- URG-PSH	Blocks TCP packets with the FIN-URG-PSH flag set.
TCP Flags with SYN-RST	Blocks TCP packets with the SYN-RST flag set.
TCP/UDP port is zero	Blocks TCP/UDP packets with the port number 0.
Fragmented ICMP v4	Blocks ICMP IPv4 fragmentation packets.
ARP MAC SA Mismatch (Src-MAC and Sender MAC of ARP Payload)	Blocks ARP packets that contain different source and sender MAC addresses.

3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Section VIII

Troubleshooting Tools and Maintenance

This section contains the following chapters:

- ❑ Chapter 41, “System Log and Syslog Client” on page 491
- ❑ Chapter 42, “Troubleshooting Tools” on page 499
- ❑ Chapter 43, “Loopback Detection” on page 507
- ❑ Chapter 44, “Switch Firmware” on page 513

Chapter 41

System Log and Syslog Client

This chapter contains the following sections:

- “Viewing the Event System Log” on page 492
- “Sending System Log Events to a Syslog Server” on page 495

Viewing the Event System Log

A managed switch is a complex piece of computer equipment that includes both hardware and software components. Multiple software features operate simultaneously, inter-operating with each other and processing large amounts of network traffic. It is often difficult to determine exactly what is happening when a switch appears not to be operating normally, or what happened when a problem occurred.

You can monitor the operations of the switch by viewing the event messages generated by the device. These events and the vital information about system activity they provide can help you identify and solve system problems.

The events are stored by the switch in an event log, in temporary memory. The events in the log are discarded whenever you reset or power cycle the switch.

Perform this procedure to view the event log:

1. Select **System** -> **System Log Settings** from the main menu. Refer to Figure 226.

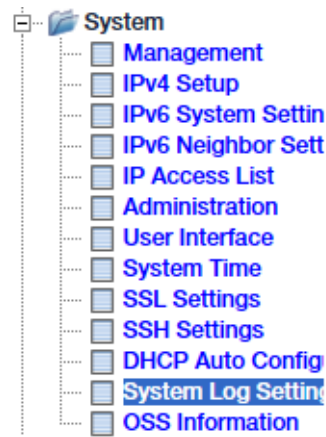


Figure 226. System Log Settings Menu Selection

The System Log Settings window is shown in Figure 227 on page 493.

System log Settings

Time Stamp

Messages Buffered Size (1-200)

Syslog Status

Syslog Server IP . . . IPv4 IPv6

Facility

Logging Level

1	local0/Info	01/01/2020 00:00:03	Port 1 link up, 1Gbps FULL duplex
2	local0/Info	01/01/2020 00:00:03	Port 2 link up, 1Gbps FULL duplex
3	local0/Info	01/01/2020 00:00:03	Port 3 link up, 1Gbps FULL duplex
4	local0/Info	01/01/2020 00:00:03	Port 4 link up, 1Gbps FULL duplex

Figure 227. System Log Settings Window

The Syslog fields are described in “Sending System Log Events to a Syslog Server” on page 495.

The window is static. The switch does not update the window. To view new events, click the Refresh button at the bottom of the window. Events have the following parts:

- Index number
- Facility number (local0)
- Severity level
- Date
- Time
- Description

At the bottom of the window are the following buttons:

- Clear - Clears all log events from the switch.
 - Refresh - Updates the window with the latest events.
2. To add or omit the time and date from event messages, select one of the following from Time Stamp:
- Enabled:** The switch adds the time and date to events. This is the default setting.
 - Disabled:** The switch omits the time and date from event messages.

3. To set the limit on the number of events the log can store, select the Messages Buffered Size field. The range is 1 to 200 messages. The default is 50 messages. Once reaching its maximum capacity, the log deletes the oldest messages as it stores new messages.
4. Click the **Apply** button.
5. Select **Save Settings to Flash** from the main menu to save your changes.

Sending System Log Events to a Syslog Server

The system log events can play an important role in diagnosing and resolving network problems. However, when problems do occur it may not always be immediately evident which device was the cause. Troubleshooting problems might require the burdensome and time consuming task of reviewing system logs on multiple managed devices.

Having the managed devices send their events to a central server for storage can simplify troubleshooting problems. Rather than having to view the events on the individual devices, you can instead view them on a central server.

That is the function of the syslog client on the switch. It allows the switch to send its event messages to a syslog server on your network, for storage.

Here are the guidelines to using the syslog client:

- You can specify one syslog server.
- The switch has to have a management IP address. For instructions, refer to Chapter 3, “IPv4 Address” on page 59 or Chapter 4, “IPv6 Address” on page 65.
- If the syslog server is not a member of the same subnet as the management IP address of the switch, the switch has to have a default gateway specifying the first hop to reaching the server. For instructions on specifying the default gateway, refer to Chapter 3, “IPv4 Address” on page 59 or Chapter 4, “IPv6 Address” on page 65.
- The event messages are transmitted when they are generated. Any event messages already in the event log are not transmitted when you enable the syslog client.
- The syslog client uses UDP port 514. You cannot change the UDP port.

Perform the following procedure to configure the syslog client:

1. Select **System** -> **System Log Settings** from the main menu. Refer to Figure 226 on page 492 and Figure 227 on page 493.
2. Configure the fields in Table 83 on page 496.

Table 83. System Log Settings Window - Syslog Client

Field	Description
Syslog Status	<p>Select one of the following:</p> <ul style="list-style-type: none"> - Enabled: Activates the Syslog client so that the switch sends events to a Syslog server. Syslog Status has to be set to Enabled for you to configure the other settings. - Disabled: Deactivates the Syslog client to stop the switch from sending events to a Syslog server. This is the default sending.
Syslog Server IP	<p>Enter the IPv4 or IPv6 address of the Syslog server on your network. You can enter only one address. The format for an IPv4 address is shown here:</p> <p>nnn nnn nnn nnn</p> <p>Each N is a decimal number from 0 to 255.</p> <p>The format for an IPv6 address is shown here:</p> <p>xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx</p> <p>Each x is a hexadecimal digit representing four bits. The eight groups of digits have to be separated by colons. Groups where all four digits are '0' can be omitted. Leading '0's in groups can also be omitted.</p> <p>As an example, the following IPv6 addresses are equivalent:</p> <p>3710:421e:09a8:0000:0000:0000:00a4:1c50</p> <p>3710:421e:9a8::a4:1c50</p>
Facility	<p>Select a facility level to add to the event messages as the switch transmits them to the syslog server. Facility levels are numerical codes that are commonly used to group entries on the syslog server according to the source network devices. You can specify only one facility level. Refer to RFC 3164 for facility code definitions. The range is local0 to local7.</p>

Table 83. System Log Settings Window - Syslog Client (Continued)

Field	Description
Logging Level	<p>Select the severity level of the log events to send to the Syslog server. The levels are listed here from highest to lowest severity:</p> <ul style="list-style-type: none"> - Alert - Critical - Warning - Info <p>Selecting a severity included the messages of that level and those levels above it. The default severity is Info. At the default setting, The switch transmits alert, critical, and warning messages. To send messages from all severities, select Info.</p>

3. Click the **Apply** button.

If you enabled the Syslog client, the switch begins transmitting log events as they occur, to the Syslog server. It does not transmit messages already stored in the event log.

4. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 42

Troubleshooting Tools

This chapter describes the following troubleshooting tools:

- ❑ “Rebooting the Switch” on page 500
- ❑ “Restoring the Factory Default Values” on page 502
- ❑ “Locking the Factory Reset Options” on page 503
- ❑ “Pinging Network Devices” on page 505

Rebooting the Switch

This section contains the procedure for rebooting the switch. You might reboot the switch if it is experiencing a problem or to discard unsaved configuration changes. You can also reboot the device by pressing the eco-Friendly button on the front panel for five to nine seconds, as explained in the *GS950 PS V2 Series Installation Guide*.

Note

Configuration settings not saved in the configuration file are discarded when the switch is rebooted. To save the current configuration, select **Save Settings to Flash** from the main menu before rebooting the switch.



Caution

The switch stops forwarding network traffic for approximately two minutes while it initializes its management software. Some network traffic may be lost.

Perform the following procedure to reboot the switch:

1. Select **Tools > Reboot** from the main menu. Refer to Figure 228.




Figure 228. Reboot Menu Selection

The Factory Default Reset/Reboot window is shown in Figure 229.

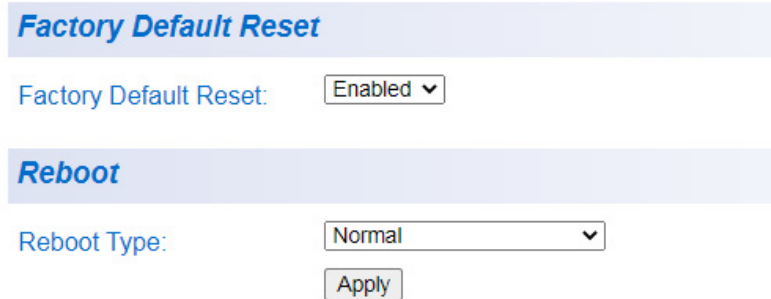


Figure 229. Factory Default Reset/Reboot Window

2. Select **Normal** from the Reboot Type menu.

Note

The Factory Default and Factory Default Except IP selections in the Reboot Type menu are described in “Restoring the Factory Default Values” on page 502.

3. Click the **Apply** button.

The switch reboots and initializes its management program, a process that takes approximately two minutes to complete. After the reboot is finished, you can log in again to resume managing the switch.

Restoring the Factory Default Values

This section contains the procedure for restoring the factory default settings to the switch. You might perform this procedure to discard the current configuration before configuring the switch with new settings.



Caution

The switch stops forwarding network traffic for approximately two minutes while it initializes the management software. Some network traffic may be lost.

Perform the following procedure to restore the default settings:

1. Select **Tools > Reboot** from the main menu. Refer to Figure 228 on page 500. The Factory Default Reset/Reboot window is shown in Figure 229 on page 500.
2. Select one of the following from the Reboot Type menu:
 - Factory Default** - Returns all the switch parameters, including the IPv4 and IPv6 addresses, to the factory default settings. The IPv4 address is return to the default 192.168.1.1 and the DHCP client is disabled. There is no default IPv6 address.
 - Factory Default Except IP** - Returns all the switch parameters, excluding the IPv4 and IPv6 addresses, to the factory default settings. If the DHCP client is enabled, it remains enabled after this reset.
3. If prompted for a password, enter the password that was used to lock the eco-Friendly button and default settings in “Disabling the eco-Friendly Button and Default Settings” on page 503. The password is case sensitive.

Note

You cannot recover or reset the password to the eco-Friendly button and default settings if it is lost or forgotten.

4. Click the **Apply** button.

The switch reboots, initializes the management program, and restores the default settings, a process that takes approximately two minutes to complete. When the reboot is finished, you can log in again to continue managing the switch. If the IPv4 address was returned to the default value 192.168.1.1, refer to “Starting the First Management Session” on page 32 for instructions.

Locking the Factory Reset Options

You can reboot the switch as well as restore the default feature settings with the eco-Friendly button on the front panel of the device. If the switch is installed in a non-secured area, you may want to disable the button to prevent unauthorized individuals from using it to reboot the device and disrupt network operations. You can also add password protection to restoring the default feature settings from the management software so that only those individuals who know the password can restore the settings.



Caution

Keep the password in a safe place. You cannot recover or reset it if it is lost or forgotten.

Disabling the eco-Friendly Button and Default Settings

Perform the following procedure to disable the eco-Friendly button on the front panel of the switch and add password protection to restoring the default feature settings:

1. Select **Tools > Reboot** from the main menu. Refer to Figure 228 on page 500. The Factory Default Reset/Reboot page is shown in Figure 229 on page 500.
2. Select **Disabled** from the **Factory Default Reset** menu. The window adds Password fields.
3. In the **New Password** field, enter a password of up to 12 alphanumeric characters in length. The password is case-sensitive. Spaces and special characters are not allowed. There is no default password.
4. Enter the password again in the **Confirm Password** field.
5. Click the **Apply** button. The switch displays the following message:

By clicking on Accept, the Factory Default Reset function will be Disabled on both the switch management software and the physical front panel ecoFriendly button. If you loose this password, ATI cannot recover it for you.

By Clicking on Cancel, the "Factory Default Reset" function will remain Enabled on both the switch management software and the physical front panel ecoFriendly button.

6. Click the **Accept** button.

7. Select **Save Settings to Flash** from the main menu to save your changes.

The eco-Friendly on the front panel of the switch is now disabled. Also, password protection is added to restoring the default settings to the switch.

Enabling the eco-Friendly Button and Default Settings

Perform the following procedure to enable the eco-Friendly button on the front panel of the switch or remove password protection from restoring the default feature settings:

1. Select **Tools > Reboot** from the main menu. Refer to Figure 228 on page 500. The Factory Default Reset/Reboot page is shown in Figure 229 on page 500.
2. Select **Enabled** from the Factory Default Reset menu. The Factory Default Reset/Reboot window displays a Password field.
3. Enter the password that locks the eco-Friendly button and default settings in “Disabling the eco-Friendly Button and Default Settings” on page 503. The password is case sensitive.
4. Click the **Apply** button.

The eco-Friendly button on the front panel of the switch is activated again and password protection is removed from restoring the default settings. The Reboot Type menu should now include options for returning the switch to the factory default values. Refer to “Restoring the Factory Default Values” on page 502.

5. Select **Save Settings to Flash** from the main menu to save your changes.

Pinging Network Devices

This section contains the procedure for instructing the switch to transmit ICMP Echo Requests to devices on the network. You might perform this procedure to test for active physical links between the switch and other network devices. This can be useful when troubleshooting network communications problems between devices. You can ping IPv4 or IPv6 devices.

ICMP Echo Requests are supported on ports of 802.1Q Tagged VLANs that are designated as Management VLANs, including the Default VLAN. They are not supported on Port-based VLANs.

Perform the following procedure to instruct the switch to issue ICMP Echo Requests from ports in 802.1 Q Tagged VLANs:

1. Select **Tools > Ping** from the main menu. Refer to Figure 230.




Figure 230. Ping Menu Selection

2. The Ping Test Settings window is shown in Figure 231.

Ping Test Settings

Destination IP Address: . . . IPv4
 IPv6

Timeout Value: (1-5) sec

Number of Ping Requests: (1-10) times

Figure 231. Ping Test Settings Window

3. Enter the fields in Table 84.

Table 84. Ping Test Settings Window

Field	Description
Destination IP Address	<p>Enter the IPv4 or IPv6 address of the remote network device. The format for an IPv4 address is shown here:</p> <p>nnn nnn nnn nnn</p> <p>Each N is a decimal number from 0 to 255.</p> <p>The format for an IPv6 address is shown here:</p> <p>xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx</p> <p>Each x is a hexadecimal digit representing four bits. The eight groups of digits have to be separated by colons. Groups where all four digits are '0' can be omitted. Leading '0's in groups can also be omitted.</p> <p>As an example, the following IPv6 addresses are equivalent:</p> <p>3710:421e:09a8:0000:0000:0000:00a4:1c50</p> <p>3710:421e:9a8::a4:1c50</p>
Timeout Value	Enter the length of time, in seconds, the switch waits for a response before assuming pings failed.
Number of Ping Requests	Enter the number of ping requests the switch is to transmit.

4. Click the **Start** button.
5. Click the **Show Ping Results** button to view the ping results. Table 85 defines the information.

Table 85. Ping Test Results Window

Field	Description
Destination IP Address	Displays the IPv4 or IPv6 address of the remote network device.
Pass	Displays the percentage of successful pings.
Average Time	Displays the average transit time (milliseconds) of the ping responses.

6. Click **Back to Ping Test** to return to the Ping Test Settings window.

Chapter 43

Loopback Detection

This chapter describes loopback detection in the following sections:

- “Configuring Loopback Detection” on page 508
- “Disabling Loopback Detection” on page 511

Configuring Loopback Detection

Loopback detection tests for shorts or crossovers on the strands of wires in the network cables on the copper ports. It automatically disables ports on which it detects loopback conditions to protect network operations from faulty cables.

Note

Loopback detection requires that spanning tree protocol be disabled. Refer to “Configuring STP and RSTP Global Settings” on page 138.

Perform the following procedure to activate and configure loopback detection:

1. Select **Bridge > Loopback Detection** from the main menu. Refer to Figure 232.




Figure 232. Loopback Detection Menu Selection

The Loopback Detection window is shown in Figure 233.

Loopback Detection

State Disabled ▾

Loopback Detection Global Settings

Interval 2 Sec(1-32767)

Recover Time: 60 Sec(0 or 60-1000000, 0 is Disabled)

Apply

Note: Disabling will turn off the function and return all values to default.

Port	Loopback Detection State	Loop Status	Action
All	Ignore ▾	-	Apply
1	Disabled ▾	Normal	Apply
2	Disabled ▾	Normal	Apply

Figure 233. Loopback Detection Window

2. Select **Enabled** from the State menu. You have to enable the feature to configure the settings.
3. To configure the settings in the Loopback Detection Global Settings section of the window, do the following:
 - a. Configure the fields in Table 86.

Table 86. Loopback Detection Global Settings

Field	Description
Interval	Enter the interval, in seconds, between tests. The range is 1 to 32767 seconds. The default is two seconds.
Recover Time	Enter the number of seconds ports remain disabled after a loopback condition is corrected. The range is 60 to 1000000 seconds. The default is 60 seconds. The default setting disables recover time. Ports that experience loopback events remain disabled until the feature is disabled, the Recover Time is changed to a different value, or the switch is reset.

- b. Click the **Apply** button.
4. To enable or disable loopback detection on the individual switch ports, do the following:
 - a. Select one of the following from the Loopback Detection State menu of the port in the table.
 - Enabled:** Enables the feature on the port. The switch disables the port, blocking it from forwarding traffic, if it detects a loopback condition.
 - Disabled:** Disables the feature, This is the default setting.
 - b. Click the **Apply** button.
 - c. Repeat this step to configure other ports.

The Loop Status column in the table displays the status of the test on the individual ports. Possible states are listed here:

- Normal:** The test is not detecting any problems with the copper cables.
- Loop:** The test disabled the port because it detected a loopback condition on the copper cable. The port remains disabled until the condition is resolved and the recovery timer expires.

5. Select **Save Settings to Flash** from the main menu to save your changes.

Disabling Loopback Detection

Perform the following procedure to disable loopback detection on the switch:

1. Select **Bridge > Loopback Detection** from the main menu. Refer to Figure 232 on page 508. The Loopback Detection window is shown in Figure 233 on page 508.
2. Select **Disabled** from the State menu.
3. Click the **Apply** button.
4. Select **Save Settings to Flash** from the main menu to save your changes.

Chapter 44

Switch Firmware

This chapter explains how to upgrade the management software on the switch in the following sections:

- ❑ “Management Software Overview” on page 514
- ❑ “Upgrading the Management Software with a Web Browser” on page 515
- ❑ “Upgrading the Management Software with TFTP” on page 518
- ❑ “Designating the Boot Management Software” on page 520

Management Software Overview

Allied Telesis may periodically release new versions of the management software for this product and post them on its web site for customers to download. Refer to the Allied Telesis web site for details.

The switch maintains two copies of the management software in flash memory. The versions are called Image1 and Image2. One of the images will be the active image and the other a backup. The switch uses the active image when it is booted or powered on.

When you download a new firmware image, the switch stores it as the inactive image. It then makes that image active and the other inactive. Here is an example. Assume that Image1 is active and Image2 is inactive. When the switch receives a new firmware image, it stores it as Image2 because that is the inactive image. After completing the download process, it makes Image2 the active image and Image 1 the inactive image.

If the switch is unable to initialize the active image or you believe there is a problem with it, you can manually instruct the switch to use the inactive image instead the next time the device is booted. Refer to “Designating the Boot Management Software” on page 520.

There are two ways to download new management software to the switch. One way is with your web browser on your management workstation, explained in “Upgrading the Management Software with a Web Browser” on page 515. The other is with a TFTP server, explained in “Upgrading the Management Software with TFTP” on page 518.

Upgrading the Management Software with a Web Browser

This section contains the procedure for upgrading the management software on the GS950 PS V2 Switch using HTTP on a web browser. Review the following guidelines:

- ❑ This procedure assumes you have already obtained the new management software from the Allied Telesis web site and stored it on your computer.
- ❑ The switch retains its current configuration when new management software is installed.
- ❑ The switch stores the new firmware as the inactive image. After downloading the new firmware, you have to designate it as the active image and reboot the switch before it will use the new firmware. Refer to “Management Software Overview” on page 514.



Caution

Do not power off the switch during the software upgrade. Interrupting the transfer may corrupt the firmware on the switch.

Note

The switch does not forward network traffic while it uploads the management software from your computer and writes it to flash memory. To minimize the disruption of the upgrade procedure to network operations, you should perform this procedure during periods of low traffic activity, such as during non-business hours.

Note

Do not close the web browser window or change to a different window until the entire procedure is finished. Interrupting the transfer may corrupt the file on the switch.

Perform the following procedure to upgrade the management software on the switch with your web browser and HTTP:

1. Select **Tools > Firmware Upgrade > via HTTP** from the main menu. Refer to Figure 234 on page 516.




Figure 234. Firmware Upgrade via HTTP Menu Selection

The Firmware Upgrade via HTTP window is shown in Figure 235.

Firmware Upgrade via HTTP

Next Boot Image ID: Image1 Image2

Running Image ID:

Image1 Version:

Image2 Version:

Firmware File: No file chosen

Note:

1. System will not reset automatically after burning image to flash.
2. Firmware will upgrade to a Not-Running image, please modify "Next Boot Image ID".

Figure 235. Firmware Upgrade via HTTP Window

The Next Boot Image ID and Running Image ID in the window are explained in “Designating the Boot Management Software” on page 520.

The Image1 Version and Image2 Version fields display the version numbers of the active and backup management software, respectively, on the switch. For background information, refer to “Management Software Overview” on page 514.

2. Click the **Choose File** button and locate the new management software file on your computer.
3. Click the **Upgrade** button to begin the upgrade process.

The switch needs to be restarted manually to download and verify the management software, and write it to flash memory. This entire process takes several minutes.

4. To resume managing the switch at the completion of the upgrade, start a new management session.

Upgrading the Management Software with TFTP

This section contains the procedure for upgrading the management software on the GS950 PS V2 Switch using TFTP. Review the following:

- ❑ This procedure assumes that you have already obtained the new management software from the Allied Telesis web site and stored it on your TFTP server.
- ❑ The switch retains its current configuration when new management software is installed.
- ❑ Start the TFTP server software before beginning the download procedure.
- ❑ The switch stores the new firmware as Image1 and renames its previous active firmware as Image2. For background information, refer to “Management Software Overview” on page 514.



Caution

Do not power off the switch during the software upgrade. Interrupting the transfer may corrupt the file on the switch.



Caution

The switch does not forward network traffic while it uploads the management software from your computer and writes it to flash memory. To minimize the disruption of the upgrade procedure to network operations, you should perform this procedure during periods of low traffic activity, such as during non-business hours.

Perform the following procedure to download new management software to the switch from a TFTP server:

1. Start the TFTP server on your network.
2. Select **Tools > Firmware Upgrade > via TFTP** from the main menu. Refer to Figure 236.




Figure 236. Firmware Upgrade via TFTP Menu Selection

The Firmware Upgrade via TFTP page is shown in Figure 237.

Firmware Upgrade via TFTP

Image1 Version: 1.00.005

Image2 Version: 1.00.011

TFTP Server IP: - - - IPv4 IPv6

Image File Name: (64 Characters Max.)

Retry Count:

Note:

1. System will not reset automatically after burning image to flash.
2. Firmware will upgrade to a Not-Running image, please modify "Next Boot Image ID".

Figure 237. Firmware Upgrade via TFTP Window

The Image1 Version and Image2 Version fields display the version numbers of the active and backup management software, respectively, on the switch. For background information, refer to “Management Software Overview” on page 514

3. Click either the **IPv4** or **IPv6** radio circle and enter the IP address of the TFTP server.
4. Enter the filename of the new management software on the TFTP server in the Image File Name field.
5. Enter In the Retry Count field the number of times the switch is to retry the firmware upgrade if it encounters a problem downloading the firmware. The range is 1 to 20.
6. Click the **Apply** button.

The switch needs to be restarted manually to download and verify the management software, and write it to flash memory. This entire process takes several minutes.

7. To resume managing the switch at the completion of the upgrade, start a new management session.

Designating the Boot Management Software

This procedure explains how to configure the switch to use the Image2 firmware as its management software the next time you reboot it. You might perform this procedure if there is a problem with the Image1 firmware. Refer to “Management Software Overview” on page 514 for background information.

Perform the following procedure to designate the Image2 firmware as the active management software the next time the switch boots:

1. Select **Tools > Firmware Upgrade > via HTTP** from the main menu. Refer to Figure 238.




Figure 238. Firmware Upgrade via HTTP Menu Selection

The Firmware Upgrade via HTTP window is shown in Figure 239.

Firmware Upgrade via HTTP

Next Boot Image ID: Image1 Image2

Running Image ID:

Image1 Version:

Image2 Version:

Firmware File: No file chosen

Note:

1. System will not reset automatically after burning image to flash.
2. Firmware will upgrade to a Not-Running image, please modify "Next Boot Image ID".

Figure 239. Firmware Upgrade via HTTP Window

The bottom portion of the window is explained in “Upgrading the Management Software with a Web Browser” on page 515.

2. Click the **Image2** radio button.
3. Click the **Apply** button.
4. Reboot the switch. Refer to “Rebooting the Switch” on page 500. The switch uses the Image2 firmware as its active management software.

Appendix A

Open Source Software

This product contains open source software. To view the copyright information, select **System** > **OSS Information** from the main menu. Refer to Figure 240.

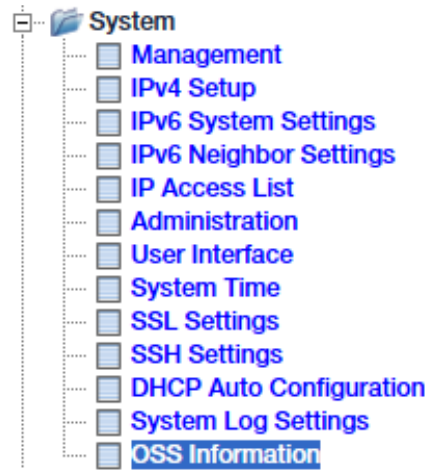


Figure 240. OSS Information Menu Selection

The OSS Information window is shown in Figure 241.



Figure 241. OSS Information Window

GPL2 (GNU General Public License, v.2)

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

< one line to give the program's name and a brief idea of what it does.>

Copyright (C) 19yy <name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) 19yy name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'. This is free software, and you are welcome to redistribute it under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program `Gnomovision' (which makes passes at compilers) written by James Hacker.

<signature of Ty Coon>, 1 April 1989

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

OpenSSH

The licences which components of this software fall under are as follows. First, we will summarize and say that all components are under a BSD

licence, or a licence more free than that.

OpenSSH contains no GPL code.

1. Copyright (c) 1995 Tatu Ylonen, Espoo, Finland All rights reserved

As far as I am concerned, the code I have written for this software can be used freely for any purpose. Any derived versions of this software must be clearly marked as such, and if the derived work is incompatible with the protocol description in the RFC file, it must be called by a name other than "ssh" or "Secure Shell".

[Tatu continues]

However, I am not implying to give any licenses to any patents or copyrights held by third parties, and the software includes parts that are not under my direct control. As far as I know, all included source code is used in accordance with the relevant license agreements and can be used freely for any purpose (the GNU license being the most restrictive); see below for details.

[However, none of that term is relevant at this point in time. All of these restrictively licenced software components which he talks about have been removed from OpenSSH, i.e.,

- RSA is no longer included, found in the OpenSSL library
- IDEA is no longer included, its use is deprecated
- DES is now external, in the OpenSSL library
- GMP is no longer used, and instead we call BN code from OpenSSL
- Zlib is now external, in a library
- The make-ssh-known-hosts script is no longer included
- TSS has been removed
- MD5 is now external, in the OpenSSL library
- RC4 support has been replaced with ARC4 support from OpenSSL
- Blowfish is now external, in the OpenSSL library

[The licence continues]

Note that any information and cryptographic algorithms used in this software are publicly available on the Internet and at any major bookstore, scientific library, and patent office worldwide. More information can be found e.g. at "<http://www.cs.hut.fi/crypto>".

The legal status of this program is some combination of all these permissions and restrictions. Use only at your own responsibility. You will be responsible for any legal consequences yourself; I am not

making any claims whether possessing or using this is legal or not in your country, and I am not taking any responsibility on your behalf.

NO WARRANTY

BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

2. The 32-bit CRC compensation attack detector in deattack.c was contributed by CORE SDI S.A. under a BSD-style license.

Cryptographic attack detector for ssh - source code

Copyright (c) 1998 CORE SDI S.A., Buenos Aires, Argentina.

All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that this copyright notice is retained.

THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES ARE DISCLAIMED. IN NO EVENT SHALL CORE SDI S.A. BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY OR CONSEQUENTIAL DAMAGES RESULTING FROM THE USE OR MISUSE OF THIS SOFTWARE.

Ariel Futoransky futo@core-sdi.com

<http://www.core-sdi.com>

3. ssh-keyscan was contributed by David Mazieres under a BSD-style license.

Copyright 1995, 1996 by David Mazieres dm@lcs.mit.edu.

Modification and redistribution in source and binary forms is permitted provided that due credit is given to the author and the OpenBSD project by leaving this copyright notice intact.

4. The Rijndael implementation by Vincent Rijmen, Antoon Bosselaers and Paulo Barreto is in the public domain and distributed with the following license:

@version 3.0 (December 2000)

Optimised ANSI C code for the Rijndael cipher (now AES)

@author Vincent Rijmen vincent.rijmen@esat.kuleuven.ac.be

@author Antoon Bosselaers antoon.bosselaers@esat.kuleuven.ac.be

@author Paulo Barreto paulo.barreto@terra.com.br

This code is hereby placed in the public domain.

THIS SOFTWARE IS PROVIDED BY THE AUTHORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

5. One component of the ssh source code is under a 3-clause BSD license, held by the University of California, since we pulled these parts from original Berkeley code.

Copyright (c) 1983, 1990, 1992, 1993, 1995

The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

6. Remaining components of the software are provided under a standard 2-term BSD licence with the following names as copyright holders:

Markus Friedl

Theo de Raadt

Niels Provos

Dug Song

Aaron Campbell

Damien Miller

Kevin Steves

Daniel Kouril

Wesley Griffin

Per Allansson

Nils Nordman

Simon Wilkinson

Portable OpenSSH additionally includes code from the following copyright holders, also under the 2-term BSD license:

Ben Lindstrom

Tim Rice

Andre Lucas

Chris Adams

Corinna Vinschen

Cray Inc.

Denis Parker

Gert Doering

Jakob Schlyter

Jason Downs

Juha Yrj Michael Stone

Networks Associates Technology, Inc.

Solar Designer

Todd C. Miller

Wayne Schroeder

William Jones

Darren Tucker

Sun Microsystems

The SCO Group

Daniel Walsh

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1.** Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2.** Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY

AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

7. Portable OpenSSH contains the following additional licenses:

a. md5crypt.c, md5crypt.h

"THE BEER-WARE LICENSE" (Revision 42): wrote this file. As long as you retain this notice you can do whatever you want with this stuff. If we meet some day, and you think this stuff is worth it, you can buy me a beer in return. Poul-Henning Kamp

b. snprintf replacement

Copyright Patrick Powell 1995

This code is based on code written by Patrick Powell (papowell@astart.com) It may be used for any purpose as long as this notice remains intact on all source code distributions

c. Compatibility code (openbsd-compat)

Apart from the previously mentioned licenses, various pieces of code in the openbsd-compat/ subdirectory are licensed as follows:

Some code is licensed under a 3-term BSD license, to the following copyright holders:

Todd C. Miller

Theo de Raadt

Damien Miller

Eric P. Allman

The Regents of the University of California

Constantin S. Svintsoff

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Some code is licensed under an ISC-style license, to the following copyright holders:

Internet Software Consortium.

Todd C. Miller

Reyk Floeter

Chad Mynhier

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND TODD C. MILLER DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL TODD C. MILLER BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Some code is licensed under a MIT-style license to the following copyright holders:

Free Software Foundation, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the 'Software'), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, distribute with modifications, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED 'AS IS', WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE ABOVE COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name(s) of the above copyright holders shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization.

\$OpenBSD: LICENCE,v 1.19 2004/08/30 09:18:08 markus Exp \$

OpenSSL

LICENSE ISSUES

=====

The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. See below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.

OpenSSL License

=====

Copyright (c) 1998-2004 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"
4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment:

*This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

=====

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com) All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

"This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)"

The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related :-).

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement::

"This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The licence and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.]

