

HPE ArubaOS-Switch Management and Configuration Guide for YA/YB.16.03



Part Number: 5200-2925a
Published: January 2017
Edition: 2

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel[®], Itanium[®], Pentium[®], Intel Inside[®], and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft[®] and Windows[®] are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe[®] and Acrobat[®] are trademarks of Adobe Systems Incorporated.

Java[®] and Oracle[®] are registered trademarks of Oracle and/or its affiliates.

UNIX[®] is a registered trademark of The Open Group.

Chapter 1 About this document	18
Chapter 2 Time Protocols	19
General steps for running a time protocol on the switch.....	19
TimeP time synchronization.....	19
SNTP time synchronization.....	19
Selecting a time synchronization protocol.....	20
Disabling time synchronization.....	20
SNTP: Selecting and configuring.....	20
Viewing and configuring SNTP (Menu).....	21
Viewing and configuring SNTP (CLI).....	23
Configuring (enabling or disabling) the SNTP mode.....	24
TimeP: Selecting and configuring.....	29
Viewing, enabling, and modifying the TimeP protocol (Menu).....	30
Viewing the current TimeP configuration (CLI).....	31
Configuring (enabling or disabling) the TimeP mode.....	32
SNTP unicast time polling with multiple SNTP servers.....	36
Displaying all SNTP server addresses configured on the switch (CLI).....	36
Adding and deleting SNTP server addresses.....	36
Adding addresses.....	36
Deleting addresses.....	37
Operating with multiple SNTP server addresses configured (Menu).....	37
SNTP messages in the Event Log.....	37
Network Time Protocol (NTP).....	37
Commands.....	38
timesync Command.....	38
timesync ntp.....	38
ntp.....	38
[no] ntp.....	39
ntp enable.....	39
ntp authentication.....	40
ntp authentication key-id.....	41
ntp max-association.....	41
ntp server.....	42
ntp server key-id.....	44
ntp ipv6-multicast.....	44
debug ntp.....	45
ntp trap.....	45
show ntp statistics.....	46
show ntp status.....	47
show ntp associations.....	47
show ntp authentication.....	48
Validation rules.....	49
Event log messages.....	51
Chapter 3 Port Status and Configuration	53
Viewing port status and configuring port parameters.....	53
Connecting transceivers to fixed-configuration devices.....	53

Viewing port configuration (Menu).....	53
Configuring ports (Menu).....	56
Viewing port status and configuration (CLI).....	57
Dynamically updating the show interfaces command (CLI/Menu).....	57
Customizing the show interfaces command (CLI).....	58
Error messages associated with the show interfaces command.....	59
Viewing port utilization statistics (CLI).....	60
Operating notes for viewing port utilization statistics.....	60
Viewing transceiver status (CLI).....	61
Operating Notes.....	61
Enabling or disabling ports and configuring port mode (CLI).....	62
Enabling or disabling flow control (CLI).....	63
Configuring a broadcast limit.....	65
Broadcast-limit.....	65
Port shutdown with broadcast storm.....	65
Viewing broadcast storm.....	66
SNMP MIB.....	67
Configuring auto-MDIX.....	69
Manual override.....	69
Configuring auto-MDIX (CLI).....	70
Using friendly (optional) port names.....	71
Configuring and operating rules for friendly port names.....	71
Configuring friendly port names (CLI).....	72
Configuring a single port name (CLI).....	72
Configuring the same name for multiple ports (CLI).....	72
Displaying friendly port names with other port data (CLI).....	73
Listing all ports or selected ports with their friendly port names (CLI).....	73
Including friendly port names in per-port statistics listings (CLI).....	74
Searching the configuration for ports with friendly port names (CLI).....	75
Uni-directional link detection (UDLD).....	76
Configuring UDLD.....	76
Configuring uni-directional link detection (UDLD) (CLI).....	77
Enabling UDLD (CLI).....	77
Changing the keepalive interval (CLI).....	78
Changing the keepalive retries (CLI).....	78
Configuring UDLD for tagged ports.....	78
Viewing UDLD information (CLI).....	79
Viewing summary information on all UDLD-enabled ports (CLI).....	79
Viewing detailed UDLD information for specific ports (CLI).....	79
Clearing UDLD statistics (CLI).....	80
Chapter 4 Power Over Ethernet (PoE/PoE+) Operation.....	81
Introduction to PoE.....	81
PoE terminology.....	81
Planning and implementing a PoE configuration.....	81
Power requirements.....	81
Assigning PoE ports to VLANs.....	82
Applying security features to PoE configurations.....	82
Assigning priority policies to PoE traffic.....	82
PoE Event Log messages.....	82
About PoE operation.....	82
Configuration options.....	83
PD support.....	83
Power priority operation.....	84
Configuring PoE operation.....	84

Disabling or re-enabling PoE port operation.....	84
Enabling support for pre-standard devices.....	84
Configuring the PoE port priority.....	85
Controlling PoE allocation.....	86
Manually configuring PoE power levels.....	87
Changing the threshold for generating a power notice.....	89
Cycling power on a port.....	90
PoE/PoE+ allocation using LLDP information.....	90
LLDP with PoE.....	90
Enabling or disabling ports for allocating power using LLDP.....	91
Enabling PoE detection via LLDP TLV advertisement.....	91
LLDP with PoE+.....	91
Overview.....	91
PoE allocation.....	91
Initiating advertisement of PoE+ TLVs.....	93
Viewing PoE when using LLDP information.....	93
Operation note.....	95
Viewing the global PoE power status of the switch.....	95
Viewing PoE status on all ports.....	96
Viewing the PoE status on specific ports.....	98

Chapter 5 Port Trunking..... 101

Overview of port trunking.....	101
Port connections and configuration.....	101
Port trunk features and operation.....	102
Fault tolerance.....	102
Trunk configuration methods.....	102
Dynamic LACP trunk.....	102
Static trunk.....	102
Viewing and configuring a static trunk group (Menu).....	106
Viewing and configuring port trunk groups (CLI).....	107
Viewing static trunk type and group for all ports or for selected ports.....	107
Viewing static LACP and dynamic LACP trunk data.....	108
Dynamic LACP Standby Links.....	109
Configuring a static trunk or static LACP trunk group.....	109
Removing ports from a static trunk group.....	110
Enabling a dynamic LACP trunk group.....	110
Removing ports from a dynamic LACP trunk group.....	111
Viewing existing port trunk groups (WebAgent).....	112
Trunk group operation using LACP.....	112
Default port operation.....	114
LACP notes and restrictions.....	115
802.1X (Port-based access control) configured on a port.....	115
Port security configured on a port.....	116
Changing trunking methods.....	116
Static LACP trunks.....	116
Dynamic LACP trunks.....	116
VLANs and dynamic LACP.....	116
Blocked ports with older devices.....	116
Spanning Tree and IGMP.....	117
Half-duplex, different port speeds, or both not allowed in LACP trunks.....	117
Dynamic/static LACP interoperation.....	118
Trunk group operation using the "trunk" option.....	118
How the switch lists trunk data.....	118
Outbound traffic distribution across trunked links.....	118

Chapter 6 Port Traffic Controls	120
VLAN-based rate-limiting	120
ICMP rate-limiting	120
Guidelines for configuring ICMP rate-limiting	121
Configuring ICMP rate-limiting	121
Using both ICMP rate-limiting and all-traffic rate-limiting on the same interface	122
Viewing the current ICMP rate-limit configuration	123
Operating notes for ICMP rate-limiting	123
Notes on testing ICMP rate-limiting	124
ICMP rate-limiting trap and Event Log messages	124
Determining the switch port number used in ICMP port reset commands	125
Configuring inbound rate-limiting for broadcast and multicast traffic	126
Operating Notes	127
Jumbo frames	128
Operating rules	128
Jumbo traffic-handling	128
Configuring jumbo frame operation	129
Overview	129
Viewing the current jumbo configuration	130
Enabling or disabling jumbo traffic on a VLAN	131
Configuring a maximum frame size	131
Configuring IP MTU	132
SNMP implementation	132
Displaying the maximum frame size	132
Operating notes for maximum frame size	132
Troubleshooting	133
A VLAN is configured to allow jumbo frames, but one or more ports drops all inbound jumbo frames	133
A non-jumbo port is generating "Excessive undersize/giant frames" messages in the Event Log	133
Chapter 7 Fault-Finder port-level link-flap	134
Overview	134
Fault-finder link-flap	134
Show fault-finder link-flap	136
Event Log	137
Restrictions	137
Chapter 8 Configuring for Network Management Applications	138
Using SNMP tools to manage the switch	138
SNMP management features	138
SNMPv1 and v2c access to the switch	139
SNMPv3 access to the switch	139
Enabling and disabling switch for access from SNMPv3 agents	140
Enabling or disabling restrictions to access from only SNMPv3 agents	140
Enabling or disabling restrictions from all non-SNMPv3 agents to read-only access	140
Viewing the operating status of SNMPv3	140
Viewing status of message reception of non-SNMPv3 messages	140
Viewing status of write messages of non-SNMPv3 messages	140
Enabling SNMPv3	140
SNMPv3 users	141
Group access levels	144

SNMPv3 communities.....	145
Viewing and configuring non-version-3 SNMP communities (Menu).....	146
Listing community names and values (CLI).....	147
SNMP notifications.....	148
Supported Notifications.....	148
General steps for configuring SNMP notifications.....	149
SNMPv1 and SNMPv2c Traps.....	149
SNMP trap receivers.....	149
SNMPv2c informs.....	151
Configuring SNMPv3 notifications (CLI).....	152
Network security notifications.....	154
Enabling Link-Change Traps (CLI).....	156
Source IP address for SNMP notifications.....	157
Viewing SNMP notification configuration (CLI).....	159
Advanced management: RMON.....	159
CLI-configured sFlow with multiple instances.....	160
Configuring sFlow (CLI).....	160
Viewing sFlow Configuration and Status (CLI).....	160
Configuring UDLD Verify before forwarding.....	162
UDLD time delay.....	162
Restrictions.....	163
UDLD configuration commands.....	163
Show commands.....	164
RMON generated when user changes UDLD mode.....	164
LLDP.....	164
General LLDP operation.....	165
LLDP-MED.....	165
Packet boundaries in a network topology.....	165
LLDP operation configuration options.....	165
Enable or disable LLDP on the switch.....	165
Enable or disable LLDP-MED.....	166
Change the frequency of LLDP packet transmission to neighbor devices.....	166
Change the Time-To-Live for LLDP packets sent to neighbors.....	166
Transmit and receive mode.....	166
SNMP notification.....	166
Per-port (outbound) data options.....	166
Remote management address.....	168
Debug logging.....	168
Options for reading LLDP information collected by the switch.....	168
LLDP and LLDP-MED standards compatibility.....	168
LLDP operating rules.....	168
Port trunking.....	169
IP address advertisements.....	169
Spanning-tree blocking.....	169
802.1X blocking.....	169
Configuring LLDP operation.....	169
Displaying the global LLDP, port admin, and SNMP notification status (CLI).....	169
Configuring Global LLDP Packet Controls.....	171
Configuring SNMP notification support.....	174
Configuring per-port transmit and receive modes (CLI).....	175
Basic LLDP per-port advertisement content.....	175
Support for port speed and duplex advertisements.....	177
Port VLAN ID TLV support on LLDP.....	177
Configuring the VLAN ID TLV.....	177
Viewing the TLVs advertised.....	178
SNMP support.....	179
LLDP-MED (media-endpoint-discovery).....	179

LLDP-MED endpoint support.....	180
LLDP-MED endpoint device classes.....	181
LLDP-MED operational support.....	181
LLDP-MED fast start control.....	181
Advertising device capability, network policy, PoE status and location data.....	182
Location data for LLDP-MED devices.....	184
Viewing switch information available for outbound advertisements.....	188
Displaying the current port speed and duplex configuration on a switch port.....	190
Viewing advertisements currently in the neighbors MIB.....	190
Displaying LLDP statistics.....	191
LLDP Operating Notes.....	193
Neighbor maximum.....	193
LLDP packet forwarding.....	193
One IP address advertisement per port.....	193
802.1Q VLAN Information.....	194
Effect of 802.1X Operation.....	194
Neighbor data can remain in the neighbor database after the neighbor is disconnected.....	194
Mandatory TLVs.....	194
LLDP and CDP data management.....	194
LLDP and CDP neighbor data.....	194
CDP operation and commands.....	195
Viewing the current CDP configuration of the switch.....	196
Viewing the current CDP neighbors table of the switch.....	196
Enabling and Disabling CDP Operation.....	197
Enabling or disabling CDP operation on individual ports.....	197
Filtering CDP information.....	197
Configuring the switch to filter untagged traffic.....	198
Displaying the configuration.....	198
Filtering PVID mismatch log messages.....	199
Generic header ID in configuration file.....	199
Introduction.....	199
Add-Ignore-Tag option.....	199
Configuration commands for the add-ignore-tag option.....	200
Show logging commands for the add-ignore-tag option.....	200
Exclusions.....	201
Chapter 9 Captive Portal for ClearPass.....	202
Requirements.....	202
Best Practices.....	202
Limitations.....	203
Features.....	203
High Availability.....	203
Load balancing and redundancy.....	203
Captive Portal when disabled.....	203
Disabling Captive Portal.....	203
Configuring Captive Portal on CPPM.....	204
Import the HP RADIUS dictionary.....	204
Create enforcement profiles.....	204
Create a ClearPass guest self-registration.....	205
Configure the login delay.....	206
Configuring the switch.....	206
Configure the URL key.....	207
Configuring a certificate for Captive Portal usage.....	207
Display Captive Portal configuration.....	208

Show certificate information.....	208
Troubleshooting.....	208
Event Timestamp not working.....	208
Cannot enable Captive Portal.....	209
Unable to enable feature.....	209
Authenticated user redirected to login page.....	209
Unable to configure a URL hash key.....	210
authentication command.....	210
show command.....	211
Debug command.....	211

Chapter 10 ZTP with AirWave Network Management..... 213

Requirements.....	213
Best Practices.....	213
Limitations.....	213
Switch configuration.....	213
Configure AirWave details in DHCP (preferred method).....	214
Configure AirWave details in DHCP (alternate method).....	219
Zero Touch Provisioning.....	226
Auto-configuration using ZTP.....	226
Disabling ZTP.....	227
Image Upgrade.....	227
Configure a switch using the CLI.....	227
Stacking and chassis switches.....	228
Troubleshooting.....	228
View AMP server messages.....	228
Validation Rules.....	228
View configuration details.....	229
amp-server.....	229
debug ztp.....	230

Chapter 11 Auto configuration upon Aruba AP detection..... 231

Auto device detection and configuration.....	231
Requirements.....	231
Limitations.....	231
Feature Interactions.....	231
Profile Manager and 802.1X.....	231
Profile Manager and LMA/WMA/MAC-AUTH.....	232
Profile manager and Private VLANs.....	232
Creating a profile and associate a device type.....	232
device-profile name.....	232
device-profile type.....	234
Rogue AP Isolation.....	234
Limitations.....	235
Feature Interactions.....	235
MAC lockout and lockdown.....	235
LMA/WMA/802.1X/Port-Security.....	236
L3 MAC.....	236
Using the Rogue AP Isolation feature.....	236
rogue-ap-isolation.....	237
rogue-ap-isolation action.....	238
rogue-ap-isolation whitelist.....	238
clear rogue-ap-isolation.....	238
Troubleshooting.....	239

Dynamic configuration not displayed when using “show running-config”	239
Switch does not detect the rogue AP TLVs	239
The <code>show run</code> command displays non-numerical value for untagged-vlan	240
Show commands	240
Validation Rules	241

Chapter 12 Link Aggregation Control Protocol—Multi-Active Detection (LACP-MAD)..... 243

LACP-MAD commands	243
Configuration command	243
show commands	243
clear command	243
LACP-MAD overview	243

Chapter 13 Scalability IP Address VLAN and Routing Maximum Values..... 245

Chapter 14 File Transfers..... 247

Overview	247
Downloading switch software	247
General software download rules	247
Using TFTP to download software from a server	247
Downloading from a server to primary flash using TFTP (Menu)	248
Troubleshooting TFTP download failures	249
Downloading from a server to flash using TFTP (CLI)	250
Using SCP and SFTP	251
Enabling SCP and SFTP	252
Disabling TFTP and auto-TFTP for enhanced security	252
Enabling SSH V2 (required for SFTP)	254
Authentication	254
SCP/SFTP operating notes	255
Troubleshooting SSH, SFTP, and SCP operations	256
Using Xmodem to download switch software from a PC or UNIX workstation	257
Downloading to primary flash using Xmodem (Menu)	257
Downloading to primary or secondary flash using Xmodem and a terminal emulator (CLI)	258
Switch-to-switch download	259
Switch-to-switch download to primary flash (Menu)	259
Downloading the OS from another switch (CLI)	260
Using AirWave to update switch software	260
Using IMC to update switch software	261
Copying software images	261
TFTP: Copying a software image to a remote host (CLI)	261
Xmodem: Copying a software image from the switch to a serially connected PC or UNIX workstation (CLI)	261
Transferring switch configurations	261
TFTP: Copying a configuration file to a remote host (CLI)	262
TFTP: Copying a configuration file from a remote host (CLI)	262
TFTP: Copying a customized command file to a switch (CLI)	262
Xmodem: Copying a configuration file to a serially connected PC or UNIX workstation (CLI)	263
Xmodem: Copying a configuration file from a serially connected PC or UNIX workstation (CLI)	264

Copying diagnostic data to a remote host, PC or UNIX workstation	265
Copying command output to a destination device (CLI).....	265
Copying Event Log output to a destination device (CLI).....	266
Copying crash data content to a destination device (CLI).....	266
Chapter 15 Monitoring and Analyzing Switch Operation.....	268
Overview.....	268
Accessing port and trunk group statistics.....	268
show interfaces.....	268
Reset port counters.....	268
clear statistics.....	269
Accessing port and trunk statistics (Menu).....	270
MAC address tables.....	270
MAC address views and searches.....	270
show mac-address.....	270
Using the menu to view and search MAC addresses.....	271
Finding the port connection for a specific device on a VLAN.....	272
Viewing and searching port-level MAC addresses.....	272
Determining whether a specific device is connected to the selected port.....	273
MSTP data.....	273
show spanning-tree.....	273
IP IGMP status.....	274
show ip igmp.....	274
VLAN information.....	275
show vlan.....	276
Configuring a source switch in a local mirroring session.....	277
Selecting all traffic on a port interface for mirroring according to traffic direction.....	278
Viewing all mirroring sessions configured on the switch.....	279
Viewing the mirroring configuration for a specific session.....	280
Using the Menu to configure local mirroring.....	281
Menu and WebAgent limits.....	281
High-level overview of the mirror configuration process.....	281
Determine the mirroring session and destination.....	281
For a local mirroring session.....	281
Classifier-based mirroring configuration.....	281
Classifier-based mirroring restrictions.....	283
Mirroring configuration examples.....	284
Maximum supported frame size.....	285
Enabling jumbo frames to increase the mirroring path MTU.....	285
Effect of downstream VLAN tagging on untagged, mirrored traffic.....	286
Operating notes for traffic mirroring.....	287
Troubleshooting traffic mirroring.....	288
Interface monitoring features.....	289
Configuring port and static trunk monitoring (Menu).....	289
Configuring port and static trunk monitoring (CLI).....	289
Displaying the monitoring configuration.....	289
Configuring the monitor port.....	290
Selecting or removing monitoring source interfaces.....	290
Chapter 16 Troubleshooting.....	292
Overview.....	292
Troubleshooting approaches.....	292
Browser or Telnet access problems.....	293
Cannot access the WebAgent.....	293

Cannot Telnet into the switch console from a station on the network.....	293
Unusual network activity.....	294
General problems.....	294
The network runs slow; processes fail; users cannot access servers or other devices.....	294
Duplicate IP addresses.....	294
Duplicate IP addresses in a DHCP network.....	294
The switch has been configured for DHCP/Bootp operation, but has not received a DHCP or Bootp reply.....	294
802.1Q Prioritization problems.....	295
Ports configured for non-default prioritization (level 1 to 7) are not performing the specified action.....	295
Addressing ACL problems.....	295
ACLs are properly configured and assigned to VLANs, but the switch is not using the ACLs to filter IP layer 3 packets.....	295
The switch does not allow management access from a device on the same VLAN.....	296
Error (Invalid input) when entering an IP address.....	296
Apparent failure to log all "deny" matches.....	296
The switch does not allow any routed access from a specific host, group of hosts, or subnet.....	296
The switch is not performing routing functions on a VLAN.....	296
Routing through a gateway on the switch fails.....	297
IGMP-related problems.....	298
IP multicast (IGMP) traffic that is directed by IGMP does not reach IGMP hosts or a multicast router connected to a port.....	298
IP multicast traffic floods out all ports; IGMP does not appear to filter traffic.....	298
LACP-related problems.....	298
Unable to enable LACP on a port with the <code>interface <port-number> lacp</code> command.....	298
Port-based access control (802.1X)-related problems.....	298
The switch does not receive a response to RADIUS authentication requests.....	299
The switch does not authenticate a client even though the RADIUS server is properly configured and providing a response to the authentication request.....	299
During RADIUS-authenticated client sessions, access to a VLAN on the port used for the client sessions is lost.....	299
The switch appears to be properly configured as a supplicant, but cannot gain access to the intended authenticator port on the switch to which it is connected.....	299
The supplicant statistics listing shows multiple ports with the same authenticator MAC address.....	299
The <code>show port-access authenticator <port-list></code> command shows one or more ports remain open after they have been configured with <code>control unauthorized</code>	299
RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch.....	300
The authorized MAC address on a port that is configured for both 802.1X and port security either changes or is re-acquired after execution of <code>aaa port-access authenticator <port-list> initialize</code>	300
A trunked port configured for 802.1X is blocked.....	300
QoS-related problems.....	300
Loss of communication when using VLAN-tagged traffic.....	301
Radius-related problems.....	301
The switch does not receive a response to RADIUS authentication requests.....	301
RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch.....	301
MSTP and fast-uplink problems.....	302
Broadcast storms appearing in the network.....	302
STP blocks a link in a VLAN even though there are no redundant links in that VLAN.....	302

Fast-uplink troubleshooting.....	302
SSH-related problems.....	302
Switch access refused to a client.....	302
Executing IP SSH does not enable SSH on the switch.....	302
Switch does not detect a client's public key that does appear in the switch's public key file (<code>show ip client-public-key</code>).....	303
An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages.....	303
Client ceases to respond ("hangs") during connection phase.....	303
TACACS-related problems.....	303
Event Log.....	303
All users are locked out of access to the switch.....	303
No communication between the switch and the TACACS+ server application.....	304
Access is denied even though the username/password pair is correct.....	304
Unknown users allowed to login to the switch.....	304
System allows fewer login attempts than specified in the switch configuration.....	304
TimeP, SNTP, or Gateway problems.....	304
The switch cannot find the time server or the configured gateway.....	304
VLAN-related problems.....	304
Monitor port.....	304
None of the devices assigned to one or more VLANs on an 802.1Q-compliant switch are being recognized.....	305
Link configured for multiple VLANs does not support traffic for one or more VLANs.....	305
Duplicate MAC addresses across VLANs.....	305
Fan failure.....	306
Viewing transceiver information.....	306
Viewing information about transceivers (CLI).....	307
MIB support.....	307
Viewing transceiver information.....	307
Information displayed with the detail parameter.....	309
Viewing transceiver information for copper transceivers with VCT support.....	312
Testing the Cable.....	312
Using the Event Log for troubleshooting switch problems.....	314
Event Log entries.....	315
Using the Menu.....	323
Using the CLI.....	324
Clearing Event Log entries.....	325
Turning event numbering on.....	325
Using log throttling to reduce duplicate Event Log and SNMP messages.....	325
Log throttle periods.....	325
Example: of event counter operation.....	327
Reporting information about changes to the running configuration.....	327
Debug/syslog operation.....	328
Debug/syslog messaging.....	328
Hostname in syslog messages.....	328
Logging origin-id.....	328
Viewing the identification of the syslog message sender.....	330
SNMP MIB.....	332
Debug/syslog destination devices.....	332
Debug/syslog configuration commands.....	333
Configuring debug/syslog operation.....	335
Viewing a debug/syslog configuration.....	336
Debug command.....	338
Debug messages.....	338
Debug destinations.....	340
Logging command.....	341

Configuring a syslog server.....	341
Adding a description for a Syslog server.....	348
Adding a priority description.....	349
Configuring the severity level for Event Log messages sent to a syslog server.....	349
Configuring the system module used to select the Event Log messages sent to a syslog server.....	350
Operating notes for debug and Syslog.....	350
Diagnostic tools.....	351
Port auto-negotiation.....	351
Ping and link tests.....	351
Ping test.....	352
Link test.....	352
Executing ping or link tests (WebAgent).....	352
Testing the path between the switch and another device on an IP network.....	352
Issuing single or multiple link tests.....	354
Tracing the route from the switch to a host address.....	354
Halting an ongoing traceroute search.....	355
A low maxttl causes traceroute to halt before reaching the destination address.....	356
If a network condition prevents traceroute from reaching the destination.....	356
Viewing switch configuration and operation.....	357
Viewing the startup or running configuration file.....	357
Viewing the configuration file (WebAgent).....	357
Viewing a summary of switch operational data.....	357
Saving show tech command output to a text file.....	358
Viewing more information on switch operation.....	359
Searching for text using pattern matching with show command.....	360
Displaying the information you need to diagnose problems.....	362
Restoring the factory-default configuration.....	363
Resetting to the factory-default configuration.....	363
Using the CLI.....	363
Using Clear/Reset.....	363
Restoring a flash image.....	364
Recovering from an empty or corrupted flash state.....	364
DNS resolver.....	365
Basic operation.....	365
Configuring and using DNS resolution with DNS-compatible commands.....	366
Configuring a DNS entry.....	367
Using DNS names with ping and traceroute: Example:.....	367
Viewing the current DNS configuration.....	369
Operating notes.....	369
Event Log messages.....	370
Chapter 17 MAC Address Management.....	371
Overview.....	371
Determining MAC addresses.....	371
Viewing the MAC addresses of connected devices.....	371
Viewing the switch's MAC address assignments for VLANs configured on the switch.....	372
Viewing the port and VLAN MAC addresses.....	373
Chapter 18 Power-Saving Features.....	375
Configuring the savepower LED option.....	375
Configuring the savepower port-low-pwr option.....	375

Chapter 19 Job Scheduler	377
Job Scheduler.....	377
Commands.....	377
Job at delay enable disable	377
Show job.....	378
Show job <Name>.....	379
Chapter 20 Virtual Technician	380
Cisco Discovery Protocol (CDP).....	380
Show cdp traffic.....	380
Clear cdp counters.....	380
Enable/Disable debug tracing for MOCANA code.....	381
Debug security	381
User diagnostic crash via Front Panel Security (FPS) button.....	381
Front panel security password-clear.....	381
Front-panel-security diagnostic-reset.....	382
[no] front-panel-security diagnostic-reset.....	382
Front-panel-security diagnostic-reset clear-button.....	383
[No] front-panel-security diagnostic-reset clear-button.....	383
Show front-panel-security.....	383
Diagnostic table.....	384
Validation rules.....	384
FPS Error Log.....	384
User initiated diagnostic crash via the serial console.....	385
Front-panel-security diagnostic-reset serial-console.....	385
[No] front-panel-security diagnostic-reset serial-console.....	386
Serial console error messages.....	386
Chapter 21 Easing Wired/Wireless Deployment feature integration	387
Overview.....	387
Configuration commands.....	387
allow-jumbo-frames.....	387
Validation rules.....	388
Default AP Profile.....	388
device-profile.....	388
Associating a device with a profile.....	389
device-profile type.....	389
Configuring the rogue-ap-isolation command.....	390
rogue-ap-isolation.....	390
Show commands.....	390
show device-profile.....	391
show command device-profile status	392
Show rogue-ap-isolation.....	392
Chapter 22 Local user roles	394
Overview.....	394
Captive-portal commands.....	396
Overview.....	396
[no] aaa authentication captive-portal profile.....	396
Validation rules.....	397
Policy commands.....	397

Overview.....	397
policy user.....	398
[no] policy user.....	398
policy resequence.....	398
Commands in the policy-user context.....	399
(policy-user)# class.....	399
User role configuration.....	399
aaa authorization user-role.....	399
Error log.....	401
captive-portal-profile.....	402
policy.....	402
reauth-period.....	402
Validation rules.....	403
VLAN commands.....	403
vlan-id.....	403
vlan-name.....	403
VLAN range commands.....	404
Applying a UDR.....	405
aaa port-access local-mac apply user-role.....	405
Show commands.....	405
show captive-portal profile.....	405
show user-role.....	406
show port-access clients.....	407
Chapter 23 Port QoS Trust Mode.....	409
Overview.....	409
Configuration commands.....	409
qos trust.....	409
qos dscp-map.....	410
Show commands.....	410
show qos trust.....	410
Validation rules.....	411
Chapter 24 Net-destination and Net-service.....	413
Net-service Overview.....	413
net-service [tcp udp port].....	413
Net-destination overview.....	414
net-destination host [position network].....	415
show net-destination.....	415
Chapter 25 Websites.....	416
Chapter 26 Support and other resources.....	417
Accessing Hewlett Packard Enterprise Support.....	417
Accessing updates.....	417
Customer self repair.....	417
Remote support.....	418
Warranty information.....	418
Regulatory information.....	418
Documentation feedback.....	419

Remote Device Deployment (TR-069)	420
Introduction.....	420
Advantages of TR-069.....	421
Zero-touch configuration process.....	421
Zero-touch configuration setup and execution.....	424
CLI commands.....	424
Configuration setup.....	424
ACS password configuration.....	425
When encrypt-credentials is off.....	425
When encrypt-credentials is on.....	425
ACS URL configuration.....	426
ACS username configuration.....	426
CPE configuration.....	426
CPE password configuration.....	426
When encrypt-credentials is on.....	427
When encrypt-credentials is off.....	427
CPE username configuration.....	427
Enable/disable CWMP.....	427
Show commands.....	427
CWMP configuration and status query.....	428
Event logging.....	429
System logging.....	429
Status/control commands.....	430
 Glossary	 431

This switch software guide is intended for network administrators and support personnel, and applies to the switch models listed on this page unless otherwise noted. This guide does not provide information about upgrading or replacing switch hardware.

Applicable Products

Aruba 2530 Switch Series (JL070A, J9772A, J9773A, J9774A, J9775A, J9776A, J9777A, J9778A, J9779A, J9780A, J9781A, J9782A, J9783A, J9853A, J9854A, J9855A, J9856A)



For successful time protocol setup and specific configuration details, you may need to contact your system administrator regarding your local configuration.

General steps for running a time protocol on the switch

Using time synchronization ensures a uniform time among interoperating devices. This helps you to manage and troubleshoot switch operation by attaching meaningful time data to event and error messages.

The switch offers TimeP, SNTP (Simple Network Time Protocol), NTP, and a `timesync` command for changing the time protocol selection (or turning off time protocol operation).



Although you can create and save configurations for all time protocols without conflicts, the switch allows only one active time protocol at any time.

In the factory-default configuration, time synchronization is disabled by default.



Because the Aruba 2530 Switch Series does not contain an RTC (real time clock) chip, Hewlett Packard Enterprise recommends configuring one of the time synchronization protocols supported. Failure to do so could result in the switch time being reset to the factory default of 01/01/1990 00:00:00 in the case of a switch reload, software upgrade, or power cycle.

TimeP time synchronization

You can either manually assign the switch to use a TimeP server or use DHCP to assign the TimeP server. In either case, the switch can get its time synchronization updates from only one designated TimeP server. This option enhances security by specifying which time server to use.

SNTP time synchronization

SNTP provides three operating modes:

- **Broadcast mode**

The switch acquires time updates by accepting the time value from the first SNTP time broadcast detected. (In this case, the SNTP server must be configured to broadcast time updates to the network broadcast address; see the documentation provided with your SNTP server application.) Once the switch detects a particular server, it ignores time broadcasts from other SNTP servers unless the configurable Poll Interval expires three consecutive times without an update received from the first-detected server.



To use Broadcast mode, the switch and the SNTP server must be in the same subnet.

- **DHCP mode**

DHCP mode is enabled by default. In DHCP mode, the SNTP server address and the timezone are provided in the DHCP address reply.

- **Unicast mode**

The switch requests a time update from the configured SNTP server. (You can configure one server using the menu interface, or up to three servers using the CLI `sntp server` command.) This option provides increased

security over the Broadcast mode by specifying which time server to use instead of using the first one detected through a broadcast.

Selecting a time synchronization protocol

Procedure

1. Select the time synchronization protocol: `TimeP`, `SNTP`, or `NTP`.
2. Enable the protocol; the choices are:
 - a. `TimeP`: `DHCP` or `Manual`
 - b. `SNTP`: `Broadcast` or `Unicast`
 - c. `NTP`: `Broadcast` or `Unicast`
3. Configure the remaining parameters for the time protocol you selected.

The switch retains the parameter settings for both time protocols even if you change from one protocol to the other. Thus, if you select a time protocol, the switch uses the parameters you last configured for the selected protocol.

Simply selecting a time synchronization protocol does not enable that protocol on the switch unless you also enable the protocol itself (step 2, above). For example, in the factory-default configuration, `TimeP` is the selected time synchronization method. However, because `TimeP` is disabled in the factory-default configuration, no time synchronization protocol is running.

Disabling time synchronization

You can use either of the following methods to disable time synchronization without changing the `TimeP`, `SNTP`, or `NTP` configuration:

- Global config level of the CLI
 - Execute `no timesync`.
- System Information screen of the Menu interface
 1. Set the `Time Sync Method` parameter to `None`.
 2. Press **[Enter]**, then **[S]** (for **Save**).

SNTP: Selecting and configuring

The following table shows the `SNTP` parameters and their operations.

Table 1: *SNTP* parameters

SNTP parameter	Operation
Time Sync Method	Used to select either <code>SNTP</code> , <code>TIMEP</code> , <code>NTP</code> , or <code>None</code> as the time synchronization method.
SNTP Mode	
<code>Disabled</code>	The Default. <code>SNTP</code> does not operate, even if specified by the Menu interface Time Sync Method parameter or the CLI <code>timesync</code> command.
<code>Unicast</code>	Directs the switch to poll a specific server for <code>SNTP</code> time synchronization. Requires at least one server address.

Table Continued

SNTP parameter	Operation
Broadcast	Directs the switch to acquire its time synchronization from data broadcast by any SNTP server to the network broadcast address. The switch uses the first server detected and ignores any others. However, if the Poll Interval expires three times without the switch detecting a time update from the original server, the switch accepts a broadcast time update from the next server it detects.
Poll Interval (seconds)	In Unicast Mode: Specifies how often the switch polls the designated SNTP server for a time update. In Broadcast Mode: Specifies how often the switch polls the network broadcast address for a time update.Value is between 30 to 720 seconds.
Server Address	Used only when the SNTP Mode is set to Unicast . Specifies the IP address of the SNTP server that the switch accesses for time synchronization updates. You can configure up to three servers; one using the menu or CLI, and two more using the CLI.
Server Version	Specifies the SNTP software version to use and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3. Default: 3; range: 1 to 7.
Priority	Specifies the order in which the configured servers are polled for getting the time. Value is between 1 and 3.

Viewing and configuring SNTP (Menu)

Procedure

1. From the Main Menu, select:
 - a. **2. Switch Configuration...**
 - b. **1. System Information**

Figure 1: System Information screen (default values)

```

===== CONSOLE - MANAGER MODE =====
Switch Configuration - System Information

System Name : HP Switch
System Contact :
System Location :

Inactivity Timeout (min) [0] : 0      MAC Age Time (sec) [300] : 300
Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled
Tftp-enable [Yes] : Yes
Time Zone [0] : 0
Daylight Time Rule [None] : None

Server Address :
Jumbo Max Frame Size [9216] : 9216
Jumbo IP MTU [9198] : 9198

Time Protocol Selection Parameter
- TIMEP
- SNTP
- None

Actions->  Cancel      Edit      Save      Help

```

2. Press **[E]** (for **Edit**).
 - Move the cursor to the **System Name** field.
3. Use the **Space** bar to move the cursor to the **Time Sync Method** field.
4. Use the **Space** bar to select **SNTP**, then move to the **SNTP Mode** field.
5. Complete one of the following options.

Option 1

- a. Use the **Space** bar to select the **Broadcast** mode.
- b. Move the cursor to the **Poll Interval** field.
- c. Go to 6. (For Broadcast mode details, see **SNTP time synchronization**)

Figure 2: Time configuration fields for SNTP with broadcast mode

```

Time Sync Method [None] : SNTP
SNTP Mode [Disabled] : Broadcast
Poll Interval (sec) [720] : 720
Tftp-enable [Yes] : Yes
Time Zone [0] : 0
Daylight Time Rule [None] : None

```

Option 2

- d. Use the **Space** bar to select the **Unicast** mode.
- e. Move the cursor to the **Server Address** field.
- f. Enter the IP address of the SNTP server you want the switch to use for time synchronization.



This step replaces any previously configured server IP address. If you will be using backup SNTP servers (requires use of the CLI), see **SNTP unicast time polling with multiple SNTP servers**.

- g. Move the cursor to the **Server Version** field. Enter the value that matches the SNTP server version running on the device you specified in the preceding step.

If you are unsure which version to use, Hewlett Packard Enterprise recommends leaving this value at the default setting of 3 and testing SNTP operation to determine whether any change is necessary.



Using the menu to enter the IP address for an SNTP server when the switch already has one or more SNTP servers configured, the switch deletes the primary SNTP server from the server list. The switch then selects a new primary SNTP server from the IP addresses in the updated list. For more on this topic, see **SNTP unicast time polling with multiple SNTP servers**.

- h. Move the cursor to the **Poll Interval** field, then go to step 6.

Figure 3: SNTP configuration fields for SNTP configured with unicast mode

```

Time Sync Method [None] : SNTP
SNTP Mode [Disabled] : Unicast           Server Address : 10.28.227.15
Poll Interval (sec) [720] : 720         Server Version [3] : 3
Tftp-enable [Yes] : Yes
Time Zone [0] : 0                       ←
Daylight Time Rule [None] : None

```

Note: The Menu interface lists only the highest priority SNTP server, even if others are configured. To view all SNTP servers configured on the switch, use the CLI **show management** command. Refer to “SNTP Unicast Time Polling with Multiple SNTP Servers” on page 1-33.

6. In the **Poll Interval** field, enter the time in seconds that you want for a Poll Interval.
(For Poll Interval operation, see **SNTP parameters**)
7. Press **Enter** to return to the Actions line, then **S** (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

Viewing and configuring SNTP (CLI)

Syntax:

```
show sntp
```

Lists both the time synchronization method (TimeP, SNTP, or None) and the SNTP configuration, even if SNTP is not the selected time protocol.

If you configure the switch with SNTP as the time synchronization method, then enable SNTP in broadcast mode with the default poll interval, `show sntp` lists the following:

SNTP configuration when SNTP is the selected time synchronization method

```
switch(config)# show sntp
```

```
SNTP Configuration
```

```
Time Sync Mode: Sntp
SNTP Mode : Unicast
Poll Interval (sec) [720] : 719
```

Priority	SNTP Server Address	Protocol Version
1	2001:db8::215:60ff:fe79:8980	7
2	10.255.5.24	3
3	fe80::123%vlan10	3

In the factory-default configuration (where TimeP is the selected time synchronization method), `show sntp` still lists the SNTP configuration, even though it is not currently in use. In **SNTP configuration when SNTP is not the selected time synchronization method** on page 23, even though TimeP is the current time synchronous method, the switch maintains the SNTP configuration.

SNTP configuration when SNTP is not the selected time synchronization method

```
switch(config)# show sntp
```

```
SNTP Configuration
```

```
Time Sync Mode: Timep
SNTP Mode : Unicast
Poll Interval (sec) [720] : 719
```

Priority	SNTP Server Address	Protocol Version
1	2001:db8::215:60ff:fe79:8980	7
2	10.255.5.24	3
3	fe80::123%vlan10	3

Syntax:

```
show management
```

This command can help you to easily examine and compare the IP addressing on the switch. It lists the IP addresses for all time servers configured on the switch, plus the IP addresses and default gateway for all VLANs configured on the switch.

Display showing IP addressing for all configured time servers and VLANs

```
switch(config)# show management

Status and Counters - Management Address Information

Time Server Address : fe80::215:60ff:fe7a:adc0%vlan10

Priority  SNTP Server Address          Protocol Version
-----  -
1         2001:db8::215:60ff:fe79:8980       7
2         10.255.5.24                        3
3         fe80::123%vlan10                   3

Default Gateway   :10.0.9.80

VLAN Name      MAC Address      | IP address
-----+-----
DEFAULT_VLAN  001279-88a100   | Disabled
VLAN10        001279-88a100   | 10.0.10.17
```

Configuring (enabling or disabling) the SNTP mode

Enabling the SNTP mode means to configure it for either broadcast or unicast mode. Remember that to run SNTP as the switch's time synchronization protocol, you must also select SNTP as the time synchronization method by using the CLI `timesync` command (or the menu interface **Time Sync Method** parameter.)

Syntax:

```
timesync sntp
```

Selects SNTP as the time protocol.

```
sntp {<broadcast | unicast>}
```

Enables the SNTP mode.

Syntax:

```
sntp server <ip-addr>
```

Required only for unicast mode.

Syntax:

```
sntp server priority <1-3>
```

Specifies the order in which the configured servers are polled for getting the time. Value is between 1 and 3.

Syntax:

```
sntp <30-720>
```

Configures the amount of time between updates of the system clock via SNTP.

Default: 720 seconds

Enabling SNTP in Broadcast Mode

Because the switch provides an SNTP polling interval (default: 720 seconds), you need only these two commands for minimal SNTP broadcast configuration:

Syntax:

```
timesync sntp
```

Selects SNTP as the time synchronization method.

Syntax:

```
sntp broadcast
```

Configures broadcast as the SNTP mode.

Example:

Suppose that time synchronization is in the factory-default configuration (TimeP is the currently selected time synchronization method.) Complete the following:

Procedure

1. View the current time synchronization.
2. Select **SNTP** as the time synchronization mode.
3. Enable **SNTP** for Broadcast mode.
4. View the SNTP configuration again to verify the configuration.

The commands and output would appear as follows:

Figure 4: *Enabling SNTP operation in Broadcast Mode*

```
switch(config)# show sntp 1
SNTP Configuration
Time Sync Mode: Timep
SNTP Mode : disabled
Poll Interval (sec) [720] :720
```

```
switch(config)# timesync sntp
```

```
switch(config)# sntp broadcast
```

```
switch(config)# show sntp 2
SNTP Configuration
Time Sync Mode: Sntp
SNTP Mode : Broadcast
Poll Interval (sec) [720] :720
```

• 1

`show sntp` displays the SNTP configuration and also shows that TimeP is the currently active time synchronization method.

• 2

`show sntp` again displays the SNTP configuration and shows that SNTP is now the currently active time synchronization method.

Enabling SNTP in unicast mode (CLI)

Like broadcast mode, configuring SNTP for unicast mode enables SNTP. However, for unicast operation, you must also specify the IP address of at least one SNTP server. The switch allows up to three unicast servers. You can use the Menu interface or the CLI to configure one server or to replace an existing unicast server with another. To add a second or third server, you must use the CLI. For more on SNTP operation with multiple servers, see [SNTP unicast time polling with multiple SNTP servers](#) on page 36

Syntax:

```
timesync sntp
```

Selects SNTP as the time synchronization method.

Syntax:

```
sntp unicast
```

Configures the SNTP mode for unicast operation.

Syntax:

```
[no] sntp server priority < 1-3 > < ip-address > [version]
```

Use the `no` version of the command to disable SNTP.

- priority** Specifies the order in which the configured SNTP servers are polled for the time.
- ip-address** An IPv4 or IPv6 address of an SNTP server.
- version** The protocol version of the SNTP server. Allowable values are 1 through 7; default is 3.

Syntax:

```
no sntp server priority <1-3> <ip-addr>
```

Deletes the specified SNTP server.



```
priority <1-3>
```

value must match what server is configured with. Deleting an SNTP server when only one is configured disables SNTP unicast operation.

Example:

To select SNTP and configure it with unicast mode and an SNTP server at 10.28.227.141 with the default server version (3) and default poll interval (720 seconds):

```
switch(config)# timesync sntp
```

Selects SNTP.

```
switch(config)# sntp unicast
```

Activates SNTP in unicast mode.

```
switch(config)# sntp server priority 1 10.28.227.141
```

Specifies the SNTP server and accepts the current SNTP server version (default: 3).

Configuring SNTP for unicast operation

```
switch(config)# show sntp
```

```
SNTP Configuration
```

```
Time Sync Mode: Sntp  
SNTP Mode : Unicast  
Poll Interval (sec) [720] : 720
```

Priority	SNTP Server Address	Protocol Version
1	2001:db8::215:60ff:fe79:8980	7
2	10.255.5.24	3
3	fe80::123%vlan10	3

In this Example:, the **Poll Interval** and the **Protocol Version** appear at their default settings.

Both IPv4 and IPv6 addresses are displayed.

Note: Protocol Version appears only when there is an IP address configured for an SNTP server.

If the SNTP server you specify uses SNTP v4 or later, use the `sntp server` command to specify the correct version number. For example, suppose you learned that SNTP v4 was in use on the server you specified above (IP address 10.28.227.141). You would use the following commands to delete the server IP address , re-enter it with the correct version number for that server.

Specifying the SNTP protocol version number

```
switch(config)# no sntp server 10.28.227.141 1  
switch(config)# sntp server 10.28.227.141 4 2  
switch(config)# show sntp
```

```
SNTP Configuration
```

```
Time Sync Mode: Sntp  
SNTP Mode : Broadcast  
Poll Interval (sec) [720] : 600
```

IP Address	Protocol Version
10.28.227.141	4 3

- ¹Deletes unicast SNTP server entry.
- ²Re-enters the unicast server with a non-default protocol version.
- ³show sntp displays the result.

Changing the SNTP poll interval (CLI)

Syntax:

```
sntp <30..720>
```

Specifies the amount of time between updates of the system clock via SNTP. The default is 720 seconds and the range is 30 to 720 seconds. (This parameter is separate from the poll interval parameter used for Timep operation.)

Example:

To change the poll interval to 300 seconds:

```
switch(config)# sntp 300
```

Changing the SNTP server priority (CLI)

You can choose the order in which configured servers are polled for getting the time by setting the server priority.

Syntax:

```
sntp server priority <1-3> <ip-address>
```

Specifies the order in which the configured servers are polled for getting the time Value is between 1 and 3.



You can enter both IPv4 and IPv6 addresses. For more information about IPv6 addresses, see the IPv6 configuration guide for your switch.

Example:

To set one server to priority 1 and another to priority 2:

```
switch(config)# sntp server priority 1 10.28.22.141
switch(config)# sntp server priority 2
                2001:db8::215:60ff:fe79:8980
```

Disabling time synchronization without changing the SNTP configuration (CLI)

The recommended method for disabling time synchronization is to use the `timesync` command.

Syntax:

```
no timesync
```

Halts time synchronization without changing your SNTP configuration.

Example:

Suppose SNTP is running as the switch's time synchronization protocol, with `broadcast` as the SNTP mode and the factory-default polling interval. You would halt time synchronization with this command:

```
switch(config)# no timesync
```

If you then viewed the SNTP configuration, you would see the following:

SNTP with time synchronization disabled

```
switch(config)# show sntp
SNTP Configuration
  Time Sync Mode: Disabled
  SNTP Mode : Broadcast
  Poll Interval (sec) [720] : 720
```

Disabling the SNTP Mode

If you want to prevent SNTP from being used even if it is selected by `timesync` (or the Menu interface's `Time Sync Method` parameter), configure the SNTP mode as disabled.

Syntax:

```
no sntp
```

Disables SNTP by changing the SNTP mode configuration to `Disabled`.

Example:

If the switch is running SNTP in unicast mode with an SNTP server at 10.28.227.141 and a server version of 3 (the default), `no sntp` changes the SNTP configuration as shown below and disables time synchronization on the switch.

Disabling time synchronization by disabling the SNTP mode

```
switch(config)# no sntp
switch(config)# show sntp

SNTP Configuration

Time Sync Mode: Sntp
SNTP Mode : disabled
Poll Interval (sec) [720] : 600

IP Address      Protocol Version
-----
10.28.227.141  3
```

Note that even though the **Time Sync Mode** is set to **Sntp**, time synchronization is disabled because `no sntp` has disabled the **SNTP Mode** parameter.

TimeP: Selecting and configuring

The following table shows TimeP parameters and their operations.

Table 2: *TimeP parameters*

TimeP parameter	Operation
Time Sync Method	Used to select either TIMEP, SNTP, NTP, or None as the time synchronization method.
TimeP Mode	
Disabled	TimeP does not operate, even if specified by the Menu interface Time Sync Method parameter or the CLI <code>timesync</code> command.
DHCP	When TimeP is selected as the time synchronization method, the switch attempts to acquire a TimeP server IP address via DHCP. If the switch receives a server address, it polls the server for updates according to the TimeP poll interval. If the switch does not receive a TimeP server IP address, it cannot perform time synchronization updates.

Table Continued

TimeP parameter	Operation
Manual	When TimeP is selected as the time synchronization method, the switch attempts to poll the specified server for updates according to the TimeP poll interval. If the switch fails to receive updates from the server, time synchronization updates do not occur.
Server Address	Used only when the TimeP Mode is set to Manual . Specifies the IP address of the TimeP server that the switch accesses for time synchronization updates. You can configure one server.

Viewing, enabling, and modifying the TimeP protocol (Menu)

Procedure

1. From the Main Menu, select:

2. Switch Configuration

1. System Information

Figure 5: System Information screen (default values)

```

===== CONSOLE - MANAGER MODE =====
Switch Configuration - System Information

System Name : HP Switch
System Contact :
System Location :

Inactivity Timeout (min) [0] : 0      MAC Age Time (sec) [300] : 300
Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled
Tftp-enable [Yes] : Yes
Time Zone [0] : 0
Daylight Time Rule [None] : None

Server Address :
Jumbo Max Frame Size [9216] : 9216
Jumbo IP MTU [9198] : 9198

Time Protocol Selection Parameter
- TIMEP (the default)
- SNTP
- None

Actions->  Cancel      Edit      Save      Help

```

2. Press **[E]** (for **Edit**).

The cursor moves to the **System Name** field.

3. Move the cursor to the **Time Sync Method** field.

4. If **TIMEP** is not already selected, use the **Space** bar to select **TIMEP**, then move to the **TIMEP Mode** field.

5. Do one of the following:

- Use the **Space** bar to select the **DHCP** mode.
 - Move the cursor to the **Poll Interval** field.
 - Go to step 6.

Enabling TIMEP or DHCP

```

Time Sync Method [None] :  TIMEP
TimeP Mode [Disabled] :   DHCP
Poll Interval (min) [720] : 720
Time Zone [0] :          0

```

Daylight Time Rule [None] : None

- Use the **Spacebar** to select the **Manual** mode.
 - Move the cursor to the **Server Address** field.
 - Enter the IP address of the TimeP server you want the switch to use for time synchronization.



This step replaces any previously configured TimeP server IP address.

- Move the cursor to the **Poll Interval** field, then go to step 6.
6. In the **Poll Interval** field, enter the time in minutes that you want for a TimeP Poll Interval.
 7. Select **[Enter]** to return to the **Actions** line, then select **[S]** (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

Viewing the current TimeP configuration (CLI)

Using different `show` commands, you can display either the full TimeP configuration or a combined listing of all TimeP, SNTP, and VLAN IP addresses configured on the switch.

Syntax:

```
show timep
```

Lists both the time synchronization method (TimeP, SNTP, or None) and the TimeP configuration, even if SNTP is not the selected time protocol. (If the TimeP Mode is set to `Disabled` or `DHCP`, the Server field does not appear.)

If you configure the switch with TimeP as the time synchronization method, then enable TimeP in DHCP mode with the default poll interval, `show timep` lists the following:

TimeP configuration when TimeP is the selected Time synchronization method

```
switch(config)# show timep
```

```
Timep Configuration
```

```
Time Sync Mode: Timep
TimeP Mode [Disabled] : DHCP      Server Address : 10.10.28.103
Poll Interval (min) [720] : 720
```

If SNTP is the selected time synchronization method, `show timep` still lists the TimeP configuration even though it is not currently in use. Even though, in this Example:, SNTP is the current time synchronization method, the switch maintains the TimeP configuration:

TimeP configuration when TimeP is not the selected time synchronization method

```
switch(config)# show timep
```

```
Timep Configuration
```

```
Time Sync Mode: Sntp
TimeP Mode [Disabled] : Manual    Server Address : 10.10.28.100
Poll Interval (min) [720] : 720
```

Syntax:

```
show management
```

Helps you to easily examine and compare the IP addressing on the switch. It lists the IP addresses for all time servers configured on the switch plus the IP addresses and default gateway for all VLANs configured on the switch.

Display showing IP addressing for all configured time servers and VLANs

```
switch(config)# show management
```

```
Status and Counters - Management Address Information
```

```
Time Server Address : 10.10.28.100
```

Priority	SNTP Server Address	Protocol Version
1	10.10..28.101	3
2	10.255.5.24	3
3	fe80::123%vlan10	3

```
Default Gateway : 10.0.9.80
```

VLAN Name	MAC Address	IP Address
DEFAULT_VLAN	001279-88a100	10.30.248.184
VLAN10	001279-88a100	10.0.10.17

Configuring (enabling or disabling) the TimeP mode

Enabling the TimeP mode means to configure it for either broadcast or unicast mode. Remember to run TimeP as the switch's time synchronization protocol, you must also select TimeP as the time synchronization method by using the CLI `timesync` command (or the menu interface **Time Sync Method** parameter).

Syntax:

```
timesync timep
```

Selects TimeP as the time synchronization method.

Syntax:

```
ip timep {<dhcp | manual>}
```

Enables the selected TimeP mode.

Syntax:

```
[no] ip timep
```

Disables the TimeP mode.

Syntax:

```
[no] timesync
```

Disables the time protocol.

Enabling TimeP in manual mode (CLI)

Like DHCP mode, configuring TimeP for `manual` mode enables TimeP. However, for manual operation, you must also specify the IP address of the TimeP server. (The switch allows only one TimeP server.)

Syntax:

```
timesync timep
```

Selects TimeP.

Syntax:

```
ip timep manual <ip-addr>
```

Activates TimeP in manual mode with a specified TimeP server.

Syntax:

```
no ip timep
```

Disables TimeP.

Enabling TimeP in DHCP Mode

Because the switch provides a TimeP polling interval (default:720 minutes), you need only these two commands for a minimal TimeP DHCP configuration:

Syntax:

```
timesync timep
```

Selects TimeP as the time synchronization method.

Syntax:

```
ip timep dhcp
```

Configures DHCP as the TimeP mode.

For example, suppose:

- Time Synchronization is configured for SNTP.
- You want to:
 - View the current time synchronization.
 - Select TimeP as the synchronization mode.
 - Enable TimeP for DHCP mode.
 - View the TimeP configuration.

Enabling TimeP in Manual Mode

Like DHCP mode, configuring TimeP for Manual Mode enables TimeP. However, for manual operation, you must also specify the IP address of the TimeP server. (The switch allows only one TimeP server.) To enable the TimeP protocol:

Syntax:

```
timesync timep
```

Selects TimeP.

Syntax:

```
ip timep manual <ip-addr>
```

Activates TimeP in manual mode with a specified TimeP server.

Syntax:

```
[no] ip timep
```

Disables TimeP.



To change from one TimeP server to another, you must use the `no ip timep` command to disable TimeP mode, the reconfigure TimeP in manual mode with the new server IP address.

Example:

To select TimeP and configure it for manual operation using a TimeP server address of 10.28.227.141 and the default poll interval (720 minutes, assuming the TimeP poll interval is already set to the default):

```
switch(config)# timesync timep
```

Selects TimeP.

```
switch(config)# ip timep manual 10.28.227.141
```

Activates TimeP in Manual mode.

Configuring TimeP for manual operation

```
switch(config)# timesync timep
switch(config)# ip timep manual 10.28.227.141
switch(config)# show timep
Timep Configuration
  Time Sync Mode: Timep
  TimeP Mode : Manual          Server Address : 10.28.227.141
  Poll Interval (min) : 720
```

Changing from one TimeP server to another (CLI)

Procedure

1. Use the `no ip timep` command to disable TimeP mode.
2. Reconfigure TimeP in Manual mode with the new server IP address.

Changing the TimeP poll interval (CLI)

Syntax:

```
ip timep {< dhcp | manual >} interval <1-9999>
```

Specifies how long the switch waits between time polling intervals. The default is 720 minutes and the range is 1 to 9999 minutes. (This parameter is separate from the `poll interval` parameter used for SNTP operation.)

Example:

To change the poll interval to 60 minutes:

```
switch(config)# ip timep interval 60
```

Disabling time synchronization without changing the TimeP configuration (CLI)

Syntax:

```
no timesync
```

Disables time synchronization by changing the `Time Sync Mode` configuration to `Disabled`. This halts time synchronization without changing your TimeP configuration. The recommended method for disabling time synchronization is to use the `timesync` command.

Example:

Suppose TimeP is running as the switch's time synchronization protocol, with `DHCP` as the TimeP mode, and the factory-default polling interval. You would halt time synchronization with this command:

```
switch(config)# no timesync
```

If you then viewed the TimeP configuration, you would see the following:

TimeP with time synchronization disabled

```
switch(config)# show timep

Timep Configuration
Time Sync Mode: Disabled
TimeP Mode : DHCP Poll Interval (min): 720
```

Disabling the TimeP mode

Syntax:

```
no ip timep
```

Disables TimeP by changing the TimeP mode configuration to `Disabled` and prevents the switch from using it as the time synchronization protocol, even if it is the selected `Time Sync Method` option.

Example:

If the switch is running TimeP in `DHCP` mode, `no ip timep` changes the TimeP configuration as shown below and disables time synchronization. Even though the `TimeSync` mode is set to TimeP, time synchronization is disabled because `no ip timep` has disabled the TimeP mode parameter.

Disabling time synchronization by disabling the TimeP mode parameter

```
switch(config)# no ip timep

switch(config)# show timep

Timep Configuration
Time Sync Mode: Timep
TimeP Mode : Disabled
```

SNTP unicast time polling with multiple SNTP servers

When running SNTP unicast time polling as the time synchronization method, the switch requests a time update from the server you configured with either the `Server Address` parameter in the menu interface, or the primary server in a list of up to three SNTP servers configured using the CLI. If the switch does not receive a response from the primary server after three consecutive polling intervals, the switch tries the next server (if any) in the list. If the switch tries all servers in the list without success, it sends an error message to the Event Log and reschedules to try the address list again after the configured `Poll Interval` time has expired.

If there are already three SNTP server addresses configured on the switch, and you want to use the CLI to replace one of the existing addresses with a new one, you must delete the unwanted address before you configure the new one.

Displaying all SNTP server addresses configured on the switch (CLI)

The System Information screen in the menu interface displays only one SNTP server address, even if the switch is configured for two or three servers. The CLI `show management` command displays all configured SNTP servers on the switch.

How to list all SNTP servers configured on the switch

```
switch(config)# show management

Status and Counters - Management Address Information

Time Server Address : fe80::215:60ff:fe7a:adc0%vlan10

Priority SNTP Server Address                                Protocol Version
-----
1 2001:db8::215:60ff:fe79:8980                             7
2 10.255.5.24                                               3
3 fe80::123%vlan10                                         3

Default Gateway : 10.0.9.80

VLAN Name      MAC Address      | IP Address
-----+-----
DEFAULT_VLAN  001279-88a100   | Disabled
VLAN10        001279-88a100   | 10.0.10.17
```

Adding and deleting SNTP server addresses

Adding addresses

As mentioned earlier, you can configure one SNTP server address using either the Menu interface or the CLI. To configure a second and third address, you must use the CLI. To configure the remaining two addresses, you would do the following:

Creating additional SNTP server addresses with the CLI

```
switch(config)# sntp server priority <1-3> 2001:db8::215:60ff:fe79:8980  
switch(config)# sntp server 10.255.5.24
```



If there are already three SNTP server addresses configured on the switch, and you want to use the CLI to replace one of the existing addresses with a new one, you must delete the unwanted address before you configure the new one.

Deleting addresses

Syntax:

```
no sntp server <ip-addr>
```

Deletes a server address. If there are multiple addresses and you delete one of them, the switch re-orders the address priority.

Example:

To delete the primary address in the above Example: and automatically convert the secondary address to primary:

```
switch(config)# no sntp server 10.28.227.141
```

Operating with multiple SNTP server addresses configured (Menu)

When you use the Menu interface to configure an SNTP server IP address, the new address writes over the current primary address, if one is configured.

SNTP messages in the Event Log

If an SNTP time change of more than three seconds occurs, the switch's Event Log records the change. SNTP time changes of less than three seconds do not appear in the Event Log.

Network Time Protocol (NTP)

All NTP communications use Coordinated Universal Time (UTC). An NTP server usually receives its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server, and then distributes this time across the network. NTP is extremely efficient; no more than one packet per minute is necessary to synchronize two machines to within a millisecond of each other.

NTP uses a stratum to describe the distance between a network device and an authoritative time source:

- A stratum 1 time server is directly attached to an authoritative time source (such as a radio or atomic clock or a GPS time source).
- A stratum 2 NTP server receives its time through NTP from a stratum 1 time server.

Before synchronizing, NTP compares the time reported by several network devices and does not synchronize with one that is significantly different, even if it is a stratum 1.

The security features of NTP can be used to avoid the accidental or malicious setting of incorrect time. One such mechanism is available: an encrypted authentication mechanism.

Though similar, the NTP algorithm is more complex and accurate than the Simple Network Time Protocol (SNTP).



Enabling this feature results in synchronizing the system clock; therefore, it may affect all sub-systems that rely on system time.

Commands

The following commands allow the user to configure NTP or show NTP configurations.

timesync Command

This command is used to configure the protocol used for network time synchronization.

Syntax

```
[no] timesync { timep | sntp | timep-or-sntp | ntp }
```

Options

`no`

Deletes all timesync configurations on the device.

`timep`

Updates the system clock using TIMEP.

`sntp`

Updates the system clock using SNTP.

`timep-or-sntp`

Updates the system clock using TIMEP or SNTP (default).

`ntp`

Updates the system clock using NTP

Example

```
switch(config)# timesync
sntp          Update the system clock using SNTP.
timep        Update the system clock using TIMEP.
timep-or-sntp Update the system clock using TIMEP or SNTP.
ntp          Update the system clock using NTP.
```

timesync ntp

This command is used to update the system clock using NTP.

Syntax

```
timesync ntp
```

Description

Update the system clock using NTP.

ntp

This command selects the operating mode of the NTP client.

Syntax

```
ntp [broadcast|unicast]
```

Options

`broadcast`

Sets ntp server to operate in broadcast mode.

`unicast`

Sets ntp server to operate in unicast mode.

Usage

The default mode is broadcast.

[no] ntp

This command disables NTP and removes all NTP configurations on the device.

Syntax

```
[no] ntp [authentication <key-id>  
| broadcast | enable | max-association  
<integer> | server  
<IP-ADDR> | trap  
<trap-name> | unicast]
```

Description

Disable NTP and removes the entire NTP configuration.

Options

authentication	Configure NTP authentication.
broadcast	Operate in broadcast mode.
enable	Enable/disable NTP.
max-association	Maximum number of Network Time Protocol (NTP) associations.
server	Configure a NTP server to poll for time synchronization.
trap	Enable/disable NTP traps.
unicast	Operate in unicast mode.

Example

```
switch(config)# no ntp  
This will delete all NTP configurations on this device. Continue [y/n]?
```

ntp enable

This command is used to enable or disable NTP on the switch.

Syntax

```
ntp enable
```

Example

```
switch(config)# ntp  
enable      Enable/disable NTP.
```

Description

Enable or disable NTP. Use [no] to disable NTP.

Restrictions

Validation	Error/Warning/Prompt
If timeSync is in SNTP or Timep when NTP is enabled.	Timesync is not configured to NTP.
When timesync is NTP and ntp is enabled and we try to change timesync to SNTP.	Disable NTP before changing timesync to SNTP or TIMEP

ntp authentication

This command is used for authentication of NTP server by the NTP client.

Syntax

```
ntp authentication key-id <KEY-ID> [authentication-mode <MODE> key-value <KEY-  
STRING>] [trusted]
```

Parameters/Options

```
key-id <id>
```

Sets the key-id for the authentication key.

Subcommands

```
authentication-mode
```

Sets the NTP authentication mode

```
key-value <KEY-STRING>
```

Sets the key-value for the authentication key.

```
[trusted]
```

Sets the authentication key as trusted.

Example

```
Switch(config)# ntp  
Authentication    Configure NTP authentication.
```

```
Switch(config)# ntp authentication  
key-id           Set the key-id for this authentication key.
```

```
Switch(config)# ntp authentication key-id  
<1-4294967295>   Set the authentication key-id.
```



```
Switch(config)# ntp authentication key-id 1
authentication-mode Set the NTP authentication mode.
trusted Set this authentication key as trusted.
```

```
Switch(config)# ntp authentication key-id 1
authentication-mode|trusted md5
Authenticate using MD5.
```

```
Switch(config)# ntp authentication key-id 1
authentication-mode|trusted md5key-value Set the NTP authentication key.
```

```
Switch(config)# ntp authentication key-id 1
authentication-mode|trusted md5 key-value
KEY Enter a string to be set as the NTP authentication key.
```

ntp authentication key-id

Syntax

```
ntp authentication key-id
<key-id> [authentication-mode [md5 | sha1]
key-value <key-value>] [trusted]
```

Description

The NTP client authenticates the NTP server.

Options

authentication-mode	Set the NTP authentication mode. <ul style="list-style-type: none">md5: Authenticate using MD5.sha1: Authenticate using SHA1.
trusted	Set this authentication key as trusted.

ntp max-association

This command is used to configure the maximum number of servers associated with this NTP client.

Syntax

```
ntp max-association
<number>
```

Options

```
max-association <number>
```

Sets the maximum number of NTP associations.

Description

Configure maximum number of servers associated with the client. Up to eight servers can be configured as the maximum.

Restrictions

The range for a maximum number of NTP associations is 1–8.

Example

```
Switch(config)# ntp
max-associations    Maximum number of NTP associations.
```

```
Switch(config)# ntp max-associations
<1-8>              Enter the number.
```

Restrictions

Validation	Error/Warning/Prompt
When the number of configured NTP servers is more than the max-associations value.	The maximum number of NTP servers allowed is <number>.
When the max-associations value is less than the (n) number of configured NTP servers.	Max-associations value cannot be less than the number of NTP servers configured.

ntp server

This command is used to configure the NTP servers.

Syntax

```
[no] ntp server
```

```
ntp server <IP-ADDR|IPv6-ADDR> [key <key-id>] [oobm] [max-poll <max-poll-val>][min-poll <min-poll-val>][burst | iburst] [version <1-4>]
```

Parameters/Options

[no]

Removes the unicast NTP configurations on the device.

Subcommands

IP-ADDR

Sets the IPv4 address of the NTP server.

IPv6-ADDR

Sets the IPv6 address of the NTP server.

oobm

Specifies that the NTP Unicast server is accessible over an OOBM interface.

key <key-id>

Specifies the authentication key.

max-poll <max-poll-val>

Configures the maximum time intervals in power of 2 seconds. Range is 4–17 (e.g., 5 would translate to 2 raised to 5 or 32).

```
min-poll <min-poll-val>
```

Configures the minimum time intervals in seconds. Range is 4–17.

```
burst
```

Enables burst mode.

```
iburst
```

Enables initial burst mode.

```
version
```

Sets version 1–4.

Usage

A maximum of 8 NTP servers can be configured.

Example

```
switch(config)# ntp
server      Allow the software clock to be synchronized by an NTP
time server.
broadcast   Operate in broadcast mode.
unicast     Operate in unicast mode.
```

```
switch(config)# ntp server
IP-ADDR     IPv4 address of the NTP server.
IPV6-ADDR   IPv6 address of the NTP server.
```

```
switch(config)# ntp server <IP-ADDR>
Key         Specify the authentication key.
```

```
switch(config)# ntp server <IP-ADDR> key key-id
Max-poll    Configure the maximum time intervals in seconds.
```

```
switch(config)# ntp server <IP-ADDR> key key-id max-poll
<4-17>     Enter an integer number.
```

```
Switch(config)# ntp server <IP-ADDR> key key-id
Min-poll    Configure the minimum time intervals in seconds.
```

```
switch(config)# ntp server <IP-ADDR> key key-id min-poll
<4-17>     Enter an integer number.
```

```
switch(config)# ntp server <IP-ADDR> key key-id prefer max-poll
<max-poll-val> min-poll <min-poll-val>
iburst     Enable initial burst (iburst) mode.
burst      Enable burst mode.
```

```
Switch(config)# ntp server IP-ADDR key key-id prefer maxpoll <number>
minpoll <number> iburst
```

Restrictions

Validation	Error/Warning/Prompt
If authentication key-id not configured	Authentication key-id has not been configured.
If Key-id is not marked as trusted	Key-id is not trusted.
When min poll value is more than max poll value	NTP max poll value should be more than min poll value.

ntp server key-id

Syntax

```
ntp server <IP-ADDR | IPV6-ADDR>
key—id <key-id> [max-poll
<max-poll-val>] [min-poll
<min-poll-val>] [burst | iburst]
```

Description

Configure the NTP server. <IP-ADDR> indicates the IPv4 address of the NTP server. <IPV6-ADDR> indicates the IPv6 address of the NTP server.

Options

burst	Enables burst mode.
iburst	Enables initial burst (iburst) mode.
key-id	Set the authentication key to use for this server.
max-poll <max-poll-val>	Configure the maximum time intervals in seconds.
min-poll <min-poll-val>	Configure the minimum time intervals in seconds.

ntp ipv6-multicast

This command is used to configure NTP multicast on a VLAN interface.

Syntax

```
ntp ipv6-multicast
```

Description

Configure the interface to listen to the NTP multicast packets.

Example

```
Switch(vlan-2)# ntp
ipv6-multicast    Configure the interface to listen to the NTP multicast packets.
```

Restrictions

Validation	Error/Warning/Prompt
If ipv6 is not enabled on vlan interface	IPv6 address not configured on the VLAN.

debug ntp

This command is used to display debug messages for NTP.

Syntax

```
debug ntp <event |
packet>
```

Options

event

Displays event log messages related to NTP.

packets

Displays NTP packet messages.

Description

Enable debug logging. Use [no] to disable debug logging.

Example

```
Switch(config)# debug ntp
event          Display event log messages related to NTP.
packet         Display NTP packet messages.
```

ntp trap

This command is used to configure NTP traps.

Syntax

```
ntp trap <trap-name>
```

Description

Enable NTP traps. Use [no] to disable NTP traps.

Options

ntp-mode-change	Trap name resulting in send notification when the NTP entity changes mode, including starting and stopping (if possible).
ntp-stratum-change	Trap name resulting in send notification when stratum level of NTP changes.
ntp-peer-change	Trap name resulting in send notification when a (new) syspeer has been selected.
ntp-new-association	Trap name resulting in send notification when a new association is mobilized.

ntp-remove-association	Trap name resulting in send notification when an association is demobilized.
ntp-config-change	Trap name resulting in send notification when the NTP configuration has changed.
ntp-leapsec-announced	Trap name resulting in send notification when a leap second has been announced.
ntp-alive-heartbeat	Trap name resulting in send notification periodically (as defined by ntpEntHeartbeatInterval) to indicate that the NTP entity is still alive.
all	Enable all traps.

Usage

The traps defined below are generated as the result of finding an unusual condition while parsing an NTP packet or a processing a timer event. Note that if more than one type of unusual condition is encountered while parsing the packet or processing an event, only the first one will generate a trap. Possible trap names are:

- 'ntpEntNotifModeChange' The notification to be sent when the NTP entity changes mode, including starting and stopping (if possible).
- 'ntpEntNotifStratumChange' The notification to be sent when stratum level of NTP changes.
- 'ntpEntNotifSyspeerChanged' The notification to be sent when a (new) syspeer has been selected.
- 'ntpEntNotifAddAssociation' The notification to be sent when a new association is mobilized.
- 'ntpEntNotifRemoveAssociation' The notification to be sent when an association is demobilized.
- 'ntpEntNotifConfigChanged' The notification to be sent when the NTP configuration has changed.
- 'ntpEntNotifLeapSecondAnnounced' The notification to be sent when a leap second has been announced.
- 'ntpEntNotifHeartbeat' The notification to be sent periodically (as defined by ntpEntHeartbeatInterval) to indicate that the NTP entity is still alive.
- 'ntpEntNotifAll' The notification to be sent when all traps have been enabled

show ntp statistics

This command is used to show NTP statistics.

Syntax

```
show ntp statistics
```

Description

Show information about NTP packets.

Examples

```
Switch(config)# show ntp statistics
```

NTP Global statistics information

```
NTP In Packets      : 100
NTP Out Packets    : 110
NTP Bad Version Packets : 4
NTP Protocol Error Packets : 0
```

switch(config)# show ntp statistics

```
NTP Global statistics information

NTP In Packets      : 100
NTP Out Packets    : 110
NTP Bad Version Packets : 4
NTP Protocol Error Packets : 0
```

show ntp status

Syntax

Description

Show the status of NTP.

```
show ntp status
```

Example

```
Switch(config)# show ntp status

NTP Status information
NTP Status      : Disabled      NTP Mode      : Broadcast
Synchronization Status : Synchronized  Peer Dispersion : 8.01 sec
Stratum Number   : 2            Leap Direction : 1
Reference Assoc Id : 1          Clock Offset  : 0.0000 sec
Reference        : 192.0.2.1    Root Delay    : 0.00 sec
Precision        : 2**7        Root Dispersion : 15.91 sec
NTP Uptime       : 01d 09h 15m   Time Resolution : 1
Drift            : 0.000000000 sec/sec

System Time      : Tue Aug 25 04:59:11 2015
Reference Time   : Mon Jan 1 00:00:00 1990
```

show ntp associations

Syntax

```
show ntp associations [detail
<IP-ADDR>]
```

Description

Show the status of configured NTP associations.

Options

detail Show the detailed status of NTP associations configured for the system.

Switch(config)# show ntp associations

```

                                NTP Associations Entries
-----
Address      St   T   When Poll  Reach   Delay   Offset   Dispersion
-----
121.0.23.1   16  u   - 1024    0       0.000   0.000     0.000
231.45.21.4  16  u   - 1024    0       0.000   0.000     0.000
55.21.56.2   16  u   - 1024    0       0.000   0.000     0.000
23.56.13.1   3   u  209 1024    377     54.936  -6.159    12.688
91.34.255.216 4   u  132 1024    377     1.391   0.978     3.860

```

Switch(config)# show ntp associations detail <IP ADDR>

NTP association information

```

IP address      : 172.31.32.2           Peer Mode      : Server
Status         : Configured, Insane, Invalid Peer Poll Intvl : 64
Stratum        : 5                Root Delay     : 137.77 sec
Ref Assoc ID   : 0                Root Dispersion : 142.75
Association Name : NTP Association 0 Reach          : 376
Reference ID   : 16.93.49.4       Delay          : 4.23 sec
Our Mode       : Client           Offset         : -8.587 sec
Our Poll Intvl : 1024             Precision      : 2**19

```

```

Dispersion      : 1.62 sec
Association In Packets : 60
Association Out Packets : 60
Association Error Packets : 0
Origin Time     : Fri Jul 3 11:39:40 2015
Receive Time    : Fri Jul 3 11:39:44 2015
Transmit Time   : Fri Jul 3 11:39:44 2015

```

```

-----
Filter Delay = 4.23 4.14 2.41 5.95 2.37 2.33 4.26 4.33
Filter Offset = -8.59 -8.82 -9.91 -8.42 -10.51 -10.77 -10.13 -10.11

```

show ntp authentication

Syntax

Description

Show the authentication status and other information about the authentication key.

```
show ntp authentication
```

Switch(config)# show ntp authentication

NTP Authentication Information

```

Key-ID      Auth Mode   Trusted
-----
67          md5         yes
7           md5         no
1           sha1        yes
2           sha1        no

```


Validation rules

Validation	Error/Warning/Prompt
If access-list name is not valid.	Please enter a valid access-list name.
If the authentication method is being set to two-factor authentication, various messages display.	<p>If both the public key and username/password are not configured: Public key and username/password should be configured for a successful two-factor authentication.</p> <p>If public key is configured and username is not configured:</p> <p>Username and password should be configured for a successful two-factor authentication.</p> <p>If the username is configured and public key is not configured:</p> <p>Public key should be configured for a successful two-factor authentication.</p> <p>If "ssh-server" certificate is not installed at the time of enabling certificate-password authentication:</p> <p>The "ssh-server" certificate should be installed for a successful two-factor authentication.</p>
If the authentication method is set to two-factor while installing the public key, a message displays.	The client public keys without username will not be considered for the two-factor authentication for the SSH session.
If the username and the key installation user for that privilege do not match, a message displays and installation is not allowed.	The username in the key being installed does not match the username configured on the switch.
This will also happen when the authentication method is set for two-factor.	
If the maximum number of <username : TA profile> associations is reached for a given TA profile, a message displays.	Maximum number of username associations with a TA profile is 10.
If secondary authentication type for two-factor authentication chosen is not "none", a message displays.	Not legal combination of authentication methods.
If the authentication method is anything other than two-factor and the two-factor authentication method options are set, a message displays.	Not legal combination of authentication methods.
If two-factor authentication is set and user tries to SSH into another system using <code>ssh <ip hostname></code> command, a message displays.	SSH client is not supported when the two-factor authentication is enabled.

Table Continued

Validation	Error/Warning/Prompt
If timeSync is in SNTP or Timep when NTP is enabled.	Timesync is not configured to NTP.
If timesync is NTP and NTP is enabled and we try to change timesync to SNTP.	Disable NTP before changing timesync to SNTP or TIMEP.
If we try to configure NTP servers more than the configured max-associations value.	The maximum number of NTP servers allowed is 2.
If we have 'n' NTP servers configured and we try to configure a max-associations value less than (n) number of NTP servers already configured.	Max-associations value cannot be less than the number of NTP servers configured.
If authentication key-id is not configured.	Authentication key-id %d has not been configured.
If key-id is not marked as trusted.	Key-id %d is not trusted.
If min poll value is more than max poll value.	NTP max poll value should be more than min poll value.
If ipv6 is not enabled on vlan interface.	IPv6 address not configured on the VLAN.

Event log messages

Cause

Event	Message
RMON_AUTH_TWO_FACTOR_AUTHEN_STATUS	<p>W 01/01/15 18:24:03 03397: auth: %s.</p> <p>Examples:</p> <p>W 01/01/15 18:24:03 03397: auth: Public key and username/password should be configured for the successful two-factor authentication.</p> <p>W 01/01/15 18:24:03 03397: auth: Username and password should be configured for the successful two-factor authentication.</p> <p>W 01/01/15 18:24:03 03397: auth: Public key should be configured for the successful two-factor authentication.</p> <p>I 01/01/15 18:24:03 03397: auth: The validation of certificate of SSH user 'user1' is successful.</p>
RMON_SSH_KEY_TWO_FACTOR_EN	<p>W 01/01/15 18:24:03 03399: ssh: %s.</p> <p>Examples:</p> <p>W 01/01/15 18:24:03 03399: ssh: The client public keys without username will not be considered for the two-factor authentication for SSH session.</p> <p>W 01/01/15 18:24:03 03399: ssh: The privilege level for the user with the SSH key conflicts with the user configured.</p>
RMON_SSH_TWO_FACTOR_AUTH_FAIL	<p>W 01/01/15 18:24:03 03398: ssh: %s.</p> <p>Examples:</p> <p>W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed due to the failure in public key authentication.</p> <p>W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed due to the failure in username/password authentication.</p> <p>W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed due to the failure in validating the client certificate.</p> <p>W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed as "ssh-server" certificate is not installed.</p>
When NTP client enabled.	NTP client is enabled.
When NTP client disabled.	NTP client is disabled.

Table Continued

Event	Message
When NTP found a new broadcast server.	A new broadcast server at %s.
When system clock was updated with new time.	The system clock time was changed by %ld sec %lu nsec. The new time is %s.
When NTP stratum was updated.	The NTP Stratum was changed from %d to %d.
When all NTP associations are cleared.	All the NTP server associations are reset.
When server is not reachable.	The NTP Server 10.1.1.2 is unreachable. (2 times in 60 seconds)
When MD5/SHA1 authentication failed.	The MD5 authentication on the NTP packet failed. The SHA1 authentication on the NTP packet failed.

Viewing port status and configuring port parameters

Connecting transceivers to fixed-configuration devices

If the switch either fails to show a link between an installed transceiver and another device or demonstrates errors or other unexpected behavior on the link, check the port configuration on both devices for a speed and/or duplex (mode) mismatch.

- To check the mode setting for a port on the switch, use either the Port Status screen in the menu interface or `show interfaces brief` in the CLI (see **Viewing port status and configuration (CLI)**).
- To display information about the transceivers installed on a switch, enter the `show tech receivers` command in the CLI (**The show tech transceivers command** on page 61).

Viewing port configuration (Menu)

The menu interface displays the configuration for ports and (if configured) any trunk groups.

From the Main Menu, select:

1. Status and Counters 4. Port Status

A switch port status screen

```

=====-- CONSOLE - MANAGER MODE -----
                Status and Counters - Port Status
-----
Port      Type      Intrusion
Alert     Enabled  Status   Mode     MDI      Flow   Bcast
-----
1         100/1000T No       Yes      Down     100FDx   Auto   off   0
2         100/1000T No       Yes      Down     1000FDx  Auto   off   0
3         100/1000T No       Yes      Down     1000FDx  Auto   off   0
4         100/1000T No       Yes      Down     1000FDx  Auto   off   0
5         100/1000T No       Yes      Down     1000FDx  Auto   off   0
6         100/1000T No       Yes      Down     1000FDx  Auto   off   0
7         100/1000T No       Yes      Down     1000FDx  Auto   off   0
8         100/1000T No       Yes      Down     1000FDx  Auto   off   0
9         100/1000T No       Yes      Down     1000FDx  Auto   off   0
10        100/1000T No       Yes      Down     1000FDx  Auto   off   0
11        100/1000T No       Yes      Down     1000FDx  Auto   off   0

Actions->  Back   Intrusion log   Help
    
```

Return to previous screen.

Use up/down arrow keys to scroll to other entries, left/right arrow keys to change action selection, and <Enter> to execute action.


Port Type	The port Type field represents the IEEE or other industry protocol in operation on that port. For example, 1000Base-SX is a gigabit protocol for gigabit operation over fiber optic cable.
Status of Ports	<p>A port can be enabled or disabled:</p> <ul style="list-style-type: none"> • Yes: Enabled, the default. This indicates the port is ready for a network connection. • No: Disabled, the port will not operate, even if properly connected to a network. Use the setting, For example, to shut the port down for diagnostic purposes or while you are making topology changes. <p>The status of a port can be up or down (Read-only):</p> <p>Up: The port senses a link beat.</p> <p>Down: The port is not enabled, has no cables connected, or is experiencing a network error. For troubleshooting information, see the installation and getting started guide for your switch.</p>
Flow Control	<p>With the port mode set to Auto (the default) and flow control on (enabled), the switch negotiates flow control on the indicated port. If the port mode is not set to Auto, or if flow control is off (disabled) on the port, then flow control is not used. Flow control must be enabled on both ends of a link.</p> <ul style="list-style-type: none"> • On: Enabled. The port uses 802.3x Link Layer Flow Control, generates flow control packets, and processes received flow control packets. • Off: Disabled (default). The port does not generate flow control packets, and drops any flow control packets it receives.
Broadcast Limit	<p>The broadcast limit specifies the percentage of the theoretical maximum network bandwidth that can be used for broadcast and multicast traffic. Any broadcast or multicast traffic exceeding that limit will be dropped. Zero (0) means the feature is disabled.</p> <p>The <code>broadcast-limit</code> command operates at the port context level to set the broadcast limit for a port on a switch.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  <p>This feature is not appropriate for networks that require high levels of IPX or RIP broadcast traffic.</p> </div>
Modes	The mode is the port's speed and duplex (data transfer operation) setting. The following table shows possible modes available, depending on the port type (copper or fiber) and port speed.

Table 3: Supported Modes

Mode	Speed and Duplex Settings
Auto-MDIX	Senses speed and negotiates with the port at the other end of the link for port operation (MDI-X or MDI). To see what the switch negotiates for the Auto setting, use the CLI <code>show interfaces brief</code> command or the menu commands "1. Status and Counters", "3. Port Status". This feature applies only to copper port switches using twisted-pair copper Ethernet cables.
MDI	Sets the port to connect with a PC using a crossover cable (Manual mode— applies only to copper port switches using twisted-pair copper Ethernet cables).
MDIX	Sets the port to connect with a PC using a straight-through cable (Manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables).

Table Continued

Mode	Speed and Duplex Settings
Auto-10	Allows the port to negotiate between half-duplex (HDx) and full-duplex (FDx) while keeping speed at 10Mbps. Also negotiates flow control (enabled or disabled). Hewlett Packard Enterprise recommends Auto-10 for links between 10/100 auto-sensing ports connected with Cat 3 cabling. (Cat 5 cabling is required for 100 Mbps links.)
Auto-100	Uses 100 Mbps and negotiates with the port at the other end of the link for other port operation features.
Auto-10-100	Allows the port to establish a link with the port at the other end at either 10 Mbps or 100 Mbps, using the highest mutual speed and duplex mode available. Only these speeds are allowed with this setting.
Auto-1000	Uses 1000 Mbps and negotiates with the port at the other end of the link for other port operation features.
10HDx	Uses 10 Mbps, Half-Duplex
100HDx	Uses 100 Mbps, Half-Duplex
10FDX	Uses 10 Mbps, Full-Duplex
100FDx	Uses 100 Mbps, Full-Duplex
1000FDx	Uses 1000 Mbps, Full-Duplex
10 GbE FDx	Uses 10 Gigabits/sec Full-Duplex

The next two tables display the protocols and modes supported for copper ports and fiber optic ports, respectively.

Table 4: Protocols and modes supported for copper ports

10/100 Mbps		Gigabit		10 Gigabit	
10/100 TX		10/100/1000-T		10GBASE-CX4	
Modes	Settings	Modes	Settings	Modes	Settings
Auto	100FDx	Auto	1000FDx	Auto	10 Gigabit FDx
10HDx	10HDx	Auto-10	10HDx		
100HDx	100HDx	Auto-100	100FDx		
10FDx	10FDx	Auto-10-100	100FDx		
100FDx	100FDX	Auto-1000	1000FDx		
		10HDx	10HDx		
		100HDX	100HDX		

Table 5: Protocols and modes supported for fiber optic ports

	100 Mbps	Gigabit	10 Gigabit
Protocols	100BASE-FX 100BASE-BX10	1000BASE-SX 1000BASE-LX 1000BASE-BX 101000BASE-LH	10GBASE-SR 10GBASE-LR 10GBASE-LRM
Modes	100HDx 100FDx	Auto1000FDx	Auto

Configuring ports (Menu)

The menu interface uses the same screen for configuring both individual ports and port trunk groups. For information on port trunk groups, see the chapter on "Port Trunking".

Procedure

1. From the Main Menu, select:

2. Switch Configuration...

2. Port/Trunk Settings

Port/trunk settings with a trunk group configured

```

===== TELNET - MANAGER MODE =====
Switch Configuration - Port/Trunk Settings

Port      Type      Enabled  Mode      Flow Ctrl  Group  Type
-----
A1        1000T    | Yes    Auto-10-100  Disable
A2        1000T    | Yes    Auto-10-100  Disable
A3        1000T    | Yes    Auto         Disable
A3        1000T    | Yes    Auto         Disable
A4        1000T    | Yes    Auto         Disable
A5        1000T    | Yes    Auto         Disable
A6        1000T    | Yes    Auto         Disable
A7        1000T    | Yes    Auto         Disable  Trk1  Trunk
A8        1000T    | Yes    Auto         Disable  Trk2  Trunk

Actions->  Cancel  Edit   Save   Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute
action.

```

2. Press [E] (for Edit).

The cursor moves to the Enabled field for the first port.

For further information on configuration options for these features, see the online help provided with this screen.

3. When you have finished making changes to the above parameters, press [Enter], then press [S] (for Save).

Viewing port status and configuration (CLI)

Use the following commands to display port status and configuration data.

Syntax:

```
show interfaces [brief | config | < port-list >]
```

brief Lists the current operating status for all ports on the switch.

config Lists a subset of configuration data for all ports on the switch; that is, for each port, the display shows whether the port is enabled, the operating mode, and whether it is configured for flow control.

<port-list> Shows a summary of network traffic handled by the specified ports.

The show interfaces brief command listing

```
switch(config)# show interfaces brief
Status and Counters - Port Status
```

Port	Type	Intrusion Alert	Enabled	Status	Mode	MDI Mode	Flow Ctrl	Bcast Limit
B1	100/1000T	No	Yes	Down	Auto-10-100	Auto	off	0
B2	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B3	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B4	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B5	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B6	100/1000T	No	Yes	Down	1000FDx	Auto	off	0

The show interfaces config command listing

```
switch(config)# show interfaces config
```

Port Settings

Port	Type	Enabled	Mode	Flow Ctrl	MDI
B1	100/1000T	Yes	Auto-10-100	Disable	Auto
B2	100/1000T	Yes	Auto	Disable	Auto
B3	100/1000T	Yes	Auto	Disable	Auto
B4	100/1000T	Yes	Auto	Disable	Auto
B5	100/1000T	Yes	Auto	Disable	Auto
B6	100/1000T	Yes	Auto	Disable	Auto

Dynamically updating the show interfaces command (CLI/Menu)

Syntax:

```
show interfaces display
```

Uses the `display` option to initiate the dynamic update of the `show interfaces` command, with the output being the same as the `show interfaces` command.



Select **Back** to exit the display.

Example:

```
switch# show interfaces display
```

When using the **display** option in the CLI, the information stays on the screen and is updated every 3 seconds, as occurs with the display using the menu feature. The update is terminated with **Cntl-C**.

You can use the arrow keys to scroll through the screen when the output does not fit in one screen.

Figure 6: *show interfaces display command with dynamically updating output*

Status and Counters - Port Counters							
Port	Total Bytes	Total Frames	Errors Rx	Drops Tx	Flow Ctrl	Bca>	Lim>
1	2,164,277	20,366	0	0	off	0	0
2	0	0	0	0	off	0	0
3	0	0	0	0	off	0	0
4	0	0	0	0	off	0	0
5	0	0	0	0	off	0	0
6	0	0	0	0	off	0	0
7	0	0	0	0	off	0	0
8	0	0	0	0	off	0	0
9	0	0	0	0	off	0	0
10	0	0	0	0	off	0	0
11	0	0	0	0	off	0	0

Actions-> **Back** Show details Reset Help

Use up/down arrow keys to scroll to other entries, left/right arrow keys to change action selection, and <Enter> to execute action.

Customizing the show interfaces command (CLI)

You can create `show` commands displaying the information that you want to see in any order you want by using the `custom` option.

Syntax:

```
show interfaces custom [port-list] column-list
```

Select the information that you want to display. Supported columns are shown in the table below.

Table 6: *Supported columns, what they display, and examples:*

Parameter column	Displays	Examples
port	Port identifier	A2
type	Port type	100/1000T
status	Port status	up or down

Table Continued

Parameter column	Displays	Examples
speed	Connection speed and duplex	1000FDX
mode	Configured mode	auto, auto-100, 100FDX
mdi	MDI mode	auto, MDIX
flow	Flow control	on or off
name	Friendly port name	
vlanid	The vlan id this port belongs to, or "tagged" if it belongs to more than one vlan	4tagged
enabled	port is or is not enabled	yes or nointrusion
intrusion	Intrusion alert status	no
bcast	Broadcast limit	0

The custom show interfaces command

```
switch(config)# show int custom 1-4 port name:4 type vlan intrusion speed enabled mdi
```

Status and Counters - Custom Port Status

Port	Name	Type	VLAN	Intrusion Alert	Speed	Enabled	MDI-mode
1	Acco	100/1000T	1	No	1000FDx	Yes	Auto
2	Huma	100/1000T	1	No	1000FDx	Yes	Auto
3	Deve	100/1000T	1	No	1000FDx	Yes	Auto
4	Lab1	100/1000T	1	No	1000FDx	Yes	Auto

You can specify the column width by entering a colon after the column name, then indicating the number of characters to display. In the above example, the Name column displays only the first four characters of the name. All remaining characters are truncated.



Each field has a fixed minimum width to be displayed. If you specify a field width smaller than the minimum width, the information is displayed at the minimum width. For example, if the minimum width for the Name field is 4 characters and you specify Name:2, the Name field displays 4 characters.

You can enter parameters in any order. There is a limit of 80 characters per line; if you exceed this limit an error displays.

Error messages associated with the show interfaces command

The following table provides information on error messages associated with the `show interfaces custom` command.

Error	Error message
Requesting too many fields (total characters exceeds 80)	Total length of selected data exceeds one line
Field name is misspelled	Invalid input: <i><input></i>
Mistake in specifying the port list	Module not present for port or invalid port: <i><input></i>
The port list is not specified	Incomplete input: custom

Note on using pattern matching with the `show interfaces custom` command

If you have included a pattern matching command to search for a field in the output of the `show int custom` command, and the `show int custom` command produces an error, the error message may not be visible and the output is empty. For example, if you enter a command that produces an error (such as `vlan` is misspelled) with the pattern matching `include` option, the output may be empty:

```
HP Switch(config)# show int custom 1-3 name vlun | include vlan1
```

It is advisable to try the `show int custom` command first to ensure there is output, and then enter the command again with the pattern matching option.

Note that in the above command, you can substitute `int` for `interface`; that is: `show int custom`.

Viewing port utilization statistics (CLI)

Use the `show interface port-utilization` command to view a real-time rate display for all ports on the switch. The example below shows a sample output from this command.

A `show interface port-utilization` command listing

```
switch(config)# show interfaces port-utilization
Status and Counters - Port Utilization
```

Port	Mode	Rx			Tx		
		Kbits/sec	Pkts/sec	Util	Kbits/sec	Pkts/sec	Util
B1	1000FDx	0	0	0	0	0	0
B2	1000FDx	0	0	0	0	0	0
B3	1000FDx	0	0	0	0	0	0
B4	1000FDx	0	0	0	0	0	0
B5	1000FDx	0	0	0	0	0	0
B6	1000FDx	0	0	0	0	0	0
B7	100FDx	624	86	00.62	496	0	00.49

Operating notes for viewing port utilization statistics

- For each port on the switch, the command provides a real-time display of the rate at which data is received (Rx) and transmitted (Tx) in terms of kilobits per second (KBits/s), number of packets per second (Pkts/s), and utilization (Util) expressed as a percentage of the total bandwidth available.
- The `show interfaces <port-list>` command can be used to display the current link status and the port rate average over a 5 minute period. Port rates are shown in bits per second (bps) for ports up to 1 Gigabit; for 10 Gigabit ports, port rates are shown in kilobits per second (Kbps).

Viewing transceiver status (CLI)

The `show interfaces transceivers` command allows you to:

- Remotely identify transceiver type and revision number without having to physically remove an installed transceiver from its slot.
- Display real-time status information about all installed transceivers, including non-operational transceivers.

The example shows sample output from the `show tech transceivers` command.



Part # column below enables you to determine the manufacturer for a specified transceiver and revision number.

The show tech transceivers command

```
switch# show tech transceivers
```

```
Transceiver Technical Information:
```

Port #	Type	Prod #	Serial #	Part #
21	1000SX	J4858B	CN605MP23K	
22	1000LX	J4859C	H11E7X	2157-2345
23	??	??	non operational	
25	10GbE-CX4	J8440A	US509RU079	
26	10GbE-CX4	J8440A	US540RU002	
27	10GbE-LR	J8437B	PPA02-2904:0017	2157-2345
28	10GbE-SR	J8436B	01591602	2158-1000
29	10GbE-ER	J8438A	PPA03-2905:0001	

```
The following transceivers may not function correctly:
```

Port #	Message
Port 23	Self test failure.

Operating Notes

The following information is displayed for each installed transceiver:

- Port number on which transceiver is installed.
- Type of transceiver.
- Product number — Includes revision letter, such as A, B, or C. If no revision letter follows a product number, this means that no revision is available for the transceiver.
- Part number — Allows you to determine the manufacturer for a specified transceiver and revision number.
- For a non-HPE switches installed transceiver (see [line 23 of "The show tech transceivers command" example](#)), no transceiver type, product number, or part information is displayed. In the Serial Number field, non-operational is displayed instead of a serial number.
- The following error messages may be displayed for a non-operational transceiver:
 - Unsupported Transceiver. (SelfTest Err#060)
 - This switch only supports revision B and above transceivers.
 - Self test failure.
 - Transceiver type not supported in this port.
 - Transceiver type not supported in this software version.
 - Not an HPE Switch Transceiver.

Enabling or disabling ports and configuring port mode (CLI)

You can configure one or more of the following port parameters.

Syntax:

```
[no] interface <port-list> [<disable|enable>]
```

Disables or enables the port for network traffic. Does not use the `no` form of the command. (Default: `enable`.)

```
speed-duplex [<auto-10|10-full|10-half|100-full|100-half|auto|auto-100|1000-full>]
```

Note that in the above Syntax:, you can substitute `int` for `interface` (for example, `int <port-list>`).

Specifies the port's data transfer speed and mode. Does not use the `no` form of the command. (Default: `auto`.)

The 10/100 auto-negotiation feature allows a port to establish a link with a port at the other end at either 10 Mbps or 100 Mbps, using the highest mutual speed and duplex mode available. Only these speeds are allowed with this setting.

Examples:

To configure port C5 for auto-10-100, enter this command:

```
switch(config)# int c5 speed-duplex auto-10-100
```

To configure ports C1 through C3 and port C6 for 100Mbps full-duplex, enter these commands:


```
switch(config)# int c1-c3,c6 speed-duplex 100-full
```

Similarly, to configure a single port with the above command settings, you could either enter the same command with only the one port identified or go to the context level for that port and then enter the command. For example, to enter the context level for port C6 and then configure that port for 100FDx:

```
switch(config)# int e c6
switch(eth-C6)# speed-duplex 100-full
```

If port C8 was disabled, and you wanted to enable it and configure it for 100FDx with flow-control active, you could do so with either of the following command sets:

Figure 7: Two methods for changing a port configuration



For more on flow control, see [Enabling or disabling flow control \(CLI\)](#) on page 63.

Enabling or disabling flow control (CLI)



You must enable flow control on both ports in a given link. Otherwise, flow control does not operate on the link and appears as `Off` in the `show interfaces brief` port listing, even if flow control is configured as enabled on the port in the switch. (See [The show interfaces brief command listing](#) example.) Also, the port (speed-duplex) mode must be set to `Auto` (the default).

To disable flow control on some ports, while leaving it enabled on other ports, just disable it on the individual ports you want to exclude.

Syntax:

```
[no] interface <port-list> flow-control
```

Enables or disables flow control packets on the port. The `no` form of the command disables flow control on the individual ports. (Default: Disabled.)

Examples:

Suppose that:

1. You want to enable flow control on ports A1-A6.
2. Later, you decide to disable flow control on ports A5 and A6.
3. As a final step, you want to disable flow control on all ports.

Assuming that flow control is currently disabled on the switch, you would use these commands:

Figure 8: Configuring flow control for a series of ports

```
switch(config)# int a1-a6 flow-control
```

```
switch(config)# show interfaces brief
```

Status and Counters - Port Status

Port	Type	Intrusion		Enabled	Status	MDI Flow Bcast		
		Alert				Mode	Mode	Ctrl Limit
A1	10GbE-T	No	Yes	Yes	Up	1000FDx	NA	on 0
A2	10GbE-T	No	Yes	Yes	Up	10GigFD	NA	on 0
A3	10GbE-T	No	Yes	Yes	Up	10GigFD	NA	on 0
A4	10GbE-T	No	Yes	Yes	Up	10GigFD	NA	on 0
A5	10GbE-T	No	Yes	Yes	Up	10GigFD	NA	on 0
A6	10GbE-T	No	Yes	Yes	Up	10GigFD	NA	on 0
A7	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0
A8	10GbE-T	No	Yes	Yes	Up	10GigFD	NA	off 0

```
switch(config)# no int a5-a6 flow-control
```

```
switch(config)# show interfaces brief
```

Status and Counters - Port Status

Port	Type	Intrusion		Enabled	Status	MDI Flow Bcast		
		Alert				Mode	Mode	Ctrl Limit
A1	10GbE-T	No	Yes	Yes	Up	1000FDx	NA	on 0
A2	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	on 0
A3	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	on 0
A4	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	on 0
A5	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0
A6	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0
A7	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0
A8	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0

```
switch(config)# no int a1-a4 flow-control
```

```
switch(config)# show interfaces brief
```

Status and Counters - Port Status

Port	Type	Intrusion		Enabled	Status	MDI Flow Bcast		
		Alert				Mode	Mode	Ctrl Limit
A1	10GbE-T	No	Yes	Yes	Down	1000FDx	NA	off 0
A2	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0
A3	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0
A4	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0
A5	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0
A6	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0
A7	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0
A8	10GbE-T	No	Yes	Yes	Down	10GigFD	NA	off 0

Configuring a broadcast limit

Broadcast-Limit on switches covered in this guide is configured on a per-port basis. You must be at the port context level for this command to work, for example:

```
switch(config)#int 1
switch(int 1)# broadcast-limit 1
```

Broadcast-limit

Syntax:

```
broadcast-limit <0-99>
```

Enables or disables broadcast limiting for outbound broadcasts on a selected port on the switch. The value selected is the percentage of traffic allowed, for example, `broadcast-limit 5` allows 5% of the maximum amount of traffic for that port. A value of zero disables broadcast limiting for that port.



You must switch to port context level before issuing the `broadcast-limit` command.

This feature is not appropriate for networks requiring high levels of IPX or RIP broadcast traffic.

Syntax:

```
show config
```

Displays the startup-config file. The broadcast limit setting appears here if enabled and saved to the startup-config file.

Syntax:

```
show running-config
```

Displays the running-config file. The broadcast limit setting appears here if enabled. If the setting is not also saved to the startup-config file, rebooting the switch returns broadcast limit to the setting currently in the startup-config file.

For example, the following command enables broadcast limiting of 1 percent of the traffic rate on the selected port on the switch:

```
switch(int 1)# broadcast-limit 1
```

For a one Gbps port this results in a broadcast traffic rate of ten Mbps.

Port shutdown with broadcast storm

A LAN broadcast storm arises when an excessively high rate of broadcast packets flood the LAN. Occurrence of LAN broadcast storm disrupts traffic and degrades network performance. To prevent LAN traffic from being disrupted, an enhancement of fault-finder commands adds new options, and the corresponding MIBs, that trigger a port disablement when a broadcast storm is detected on that port.

Under this enhancement, the CLI commands given only supports broadcast traffic and not multicast and unicast types of traffic.

The waiting period range for re-enabling ports is 0 to 604800 seconds. The default waiting period to re-enable a port is zero which prevents the port from automatic re-enabling.



Avoid port flapping when choosing the waiting period by considering the time to re-enable carefully.

Use the following commands to configure the broadcast-storm on a port.

Syntax:

```
[no] fault-finder broadcast-storm [ethernet] <port-list> action [warn|warn-and-disable <seconds>] [percent <percent>|pps <rate>]
```

To remove the current configuration of broadcast-storm on a port, use:

Syntax:

```
no fault-finder broadcast-storm [ethernet] <port-list>
```

broadcast-storm	Configure broadcast storm control.
pps	Rising threshold level in number of broadcast packets per second.
percent	Rising threshold level as a percentage of bandwidth of the port. The percentage is calculated on 64 byte packet size.
warn	Log the event only.
warn-and-disable	Log the event and disable the port.
seconds	Re-enable the port after waiting for the specified number of seconds. Default is not to re-enable.

Configuration examples:

```
switch(config)# fault-finder broadcast-storm [ethernet] <A1> action [warn-and-disable <65535>]< percent 10>
```

```
switch(config)# fault-finder broadcast-storm [ethernet] <A2> action [warn-and-disable <pps 100>
```

```
switch(config)# fault-finder broadcast-storm [ethernet] <A22> action [warn] <pps 100>
```

Viewing broadcast storm

Use the following command to display the broadcast-storm-control configuration.

Syntax:

```
show fault-finder broadcast-storm [[ethernet] port-list]
```

Examples:

```
switch# show fault-finder broadcast-storm [A1]
```

Port	Bcast Storm	Port Status	Rising Threshold	Action	Disable Timer	Disable Timer Left
A1	Yes	Down	10%	warn-and-disable	65535	—

switch (config)# show fault-finder broadcast-storm

Port	Bcast Storm	Port Status	Rising Threshold	Action	Disable Timer	Disable Timer Left
A1	Yes	Down	200 pps	warn-and-disable	10	9

switch (config)# show fault-finder broadcast-storm A1

Port	Bcast Storm	Port Status	Rising Threshold	Action	Disable Timer	Disable Timer Left
A1	No	Up	—	none	—	—

switch (config)# show fault-finder broadcast-storm

Port	Bcast Storm	Port Status	Rising Threshold	Action	Disable Timer	Disable Timer Left
A1	Yes	Up	75%	warn	—	—

SNMP MIB

SNMP support will be provided through the following MIB objects:

hpicfFfBcastStormControlPortConfig OBJECT IDENTIFIER

:: = { hpicfFaultFinder 5 }

hpicfFfBcastStormControlPortConfigTable OBJECT-TYPE

- syntax sequence: **HpicfFfBcastStormControlPortConfigEntry**
- max-access: not-accessible
- status: current
- description: This table provides information about broadcast storm control configuration of all ports. ::= {hpicfFfBcastStormControlPortConfig 1}

hpicfFfBcastStormControlPortConfigEntry OBJECT-TYPE

- syntax **HpicfFfBcastStormControlPortConfigEntry**
- max-access: not-accessible
- status: current
- description: This object provides information about broadcast storm control configuration of each port.
- index: {hpicfffbcaststormcontrolportindex} ::= {hpicfFfBcastStormControlPortConfigTable 1}

hpicfFfBcastStormControlPortConfigEntry ::=

- Syntax sequence: hpicfFfBcastStormControlPortIndex InterfaceIndex,

hpicfFfBcastStormControlMode Integer,
hpicfFfBcastStormControlRisingpercent Integer32,
hpicfFfBcastStormControlRisingpps Integer32,
hpicfFfBcastStormControlAction Integer,
hpicfFfBcastStormControlPortDisableTimer Unsigned32

hpicfFfBcastStormControlPortIndex OBJECT-TYPE

- Syntax: Interfaceindex
- max-access: not-accessible
- status: current
- description: The Index Value Which Uniquely Identifies A Row In The Interfaces Table.

::= {**hpicfFfBcastStormControlPortConfigEntry 1**}

hpicfFfBcastStormControlMode OBJECT-TYPE

- Syntax Integer: disabled(1), **Bcastrisinglevelpercent**(2), **Bcastrisinglevelpps**(3)
- max-access: read-write
- status: current
- description: The broadcast storm control mode of a port. A value of disable (1) indicates that no rising threshold value is set for broadcast storm traffic on this port. A value of **bcastrisinglevelpercent** (2) indicates that the rising threshold rate for broadcast storm traffic is configured in percentage of port bandwidth. A value of **bcastrisinglevelpps** (3) indicates that the rising threshold rate for broadcast storm traffic is configured in packets per second.
- DEFVAL: disabled

::= {**hpicfFfBcastStormControlPortConfigEntry 2**}

hpicfFfBcastStormControlRisingpercent OBJECT-TYPE

- Syntax Integer32 (1..100)
- max-access: read-write
- status: current
- description: This Is The Rising Threshold Level in percent of bandwidth of the port. **hpicfFfBcastStormControlAction** occurs when broadcast traffic reaches this level.

::= {**hpicfFfBcastStormControlPortConfigEntry 3**}

hpicfFfBcastStormControlRisingpps OBJECT-TYPE

- Syntax Integer32 (1..10000000)
- max-access: read-write
- status: current
- description: This object indicates the rising threshold for broadcast storm control. This value is in packets-per-second of received broadcast traffic. **hpicfffbcaststormcontrolaction** object takes action when broadcast traffic reaches this level.

::= {**hpicfFfBcastStormControlPortConfigEntry 4**}

hpicfFfBcastStormControlAction OBJECT-TYPE

- Syntax integer: none(1), warn(2), warnanddisable(3)
- max-access: read-write
- status: current
- Description: This object defines the action taken by the switch when a broadcast storm occurs on a port. A value of none (1) indicates that no action is performed. A value of warn (2) indicates that an event is logged when broadcast traffic crosses the threshold value set on that port. A value of warn-and-disable (3) indicates

that the port is disabled and an event is logged as soon as the broadcast traffic reaches the threshold value set on that port.

- DEFVAL: none

::= {hpicfFfBcastStormControlPortConfigEntry 5}

hpicfFfBcastStormControlPortDisableTimer OBJECT-TYPE

- Syntax Unsigned32 (0..604800)
- Units: seconds
- max-access: read-write
- status: current
- Description: This object specifies the time period for which the port remains in disabled state. A port is disabled when broadcast traffic reaches the threshold value set on that port. This time period is specified in seconds. The default value is zero which means that the port remains disabled and is not enabled again.
- DEFVAL {0}

::= {hpicfFfBcastStormControlPortConfigEntry 6}

Configuring auto-MDIX

Copper ports on the switch can automatically detect the type of cable configuration (MDI or MDI-X) on a connected device and adjust to operate appropriately.

This means you can use a "straight-through" twisted-pair cable or a "crossover" twisted-pair cable for any of the connections—the port makes the necessary adjustments to accommodate either one for correct operation. The following port types on your switch support the IEEE 802.3ab standard, which includes the "Auto MDI/MDI-X" feature:

- 10/100-TX xl module ports
- 100/1000-T xl module ports
- 10/100/1000-T xl module ports

Using the above ports:

- If you connect a copper port using a straight-through cable on a switch to a port on another switch or hub that uses MDI-X ports, the switch port automatically operates as an MDI port.
- If you connect a copper port using a straight-through cable on a switch to a port on an end node—such as a server or PC—that uses MDI ports, the switch port automatically operates as an MDI-X port.

Auto-MDIX was developed for auto-negotiating devices, and was shared with the IEEE for the development of the IEEE 802.3ab standard. Auto-MDIX and the IEEE 802.3ab Auto MDI/MID-X feature are completely compatible. Additionally, Auto-MDIX supports operation in forced speed and duplex modes.

For more information on this subject, see the IEEE 802.3ab standard reference. For more information on MDI-X, the installation and getting started guide for your switch.

Manual override

If you require control over the MDI/MDI-X feature, you can set the switch to either of these non-default modes:

- Manual MDI
- Manual MDI-X

The table below shows the cabling requirements for the MDI/MDI-X settings.

Table 7: Cable types for auto and manual MDI/MDI-X settings

Setting	MDI/MDI-X device type	
	PC or other MDI device type	Switch, hub, or other MDI-X device
Manual MDI	Crossover cable	Straight-through cable
Manual MDI-X	Straight-through cable	Crossover cable
Auto-MDI-X (the default)	Either crossover or straight-through cable	

The AutoMDIX features apply only to copper port switches using twisted-pair copper Ethernet cables.

Configuring auto-MDIX (CLI)

The auto-MDIX features apply only to copper port switches using twisted-pair copper Ethernet cables. For information about auto-MDIX, see [Configuring auto-MDIX](#) on page 69.

Syntax:

```
interface <port-list> mdix-mode < {auto-mdix | mdi | mdix}>
```

auto-mdix	The automatic, default setting. This configures the port for automatic detection of the cable (either straight-through or crossover).
mdi	The manual mode setting that configures the port for connecting to either a PC or other MDI device with a crossover cable, or to a switch, hub, or other MDI-X device with a straight-through cable.
mdix	The manual mode setting that configures the port for connecting to either a switch, hub, or other MDI-X device with a crossover cable, or to a PC or other MDI device with a straight-through cable.

Syntax:

```
show interfaces config
```

Lists the current per-port Auto/MDI/MDI-X configuration.

Syntax:

```
show interfaces brief
```

- Where a port is linked to another device, this command lists the MDI mode the port is currently using.
- In the case of ports configured for `Auto` (`auto-mdix`), the MDI mode appears as either `MDI` or `MDIX`, depending upon which option the port has negotiated with the device on the other end of the link.
- In the case of ports configured for `MDI` or `MDIX`, the mode listed in this display matches the configured setting.
- If the link to another device was up, but has gone down, this command shows the last operating MDI mode the port was using.
- If a port on a given switch has not detected a link to another device since the last reboot, this command lists the MDI mode to which the port is currently configured.

The `show interfaces config` displays the following data when port A1 is configured for `auto-mdix`, port A2 is configured for `mdi`, and port A3 is configured for `mdix`:

Displaying the current MDI configuration

```
switch(config)# show interfaces config
```

Port Settings

Port	Type	Enabled	Mode	Flow Ctrl	MDI
A1	10GbE-T	Yes	Auto	Disable	Auto
A2	10GbE-T	Yes	Auto	Disable	MDI
A3	10GbE-T	Yes	Auto	Disable	MDIX
A4	10GbE-T	Yes	Auto	Disable	Auto
A5	10GbE-T	Yes	Auto	Disable	Auto
A6	10GbE-T	Yes	Auto	Disable	Auto
A7	10GbE-T	Yes	Auto	Disable	Auto
A8	10GbE-T	Yes	Auto	Disable	Auto

Displaying the current MDI operating mode

```
switch(config)# show interfaces brief
```

Status and Counters - Port Status

Port	Type	Intrusion Alert	Enabled	Status	Mode	MDI Mode	Flow Ctrl	Bcast Limit
A1	10GbE-T	No	Yes	Up	1000FDx	MDIX	off	0
A2	10GbE-T	No	Yes	Down	10GigFD	MDI	off	0
A3	10GbE-T	No	Yes	Down	10GigFD	MDIX	off	0
A4	10GbE-T	No	Yes	Down	10GigFD	Auto	off	0
A5	10GbE-T	No	Yes	Down	10GigFD	Auto	off	0
A6	10GbE-T	No	Yes	Down	10GigFD	Auto	off	0
A7	10GbE-T	No	Yes	Down	10GigFD	Auto	off	0
A8	10GbE-T	No	Yes	Down	10GigFD	Auto	off	0

Using friendly (optional) port names

This feature enables you to assign alphanumeric port names of your choosing to augment automatically assigned numeric port names. This means you can configure meaningful port names to make it easier to identify the source of information listed by some `show` commands. (Note that this feature **augments** port numbering, but **does not replace** it.)

Configuring and operating rules for friendly port names

- At either the global or context configuration level, you can assign a unique name to a port. You can also assign the same name to multiple ports.
- The friendly port names you configure appear in the output of the `show name [port-list]`, `show config`, and `show interface <port-number>` commands. They do not appear in the output of other `show` commands or in Menu interface screens. (See [Displaying friendly port names with other port data \(CLI\)](#) on page 73.)
- Friendly port names are not a substitute for port numbers in CLI commands or Menu displays.
- Trunking ports together does not affect friendly naming for the individual ports. (If you want the same name for all ports in a trunk, you must individually assign the name to each port.)
- A friendly port name can have up to 64 contiguous alphanumeric characters.

- Blank spaces within friendly port names are not allowed, and if used, cause an **invalid input** error. (The switch interprets a blank space as a name terminator.)
- In a port listing, **not assigned** indicates that the port does not have a name assignment other than its fixed port number.
- To retain friendly port names across reboots, you must save the current running-configuration to the startup-config file after entering the friendly port names. (In the CLI, use the `write memory` command.)

Configuring friendly port names (CLI)

For detailed information about friendly port names, see [Using friendly \(optional\) port names](#) on page 71.

Syntax:

```
interface <port-list> name <port-name-string>
```

Assigns a port name to port-list.

Syntax:

```
no interface <port-list> name
```

Deletes the port name from <port-list>.

Configuring a single port name (CLI)

Suppose that you have connected port A3 on the switch to Bill Smith's workstation, and want to assign Bill's name and workstation IP address (10.25.101.73) as a port name for port A3:

Configuring a friendly port name

```
switch(config)# int A3 name
Bill_Smith@10.25.101.73
switch(config)# write mem
switch(config)# show name A3
```

```
Port Names
Port : A3
Type : 10/100TX
```

Configuring the same name for multiple ports (CLI)

Suppose that you want to use ports A5 through A8 as a trunked link to a server used by a drafting group. In this case you might configure ports A5 through A8 with the name "Draft-Server:Trunk."

Configuring one friendly port name on multiple ports

```
switch(config)# int a5-a8 name Draft-Server:Trunk
switch(config)# write mem
switch(config)# show name a5-a8
```

```
Port Names

Port : A5
Type : 10GbE-T
Name : Draft-Server:Trunk

Port : A6
Type : 10GbE-T
```



```
Name : Draft-Server:Trunk

Port : A7
Type : 10GbE-T
Name : Draft-Server:Trunk

Port : A8
Type : 10GbE-T
Name : Draft-Server:Trunk
```

Displaying friendly port names with other port data (CLI)

You can display friendly port name data in the following combinations:

Syntax:

```
show name
```

Displays a listing of port numbers with their corresponding friendly port names and also quickly shows you which ports do not have friendly name assignments. (`show name` data comes from the running-config file.)

Syntax:

```
show interface <port-number>
```

Displays the friendly port name, if any, along with the traffic statistics for that port. (The friendly port name data comes from the running-config file.)

Syntax:

```
show config
```

Includes friendly port names in the per-port data of the resulting configuration listing. (`show config` data comes from the startup-config file.)

Listing all ports or selected ports with their friendly port names (CLI)

Syntax:

```
show name [port-list]
```

Lists the friendly port name with its corresponding port number and port type. The `show name` command without a port list shows this data for all ports on the switch.

Friendly port name data for all ports on the switch

```
switch(config)# show name
Port Names
  Port   Type           Name
-----
  A1     10GbE-T
  A2     10GbE-T
  A3     10GbE-T   Bill_Smith@10.25.101.73
  A4     10GbE-T
  A5     10GbE-T   Draft-Server:Trunk
  A6     10GbE-T   Draft-Server:Trunk
```

A7	10GbE-T	Draft-Server:Trunk
A8	10GbE-T	Draft-Server:Trunk

Friendly port name data for specific ports on the switch

```
switch(config)# show name A3-A5
Port Names
  Port : A3
  Type : 10GbE-T
  Name : Bill_Smith@10.25.101.73
  Port : A4
  Type : 10GbE-T
  Name :
  Port : A5
  Type : 10GbE-T
  Name : Draft-Server:Trunk
```

Including friendly port names in per-port statistics listings (CLI)

Syntax:

```
show interface <port-number>
```

Includes the friendly port name with the port's traffic statistics listing. A friendly port name configured to a port is automatically included when you display the port's statistics output.

If you configure port A1 with the name "O'Connor_10.25.101.43," the `show interface` output for this port appears similar to the following:

A friendly port name in a per-port statistics listing

```
switch(config)# show interface a1
Status and Counters - Port Counters for port A1

Name      : O'Connor@10.25.101.43
MAC Address      : 001871-b995ff
Link Status      : Up
Totals (Since boot or last clear) :
  Bytes Rx      : 2,763,197          Bytes Tx      : 22,972
  Unicast Rx    : 2044              Unicast Tx    : 128
  Bcast/Mcast Rx : 23,456          Bcast/Mcast Tx : 26
Errors (Since boot or last clear) :
  FCS Rx       : 0                 Drops Tx      : 0
  Alignment Rx : 0                 Collisions Tx : 0
  Runts Rx     : 0                 Late Colln Tx : 0
  Giants Rx    : 0                 Excessive Colln : 0
  Total Rx Errors : 0              Deferred Tx   : 0
Others (Since boot or last clear) :
  Discard Rx    : 0                 Out Queue Len : 0
  Unknown Protos : 0
Rates (5 minute weighted average) :
  Total Rx (bps) : 3,028,168        Total Tx (bps) : 1,918,384
  Unicast Rx (Pkts/sec) : 5          Unicast Tx (Pkts/sec) : 0
  B/Mcast Rx (Pkts/sec) : 71         B/Mcast Tx (Pkts/sec) : 0
  Utilization Rx : 00.30 %           Utilization Tx : 00.19 %
```

For a given port, if a friendly port name does not exist in the running-config file, the Name line in the above command output appears as:

Name :

Searching the configuration for ports with friendly port names (CLI)

This option tells you which friendly port names have been saved to the startup-config file. (`show config` does not include ports that have only default settings in the startup-config file.)

Syntax:

```
show config
```

Includes friendly port names in a listing of all interfaces (ports) configured with non-default settings. Excludes ports that have neither a friendly port name nor any other non-default configuration settings.

See **Listing of the startup-config file with a friendly port name configured (and saved)** on page 75 to configure port A1 with a friendly port name. Notice that the command sequence saves the friendly port name for port A1 in the startup-config file. The name entered for port A2 is not saved because it was executed after `write memory`.

Listing of the startup-config file with a friendly port name configured (and saved)

```
switch(config)# int A1 name Print_Server@10.25.101.43
switch(config)# write mem
switch(config)# int A2 name Herbert's_PC

switch(config)# show config

Startup configuration:
; J9091A Configuration Editor; Created on release xx.15.05.xxxx
hostname "HPSwitch"
interface AQ
  name "Print_Server@10.25.101.43"
exit

snmp-server community "public" Unrestricted
.
.
.
```


Uni-directional link detection (UDLD)

Uni-directional link detection (UDLD) monitors a link between two switches and blocks the ports on both ends of the link if the link fails at any point between the two devices. This feature is particularly useful for detecting failures in fiber links and trunks. **Figure 9: UDLD Example:** on page 76 shows an Example:.

Figure 9: UDLD Example:

Scenario 1 (No UDLD): Without UDLD, the switch ports remain enabled despite the link failure. Traffic continues to be load-balanced to the ports connected to the failed link.

Scenario 2 (UDLD-enabled): When UDLD is enabled, the feature blocks the ports connected to the failed link.



In this Example:, each switch load balances traffic across two ports in a trunk group. Without the UDLD feature, a link failure on a link that is not directly attached to one of the HPE switches remains undetected. As a result, each switch continue to send traffic on the ports connected to the failed link. When UDLD is enabled on the trunk ports on each switch, the switches detect the failed link, block the ports connected to the failed link, and use the remaining ports in the trunk group to forward the traffic.

Similarly, UDLD is effective for monitoring fiber optic links that use two uni-direction fibers to transmit and receive packets. Without UDLD, if a fiber breaks in one direction, a fiber port may assume the link is still good (because the other direction is operating normally) and continue to send traffic on the connected ports. UDLD-enabled ports; however, will prevent traffic from being sent across a bad link by blocking the ports in the event that either the individual transmitter or receiver for that connection fails.

Ports enabled for UDLD exchange health-check packets once every five seconds (the link-keepalive interval). If a port does not receive a health-check packet from the port at the other end of the link within the keepalive interval, the port waits for four more intervals. If the port still does not receive a health-check packet after waiting for five intervals, the port concludes that the link has failed and blocks the UDLD-enabled port.

When a port is blocked by UDLD, the event is recorded in the switch log or via an SNMP trap (if configured); and other port blocking protocols, like spanning tree or meshing, will not use the bad link to load balance packets. The port will remain blocked until the link is unplugged, disabled, or fixed. The port can also be unblocked by disabling UDLD on the port.

Configuring UDLD

When configuring UDLD, keep the following considerations in mind:

- UDLD is configured on a per-port basis and must be enabled at both ends of the link. See the note below for a list of switches that support UDLD.
- To configure UDLD on a trunk group, you must configure the feature on each port of the group individually. Configuring UDLD on a trunk group's primary port enables the feature on that port only.
- Dynamic trunking is not supported. If you want to configure a trunk group that contains ports on which UDLD is enabled, you must remove the UDLD configuration from the ports. After you create the trunk group, you can re-add the UDLD configuration.

Configuring uni-directional link detection (UDLD) (CLI)

For detailed information about UDLD, see [Uni-directional link detection \(UDLD\)](#) on page 76.

Syntax:

```
[no] interface <port-list> link-keepalive
```

Enables UDLD on a port or range of ports.

To disable this feature, enter the `no` form of the command.

Default: UDLD disabled

Syntax:

```
link-keepalive interval <interval>
```

Determines the time interval to send UDLD control packets. The *interval* parameter specifies how often the ports send a UDLD packet. You can specify from 10 to 100, in 100-ms increments, where 10 is 1 second, 11 is 1.1 seconds, and so on.

Default: 50 (5 seconds)

Syntax:

```
link-keepalive retries <num>
```

Determines the maximum number of retries to send UDLD control packets. The *num* parameter specifies the maximum number of times the port will try the health check. You can specify a value from 3 to 10.

Default: 5

Syntax:

```
[no] interface <port-list> link-keepalive vlan <vid>
```

Assigns a VLAN ID to a UDLD-enabled port for sending tagged UDLD control packets. Under default settings, untagged UDLD packets can still be transmitted and received on tagged only ports; however, a warning message is logged.

The `no` form of the command disables UDLD on the specified ports.

Default: UDLD packets are untagged; tagged-only ports transmit and receive untagged UDLD control packets

Enabling UDLD (CLI)

UDLD is enabled on a per-port basis.

Example:

To enable UDLD on port a1, enter:

```
switch(config)#interface a1 link-keepalive
```

To enable the feature on a trunk group, enter the appropriate port range. For example:

```
switch(config)#interface a1-a4 link-keepalive
```



When at least one port is UDLD-enabled, the switch will forward out UDLD packets that arrive on non-UDLD-configured ports out of all other non-UDLD-configured ports in the same vlan. That is, UDLD control packets will “pass through” a port that is not configured for UDLD. However, UDLD packets will be dropped on any blocked ports that are not configured for UDLD.

Changing the keepalive interval (CLI)

By default, ports enabled for UDLD send a link health-check packet once every 5 seconds. You can change the interval to a value from 10 to 100 deciseconds, where 10 is 1 second, 11 is 1.1 seconds, and so on.

Example:

To change the packet interval to seven seconds, enter the following command at the global configuration level:

```
switch(config)# link-keepalive interval 70
```

Changing the keepalive retries (CLI)

By default, a port waits 5 seconds to receive a health-check reply packet from the port at the other end of the link. If the port does not receive a reply, the port tries four more times by sending up to four more health-check packets. If the port still does not receive a reply after the maximum number of retries, the port goes down.

You can change the maximum number of keepalive attempts to a value from 3 to 10.

Example:

To change the maximum number of attempts to four, enter the following command at the global configuration level:

```
switch(config)# link-keepalive retries 4
```

Configuring UDLD for tagged ports

The default implementation of UDLD sends the UDLD control packets untagged, even across tagged ports. If an untagged UDLD packet is received by a non-HPE switch, that switch may reject the packet. To avoid such an occurrence, you can configure ports to send out UDLD control packets that are tagged with a specified VLAN.

To enable ports to receive and send UDLD control packets tagged with a specific VLAN ID, enter a command such as the following at the interface configuration level:

```
switch(config)#interface link-keepalive vlan 22
```



- You must configure the same VLANs that will be used for UDLD on all devices across the network; otherwise, the UDLD link cannot be maintained.
- If a VLAN ID is not specified, UDLD control packets are sent out of the port as untagged packets.
- To re-assign a VLAN ID, re-enter the command with the new VLAN ID number. The new command overwrites the previous command setting.
- When configuring UDLD for tagged ports, you may receive a warning message if there are any inconsistencies with the VLAN configuration of the port.

See **Status and parameters for each port type** for potential problems.

Viewing UDLD information (CLI)

Syntax:

```
show link-keepalive
```

Displays all the ports that are enabled for `link-keepalive`.

Syntax:

```
show link-keepalive statistics
```

Displays detailed statistics for the UDLD-enabled ports on the switch.

Syntax:

```
clear link-keepalive statistics
```

Clears UDLD statistics. This command clears the packets sent, packets received, and transitions counters in the `show link-keepalive statistics display`.

Viewing summary information on all UDLD-enabled ports (CLI)

Enter the `show link-keepalive` command.

Example:

Figure 10: *Example: of show link-keepalive command*

```
HP Switch(config)# show link-keepalive
```

Total link-keepalive enabled ports: 4
Keepalive Retries: 3 Keepalive Interval: 1 sec

Port	Enabled	Physical Status	Keepalive Status	Adjacent Switch	UDLD VLAN
1	Yes	up	up	00d9d-f9b700	200
2	Yes	up	up	01560-7b1600	
3	Yes	down	off-line		
4	Yes	up	failure		
5	No	down	off-line		

Port 1 is UDLD-enabled, and tagged for a specific VLAN.

Port 3 is UDLD-enabled, but has no physical connection.

Port 4 is connected, but is blocked due to a link-keepalive failure

Port 5 has been disabled by the System Administrator.

Viewing detailed UDLD information for specific ports (CLI)

Enter the `show link-keepalive statistics` command.

Example:

Figure 11: Example: of `show link-keepalive statistics` command

```
HP Switch(config)# show link-keepalive statistics
```

Port:	1	Neighbor MAC Addr:	0000a1-b1c1d1
Current State:	up	Neighbor Port:	5
Ulld Packets Sent:	1000	State Transitions:	2
Ulld Packets Received:	1000	Link-vlan:	1
Port Blocking:	no		

Ports 1 and 2 are UDLD-enabled and show the number of health check packets sent and received on each port.

Port:	2	Neighbor MAC Addr:	000102-030405
Current State:	up	Neighbor Port:	6
Ulld Packets Sent:	500	State Transitions:	3
Ulld Packets Received:	450	Link-vlan:	200
Port Blocking:	no		

Port:	3	Neighbor MAC Addr:	n/a
Current State:	off line	Neighbor Port:	n/a
Ulld Packets Sent:	0	State Transitions:	0
Ulld Packets Received:	0	Link-vlan:	1
Port Blocking:	no		

Port 4 is shown as blocked due to a link-keepalive failure

Port:	4	Neighbor MAC Addr:	n/a
Current State:	failure	Neighbor Port:	n/a
Ulld Packets Sent:	128	State Transitions:	8
Ulld Packets Received:	50	Link-vlan:	1
Port Blocking:	yes		

Clearing UDLD statistics (CLI)

Enter the following command:

```
switch# clear link-keepalive statistics
```

This command clears the packets sent, packets received, and transitions counters in the `show link keepalive statistics` display (see [Figure 11: Example: of show link-keepalive statistics command](#) on page 80 for an Example:).

Introduction to PoE

PoE technology allows IP telephones, wireless LAN access points, and other appliances to receive power and transfer data over existing ethernet LAN cabling. For more information about PoE technology, see the PoE/PoE+ planning and implementation guide, which is available on the HPE Networking website at <http://www.hpe.com/networking>. Enter your Switch number.

Additionally, PoE+ provides more power-management capability, allowing the switch to have more power available for more PDs. Power can be allocated exactly and automatically according to what the PD actually requires at a given time.

PoE terminology

Power-over-ethernet (PoE) and Power-over-ethernet plus (PoE+ or POEP) operate similarly in most cases. Any differences between PoE and PoE+ operation are noted; otherwise, the term "PoE" is used to designate both PoE and PoE+ functionality.

Planning and implementing a PoE configuration

This section provides an overview of some considerations for planning a PoE application. For additional information on this topic, refer to the HPE PoE/PoE+ planning and implementation guide which is available on the HPE Networking web site at <http://www.hpe.com/networking>.

Some of the elements you may want to consider for a PoE installation include:

- Port assignments to VLANs
- Use of security features
- Power requirements

This section can help you to plan your PoE installation. If you use multiple VLANs in your network, or if you have concerns about network security, you should read the first two topics. If your PoE installation comes close to (or is likely to exceed) the system's ability to supply power to all devices that may request it, then you should also read the third topic. (If it is unlikely that your installation will even approach a full utilization of the PoE power available, then you may find it unnecessary to spend much time on calculating PoE power scenarios.)

Power requirements

To get the best PoE performance, you should provide enough PoE power to exceed the maximum amount of power that is needed by all the PDs that are being used.

By connecting an external power supply you can optionally provision more PoE wattage per port and or supply the switch with redundant 12V power to operate should an internal power supply fail. A Power Supply Shelf (external power supply) can also be connected to these switches to provide extra or redundant PoE power.

See the HPE PoE/PoE+ planning and implementation guide for detailed information about the PoE/PoE+ power requirements for your switch.

Assigning PoE ports to VLANs

If your network includes VLANs, you may want to assign various PoE-configured ports to specific VLANs. For example, if you are using PoE telephones in your network, you may want to assign ports used for telephone access to a VLAN reserved for telephone traffic.

Applying security features to PoE configurations

You can utilize security features built into the switch to control device or user access to the network through PoE ports in the same way as non-PoE ports.

MAC Address Security: Using Port Security, you can configure each switch port with a unique list of MAC addresses for devices that are authorized to access the network through that port. For more information, refer to the titled “Configuring and Monitoring Port Security” in the access security guide for your switch.

Assigning priority policies to PoE traffic

You can use the configurable QoS (Quality of Service) features in the switch to create prioritization policies for traffic moving through PoE ports. The available classifiers and their order of precedence are show in the table below.

Table 8: *Classifiers for prioritizing outbound packets*

Priority	QoS classifier
1	UDP/TCP application type (port)
2	Device priority (destination or source IP address)
3	IP type of service (ToS) field (IP packets only)
4	VLAN priority
5	Incoming source-port on the switch
6	Incoming 802.1 priority (present in tagged VLAN environments)

For more on this topic, refer to the titled “Quality of Service: Managing Bandwidth More Effectively” in the *Advanced Traffic Management Guide* for your switch.

PoE Event Log messages

Please see the event log message reference guide for information about Event Log messages. To see these manuals, go to <http://www.hpe.com/networking>. Auto search the model number for your switch, for Example: “HPE Switch 2530”, then select the device from the list and click on **Product manuals**. Click on the “User guide” link under **Manuals**.

About PoE operation

Using the commands described in this chapter, you can:

- Enable or disable PoE operation on individual ports.
- Monitor PoE status and performance per module.

- Configure a non-default power threshold for SNMP and Event Log reporting of PoE consumption on either all PoE ports on the switch or on all PoE ports in one or more PoE modules.
- Specify the port priority you want to use for provisioning PoE power in the event that the PoE resources become oversubscribed.

Power-sourcing equipment (PSE) detects the power needed by a powered device (PD) before supplying that power, a detection phase referred to as "searching." If the PSE cannot supply the required amount of power, it does not supply any power. For PoE using a Type 1 device, a PSE will not supply any power to a PD unless the PSE has at least 17 watts available. For example, if a PSE has a maximum available power of 382 watts and is already supplying 378 watts, and is then connected to a PD requiring 10 watts, the PSE will not supply power to the PD.

For PoE+ using Type 2 devices, the PSE must have at least 33 watts available.

Configuration options

In the default configuration, PoE support is enabled on the ports in a PoE module installed on the switch. The default priority for all ports is **low** and the default power notification threshold is **80%**. Using the CLI, you can:

- Disable or re-enable PoE operation on individual PoE ports
- Enable support for pre-standard devices
- Change the PoE priority level on individual PoE ports
- Change the threshold for generating a power level notice
- Manually allocate the amount of PoE power for a port by usage, value, or class
- Allocate PoE power based on the link-partner's capabilities via LLDP



The ports support standard networking links and PoE links. You can connect either a non-PoE device or a PD to a port enabled for PoE without reconfiguring the port.

PD support

To best utilize the allocated PoE power, spread your connected PoE devices as evenly as possible across modules. Depending on the amount of power delivered to a PoE module, there may or may not always be enough power available to connect and support PoE operation on all ports in the module. When a new PD connects to a PoE module and the module does not have enough power left for that port, if the new PD connects to a port "X" that has a:

- **Higher**

PoE priority than another port "Y" that is already supporting another PD, the power is removed from port "Y" and delivered to port "X." In this case the PD on port "Y" loses power and the PD on port "X" receives power.

- **Lower**

priority than all other PoE ports currently providing power to PDs, power is not supplied to port "X" until one or more PDs using higher priority ports are removed.

In the default configuration (*usage*), when a PD connects to a PoE port and begins operating, the port retains only enough PoE power to support the PD's operation. Unused power becomes available for supporting other PD connections. However, if you configure the `poe-allocate-by` option to either `value` or `class`, all of the power configured is allocated to the port.

For PoE (not PoE+), while 17 watts must be available for a PoE module on the switch to begin supplying power to a port with a PD connected, 17 watts per port is not continually required if the connected PD requires less power. For example, with 20 watts of PoE power remaining available on a module, you can connect one new PD without losing power to any connected PDs on that module. If that PD draws only 3 watts, 17 watts remain available, and you can connect at least one more PD to that module without interrupting power to any other PoE devices connected to the same module. If the next PD you connect draws 5 watts, only 12 watts remain unused. With only 12 unused watts available, if you then connect yet another PD to a higher-priority PoE port, the lowest-priority port

on the module loses PoE power and remains unpowered until the module once again has 17 or more watts available. (For information on power priority, see [Power priority operation](#) on page 84.)

For PoE+, there must be 33 watts available for the module to begin supplying power to a port with a PD connected.

Disconnecting a PD from a PoE port makes that power available to any other PoE ports with PDs waiting for power. If the PD demand for power becomes greater than the PoE power available, power is transferred from the lower-priority ports to the higher-priority ports. (Ports not currently providing power to PDs are not affected.)

Power priority operation

If a PSE can provide power for all connected PD demand, it does not use its power priority settings to allocate power. However, if the PD power demand oversubscribes the available power, the power allocation is prioritized to the ports that present a PD power demand. This causes the loss of power from one or more lower-priority ports to meet the power demand on other, higher-priority ports. This operation occurs regardless of the order in which PDs connect to the module's PoE-enabled ports.

Power allocation is prioritized according to the following methods:

- **Priority class** method Assigns a power priority of **low** (the default), **high**, or **critical** to each enabled PoE port.
- **Port-number priority** method A lower-numbered port has priority over a higher-numbered port within the same configured priority class, for example, port A1 has priority over port A5 if both are configured with **high** priority.

Configuring PoE operation

Disabling or re-enabling PoE port operation

Syntax:

```
[no] interface <port-list> power-over-ethernet
```

Re-enables PoE operation on <port-list> and restores the priority setting in effect when PoE was disabled on <port-list>.

The `no` form of the command disables PoE operation on <port-list>.

Default: All PoE ports are initially enabled for PoE operation at Low priority. If you configure a higher priority, this priority is retained until you change it.



For PoE, disabling all ports allows the 22 watts of minimum PoE power or the 38 watts for PoE+ power allocated for the module to be recovered and used elsewhere. You must disable ALL ports for this to occur.

Enabling support for pre-standard devices

The HPE switches covered in this guide also support some pre-802.3af devices. For a list of the supported devices, see the FAQ for your switch model.

Syntax:

```
[no] power-over-ethernet pre-std-detect
```

Detects and powers pre-802.3af standard devices.



The default setting for the `pre-std-detect` PoE parameter has changed.

Configuring the PoE port priority

Syntax:

```
interface <port-list> power-over-ethernet [critical | high | low]
```

Reconfigures the PoE priority level on <port-list>. For a given level, ports are prioritized by port number in ascending order. For example, if ports 1-24 have a priority level of critical, port 1 has priority over ports 2-24.

If there is not enough power available to provision all active PoE ports at a given priority level, the lowest-numbered port at that level is provisioned first. For chassis switches, the lowest-numbered port at that level starting with module A, then B, C, and so on is provisioned. PoE priorities are invoked only when all active PoE ports cannot be provisioned (supplied with PoE power)

Critical	Specifies the highest-priority PoE support for <port-list>. The active PoE ports at this level are provisioned before the PoE ports at any other level are provisioned.
High	Specifies the second priority PoE support for <port-list>. The active PoE ports at this level are provisioned before the <code>Low</code> priority PoE ports are provisioned.
Low	(Default) Specifies the third priority PoE support for <port-list>. The active PoE ports at this level are provisioned only if there is power available after provisioning any active PoE ports at the higher priority levels.

The following table shows some examples of PoE priority configuration.

Table 9: PoE priority operation on a PoE module

Port	Priority setting	Configuration command ¹ and resulting operation with PDs connected to ports C3 through C24
C3 - C17	Critical	<p>In this Example:, the following CLI command sets ports C3 to C17 to Critical:</p> <pre>switch(config)# interface c3-c17 power-over-ethernet critical</pre> <p>The critical priority class always receives power. If there is not enough power to provision PDs on all ports configured for this class, no power goes to ports configured for high and low priority. If there is enough power to provision PDs on only some of the critical-priority ports, power is allocated to these ports in ascending order, beginning with the lowest-numbered port in the class, which, in this case, is port 3.</p>
C18 - C21	high	<p>In this Example:, the following CLI command sets ports C19 to C22 to high:</p> <pre>switch(config)# interface c19-c22 power-over-ethernet high</pre> <p>The high priority class receives power only if all PDs on ports with a critical priority setting are receiving power. If there is not enough power to provision PDs on all ports with a high priority, no power goes to ports with a low priority. If there is enough power to provision PDs on only some of the high-priority ports, power is allocated to these ports in ascending order, beginning, in this Example:, with port 18, until all available power is in use.</p>
C22 - C24	low	<p>In this Example:, the CLI command sets ports C23 to C24 to low²:</p> <pre>switch(config)# interface c23-c24 power-over-ethernet low</pre> <p>This priority class receives power only if all PDs on ports with high and critical priority settings are receiving power. If there is enough power to provision PDs on only some low- priority ports, power is allocated to the ports in ascending order, beginning with the lowest-numbered port in the class (port 22, in this case), until all available power is in use.</p>
C1 - C2	N/A	<p>In this Example:, the CLI command disables PoE power on ports C1 to C2:</p> <pre>switch(config)# no interface c1-c2 power-over-ethernet</pre> <p>There is no priority setting for the ports in this Example:.</p>

¹

Configuring PoE operation

² In the default PoE configuration, the ports are already set to **low** priority. In this case, the command is not necessary.

Controlling PoE allocation

Syntax:

[no] int <port-list> poe-allocate-by [usage | class | value]

Allows you to manually allocate the amount of PoE power for a port by either its class or a defined value.

The default option for PoE allocation is `usage`, which is what a PD attached to the port is allocated. You can override this value by specifying the amount of power allocated to a port by using the `class` or `value` options.

<code>usage</code>	(Default) The automatic allocation by a PD.
<code>class</code>	Uses the power ramp-up signature of the PD to identify which power class the device will be in. Classes and their ranges are shown in the following table.
<code>value</code>	A user-defined level of PoE power allocated for that port.



The allowable PD requirements are lower than those specified for PSEs to allow for power losses along the Cat-5 cable.

Table 10: *Power classes and their values*

Power class	Value
0	Depends on cable type and PoE architecture. Maximum power level output of 15.4 watts at the PSE. This is the default class; if there is not enough information about the load for a specific classification, the PSE classifies the load as class 0 (zero).
1	Requires at least 4 watts at the PSE.
2	Requires at least 7 watts at the PSE.
3	15.4 watts
4	For PoE+Maximum power level output of 30 watts at the PSE.

Example:

To allocate by class for ports 6 to 8:

```
switch(config)# int 6-8 PoE-allocate-by class
```

Manually configuring PoE power levels

You can specify a power level (in watts) allocated for a port by using the `value` option. This is the maximum amount of power that will be delivered.

To configure a port by value:

Procedure

1. Set the PoE allocation by entering the `poe-allocate-by value` command:

```
switch(config) # int A6 poe-allocate-by value
```

2. or in interface context:

```
switch(eth-A6) # poe-allocate-by value
```

3. Select a value:

```
switch(config) # int A6 poe-value 15
```

4. or in interface context:

```
switch(eth-A6) # poe-value 15
```


To view the settings, enter the `show power-over-ethernet` command, shown in **Figure 12: PoE allocation by value and the maximum power delivered** on page 89.

Figure 12: *PoE allocation by value and the maximum power delivered*

```
switch(config)# show power-over-ethernet A6

Status and Counters - Port Power Status for port A7

Power Enable    : Yes
                LLDP Detect   : enabled
Priority        : low         Configured Type :
AllocateBy     : value       Value          : 15 W 1
Detection Status : Delivering Power Class    : 2

Over Current Cnt : 0         MPS Absent Cnt : 0
Power Denied Cnt : 0         Short Cnt       : 0

Voltage        : 55.1 V      Current         : 154 mA
Power          : 8.4 W
```

- Maximum power delivered.

If you set the PoE maximum value to less than what the PD requires, a fault occurs, as shown in **Figure 13: PoE power value set too low for the PD** on page 89.

Figure 13: *PoE power value set too low for the PD*

```
switch(config)# int A7 poe-value 4

switch(config)# show power-over-ethernet A7

Status and Counters - Port Power Status for port A7

Power Enable    : Yes
                LLDP Detect   : enabled
Priority        : low         Configured Type :
AllocateBy     : value       Value          : 4 W
Detection Status : fault 2    Power Class    : 2

Over Current Cnt : 1         MPS Absent Cnt : 0
Power Denied Cnt : 2         Short Cnt       : 0

Voltage        : 55.1 V      Current         : 154 mA
Power          : 8.4 W
```

- 'Fault' appears when the PoE power value is set too low.

Changing the threshold for generating a power notice

By default, PoE support is enabled on the switch's 10/100Base-TX ports, with the power priority set to **Low** and the power threshold set to **80 (%)**. The following commands allow you to adjust these settings.

¹ Maximum power delivered.

² 'Fault' appears when the PoE power value is set too low.

Syntax:

```
power threshold <1-99>
```

The power threshold is a configurable percentage of the total PoE power available on the switch. When PoE consumption exceeds the threshold, the switch automatically generates an SNMP trap and also sends a message to the Event Log. For example, if the power threshold is set to 80% (the default), and an increasing PoE power demand crosses this threshold, the switch sends an SNMP trap and generates this Event Log message:

```
PoE usage has exceeded threshold of 80 %.
```

If the switch is configured for debug logging, it also sends the same message to the configured debug destination(s).

The switch automatically invokes the power threshold at the global configuration level with a default setting of 80%. You can configure the power threshold to a value in the range of 1% to 99%.

If an increasing PoE power load (1) exceeds the configured power threshold (which triggers the log message and SNMP trap), and then (2) later begins decreasing and drops below the threshold again, the switch generates another SNMP trap, plus a message to the Event Log and any configured Debug destinations. To continue the above Example:

```
PoE usage is below configured threshold of 80 %.
```

Cycling power on a port

Simply disabling a PoE port does not affect power delivery through that port. To cycle the power on a PD receiving power from a PoE port on the switch, disable, then re-enable the power to that port.

Syntax:

```
[no] interface [e] <port-list> power
```

Re-enables PoE operation on <port-list> and restores the priority setting in effect when PoE was disabled on <port-list>. The [no] form of the command disables PoE operation on <port-list>. (Default: All 10/100Base-TX ports on the switch enabled for PoE operation at **Low** priority.)

For example, to cycle the power on a PoE device connected to port 1 on a switch covered in this guide:

```
switch(config)# no interface 1 power  
switch(config)# interface 1 power
```

PoE/PoE+ allocation using LLDP information

LLDP with PoE

When using PoE, enabling `poe-lldp-detect` allows automatic power configuration if the link partner supports PoE. When LLDP is enabled, the information about the power usage of the PD is available, and the switch can then comply with or ignore this information. You can configure PoE on each port according to the PD (IP phone, wireless device, and so on) specified in the LLDP field. The default configuration is for PoE information to be ignored if detected through LLDP.



Detecting PoE information via LLDP affects only power delivery; it does not affect normal Ethernet connectivity.

Enabling or disabling ports for allocating power using LLDP

Syntax:

```
int <port-list> poe-lldp-detect [enabled | disabled]
```

Enables or disables ports for allocating PoE power based on the link-partner's capabilities via LLDP.

Default: Disabled

Example:

You can enter this command to enable LLDP detection:

```
switch(config) # int A7 poe-lldp-detect enabled
```

or in interface context:

```
switch(eth-A7) # poe-lldp-detect enabled
```

For more information on PoE/PoE+ and LLDP, see [PoE/PoE+ allocation using LLDP information](#) on page 90.

Enabling PoE detection via LLDP TLV advertisement

Use this command and insert the desired port or ports:

```
switch(config) # lldp config <port-number> medTlvenable poe
```

For more information on LLDP, see [PoE/PoE+ allocation using LLDP information](#) on page 90.

LLDP with PoE+

Overview

The DLC for PoE provides more exact control over the power requirement between a PSE and PD. The DLC works in conjunction with the PLC and is mandatory for any Type-2 PD that requires more than 12.95 watts of input power.



DLC is defined as part of the IEEE 802.3at standard.

You can implement the power negotiation between a PSE and a PD at the physical layer or at the data link layer. After the link is powered at the physical layer, the PSE can use LLDP to query the PD repeatedly to discover the power needs of the PD. Communication over the data link layer allows finer control of power allotment, which makes it possible for the PSE to supply dynamically the power levels needed by the PD. Using LLDP is optional for the PSE but mandatory for a Type 2 PD that requires more than 12.95 watts of power.

If the power needed by the PD is not available, that port is shut off.

PoE allocation

There are two ways LLDP can negotiate power with a PD:

- **Using LLDP MED TLVs**

Disabled by default. Can be enabled using the `int <port-list> PoE-lldp-detect [enable|disable]` command, as shown below. LLDP MED TLVs sent by the PD are used to negotiate power only if

the LLDP PoE+ TLV is disabled or inactive; if the LLDP PoE+ TLV is sent as well (not likely), the LLDP MED TLV is ignored.

- **Using LLDP PoE+ TLVs**

Enabled by default. The LLDP PoE+ TLV is always advertised unless it has been disabled (enable it by using the `lldp config <port-list> dot3TlvEnable poe_config` command.) For the Command syntax, see [Initiating advertisement of PoE+ TLVs](#) on page 93. It always takes precedence over the LLDP MED TLV.

Enabling `PoE-lldp-detect` allows the data link layer to be used for power negotiation. When a PD requests power on a PoE port, LLDP interacts with PoE to see if there is enough power to fulfill the request. Power is set at the level requested. If the PD goes into power-saving mode, the power supplied is reduced; if the need for power increases, the amount supplied is increased. PoE and LLDP interact to meet the current power demands.

Syntax:

```
int <port-list> poe-lldp-detect [enabled | disabled]
```

Allows the data link layer to be used for power negotiation between a PD on a PoE port and LLDP.

Default: Disabled

Example:

You can enter this command to enable LLDP detection:

```
switch(config) # int 7 PoE-lldp-detect enabled
```

or in interface context:

```
switch(eth-7) # PoE-lldp-detect enabled
```



Detecting PoE information via LLDP affects only power delivery; it does not affect normal Ethernet connectivity.

You can view the settings by entering the `show power-over-ethernet brief` command, as shown in [Port with LLDP configuration information obtained from the device](#) on page 92.

Port with LLDP configuration information obtained from the device

```
switch(config)# show power-over-ethernet brief
```

Status and Counters - Port Power Status

POE Port	Power Enable	Power Priority	Alloc By	Alloc Power	Actual Power	Configured Type	Detection Status	Power Class
A1	Yes	low	usage	17 W	0.0 W	Phone1	Delivering	1
A2	Yes	low	usage	17 W	0.0 W		Searching	0
A3	Yes	low	usage	17 W	0.0 W		Searching	0
A4	Yes	low	usage	17 W	0.0 W		Searching	0
A5	Yes	low	usage	17 W	0.0 W		Searching	0
A6	Yes	low	usage	17 W	8.4 W		Delivering	0

Initiating advertisement of PoE+ TLVs

Syntax:

```
lldp config <port-list> dot3TlvEnable poe_config
```

Enables advertisement of data link layer power using PoE+ TLVs. The TLV is processed only after the physical layer and the data link layer are enabled. The TLV informs the PSE about the actual power required by the device.

Default: Enabled



If LLDP is disabled at runtime, and a PD is using PoE+ power that has been negotiated through LLDP, there is a temporary power drop; the port begins using PoE+ power through the PLC. This event is recorded in the Event Log. An Example: message would look like the following:

```
W 08/04/13 13:35:50 02768 ports: Port A1 PoE power dropped.  
Exceeded physical classification for a PoE Type1 device (LLDP process  
disabled)
```

When LLDP is enabled again, it causes a temporary power drop. This event is also recorded in the Event Log. An Example: message looks like the following:

```
W 08/04/13 13:36:31 02771 ports: Port A1 PoE power dropped.  
Exceeded physical classification due to change in classification type  
(LLDP process enabled)
```

Viewing PoE when using LLDP information

Syntax:

```
show lldp config <port-list>
```

Displays the LLDP port configuration information, including the TLVs advertised.

LLDP port configuration information with PoE

```
switch(config)# show lldp config 4  
  
LLCP Port Configuration Detail  
  
Port : 4  
AdminStatus [Tx_Rx] : Tx_Rx  
NotificationsEnabled [False] : False  
Med Topology Trap Enabled [False] : False  
  
TLVS Advertised:  
* port_descr  
* system_name  
* system_descr  
* system_cap  
  
* capabilities  
* network_policy  
* location_id  
* poe  
  
* macphy_config  
* poeplus_config
```

IpAddress Advertised:

Local power information on page 94 shows an Example: of the local device power information using the `show lldp info local-device <port-list>` command.

Local power information

```
switch(config)# show lldp info local-device A1
```

LLCP Local Port Information Detail

```
Port      : A1
PortType  : local
PortId    : 1
PortDesc  : A1
Pvid      : 1
```

Poe Plus Information Detail

```
Poe Device Type      : Type2 PSE
Power Source         : Primary
Power Priority        : low
PD Requested Power Value : 20 Watts
PSE Actual Power Value  : 20 Watts
```

Remote power information on page 94 shows the remote device power information using the `show lldp info remote-device <port-list>` command.

Remote power information

```
switch(config)# show lldp info remote-device A3
```

LLCP Remote Device Information Detail

```
Local Port   : A3
ChassisType  : mac-address
ChassisId    : 00 16 35 ff 2d 40
PortType     : local
PortId       : 23
SysName      : HPSwitch
System Descr : HP Switch, revision YA.14.xx
PortDescr    : 23
Pvid         : 55
```

```
System Capabilities Supported : bridge, router
System Capabilities Enabled   : bridge
```

Remote Management Address

```
Type      : ipv4
Address   : 10.0.102.198
```

Poe Plus Information Detail

```
Poe Device Type      : Type2 PD
Power Source         : Only PSE
Power Priority        : low
PD Requested Power Value : 20 Watts
PSE Actual Power Value  : 20 Watts
```

Operation note

The advertisement of power with TLVs for LLDP PoE+ is enabled by default. If LLDP is disabled at runtime and a PD is using PoE+ power that has been negotiated through LLDP, there will be a temporary power drop. The port will begin using PoE+ power through the PLC. This event is recorded in the event log. An Example: message would look like the following:

```
W 08/04/13 13:35:50 02768 ports: Port A1 PoE power dropped.  
Exceeded physical classification for a PoE Type1 device  
(LLDP process disabled)
```

When LLDP is enabled again, it causes a temporary power drop. This event is also recorded in the event log. An Example: message looks like the following:

```
W 08/04/13 13:36:31 02771 ports: Port A1 PoE power dropped.  
Exceeded physical classification due to change in  
classification type (LLDP process enabled)
```

Viewing the global PoE power status of the switch

Syntax:

```
show power-over-ethernet
```

Displays the switch's global PoE power status, including:

- **Pre-standard Detect**
Shows whether PoE for pre-802.3af-standard powered devices is enabled on the switch. (Default: **Off**; shows **On** when PoE for pre-802.3af-standard powered devices has been enabled.)
- **Operational Status**
Indicates whether PoE power is available on the switch. (Default: **On**; shows **Off** if PoE power is not available. Shows **Faulty** if internal or external PoE power is oversubscribed or faulty.)
- **Usage Threshold (%)**
Lists the configured percentage of available PoE power provisioning the switch must exceed to generate a usage notice in the form of an Event Log message and an SNMP trap. If this event is followed by a drop in power provisioning below the threshold, the switch generates another SNMP trap and Event Log message. Event Log messages are also sent to any optionally configured debug destinations. (Default: 80%)
- **Total Available Power**
Lists the maximum PoE wattage available to provision active PoE ports on the switch. This is the amount of usable power for PDs.
- **Total Failover Power**
Lists the amount of PoE power available in the event of a single power supply failure. This is the amount of power the switch can maintain without dropping any PDs.
- **Total Redundancy Power**
Indicates the amount of PoE power held in reserve for redundancy in case of a power supply failure.
- **Total Remaining Power**
The amount of PoE power still available.

<code>brief</code>	Displays PoE information for each port. See Viewing PoE status on all ports on page 96.
<code><port-list></code>	Displays PoE information for the ports in port-list. See Viewing the PoE status on specific ports on page 98.

The `show power-over-ethernet` displays data similar to that shown in [Output for the show power-over-ethernet command](#) on page 96.

Output for the show power-over-ethernet command

```
switch(config)# show power-over-ethernet

Status and Counters - System Power Status

Pre-standard Detect   : On
System Power Status   : No redundancy
PoE Power Status      : No redundancy

Chassis power-over-ethernet

Total Available Power : 600 W
Total Failover Power  : 300 W
Total Redundancy Power : 0 W
Total Used Power      : 9 W +/- 6W
Total Remaining Power : 591 W

Internal Power
 1 300W/POE /Connected.
 2 300W/POE /Connected.
 3 Not Connected.
 4 Not Connected.
External Power
EPS1 /Not Connected.
EPS2 /Not Connected.
```

Viewing PoE status on all ports

Syntax:

```
show power-over-ethernet brief
```

Displays the port power status:

PoE Port	Lists all PoE-capable ports on the switch.
Power Enable	Shows Yes for ports enabled to support PoE (the default) and No for ports on which PoE is disabled.
Power Priority	Lists the power priority (Low , High , and Critical) configured on ports enabled for PoE. (For more information on this topic, see Configuring PoE operation on page 84.)
Alloc by	Displays how PoE is allocated (usage , class , value).
Alloc Power	The maximum amount of PoE power allocated for that port (expressed in watts).Default: 17 watts for PoE; 33 watts for PoE+.
Actual Power	The power actually being used on that port.
Configured Type	If configured, shows the user-specified identifier for the port. If not configured, this field is empty.

Table Continued

Detection Status	<ul style="list-style-type: none"> • Searching: The port is trying to detect a PD connection. • Delivering: The port is delivering power to a PD. • Disabled: On the indicated port, either PoE support is disabled or PoE power is enabled but the PoE module does not have enough power available to supply the port's power needs. • Fault: The switch detects a problem with the connected PD. • Other Fault: The switch has detected an internal fault that prevents it from supplying power on that port.
Power Class	<p>Shows the 802.3af power class of the PD detected on the indicated port. Classes include:</p> <p>0: 0.44 to 12.95 watts can be drawn by the PD. Default class.</p> <p>1: 0.44 to 3.84 watts</p> <p>2: 3.84 to 6.49 watts</p> <p>3: 6.49 to 12.95 watts</p> <p>4: For PoE+; up to 25.5 watts can be drawn by the PD</p>

The `show power-over-ethernet brief` displays this output:

Output for the `show power-over-ethernet brief` command

```
switch(config)# show power-over-ethernet brief
```

```
Status and Counters - System Power Status
```

```
System Power Status   : No redundancy
PoE Power Status      : No redundancy
```

```
Available: 600 W Used: 9 W Remaining: 591 W
```

```
Module A Power
```

```
Available: 408 W Used: 9 W Remaining: 399 W
```

POE Port	Power Enable	Power Priority	Alloc By	Alloc Power	Actual Power	Configured Type	Detection Status	Power Class
A1	Yes	low	usage	17 W	0.0 W		Searching	0
A2	Yes	low	usage	17 W	0.0 W		Searching	0
A3	Yes	low	usage	17 W	0.0 W		Searching	0
A4	Yes	low	usage	17 W	0.0 W		Searching	0
A5	Yes	low	usage	17 W	0.0 W		Searching	0
A6	Yes	low	usage	17 W	8.4 W		Delivering	2
A7	Yes	low	usage	17 W	0.0 W		Searching	0
A8	Yes	low	usage	17 W	0.0 W		Searching	0
A9	Yes	low	usage	17 W	0.0 W		Searching	0

You can also show the PoE information by **slot**:

Showing the PoE information by slot

```
switch(config)# show power-over-ethernet slot A
```

```
Status and Counters - System Power Status for slot A
```

Maximum Power	: 408 W	Operational Status	: On
Power In Use	: 9 W +/- 6 W	Usage Threshold (%)	: 80

Viewing the PoE status on specific ports

Syntax:

```
show power-over-ethernet <port-list>
```

Displays the following PoE status and statistics (since the last reboot) for each port in <port-list>:

Power Enable	Shows Yes for ports enabled to support PoE (the default) and No for ports on which PoE is disabled. For ports on which power is disabled, this is the only field displayed by <code>show power-over-ethernet port-list</code> .
Priority	Lists the power priority (Low , High , and Critical) configured on ports enabled for PoE. (For more on this topic, see Configuring PoE operation on page 84.)
Allocate by	How PoE is allocated (usage , class , value).
Detection Status	<ul style="list-style-type: none"> • Searching: The port is trying to detect a PD connection. • Delivering: The port is delivering power to a PD. • Disabled: On the indicated port, either PoE support is disabled or PoE power is enabled but the PoE module does not have enough power available to supply the port's power needs. • Fault: The switch detects a problem with the connected PD. • Other Fault: The switch has detected an internal fault that prevents it from supplying power on that port.
Over Current Cnt	Shows the number of times a connected PD has attempted to draw more than 15.4 watts for PoE or 24.5 watts for PoE+. Each occurrence generates an Event Log message.
Power Denied Cnt	Shows the number of times PDs requesting power on the port have been denied because of insufficient power available. Each occurrence generates an Event Log message.
Voltage	The total voltage, in volts, being delivered to PDs.
Power	The total power, in watts, being delivered to PDs.
LLDP Detect	Port is enabled or disabled for allocating PoE power, based on the link-partner's capabilities via LLDP.

Table Continued

Configured Type	If configured, shows the user-specified identifier for the port. If not configured, the field is empty.
Value	The maximum amount of PoE power allocated for that port (expressed in watts). Default: 17 watts for PoE; 33 watts for PoE+
Power Class	Shows the power class of the PD detected on the indicated port. Classes include: 0: 0.44 to 12.95 watts 1: 0.44 to 3.84 watts 2: 3.84 to 6.49 watts 3: 6.49 to 12.95 watts 4: For PoE+; up to 25.5 watts can be drawn by the PD
MPS Absent Cnt	Shows the number of times a detected PD has no longer requested power from the port. Each occurrence generates an Event Log message. ("MPS" refers to the "maintenance power signature.")
Short Cnt	Shows the number of times the switch provided insufficient current to a connected PD.
Current	The total current, in mA, being delivered to PDs.

If you want to view the PoE status of ports A6 and A7, you would use `show power-over-ethernet A6-A7` to display the data:

Output for the `show power-over-ethernet <port-list>` command

```
switch(config)# show power-over-ethernet slot A6-A7

Status and Counters - Port Power Status for port A6

Power Enable      : Yes
Priority           : low
AllocateBy        : value
Detection Status  : Delivering
LLDP Detect       : enabled
Configured Type   :
Value             : 17 W
Power Class       : 2
Over Current Cnt  : 0
Power Denied Cnt : 0
MPS Absent Cnt   : 0
Short Cnt         : 0
Voltage           : 55.1 V
Power             : 8.4 W
Current           : 154 mA

Status and Counters - Port Power Status for port A7

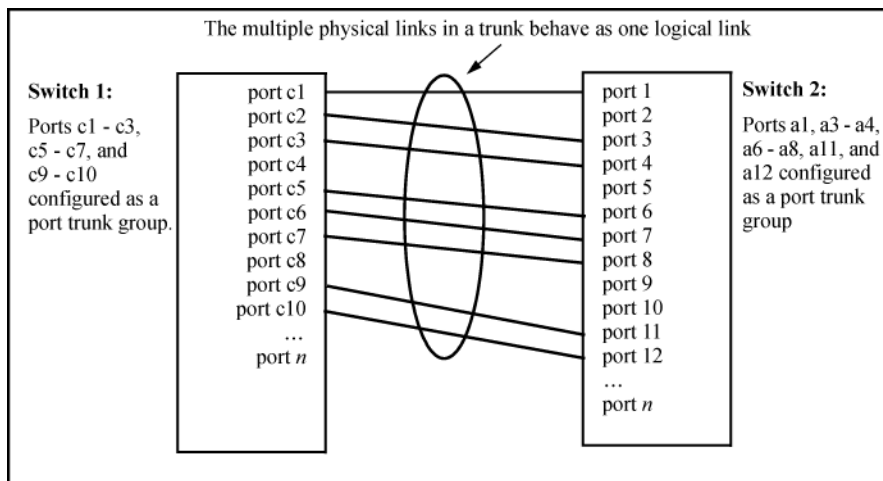
Power Enable      : Yes
Priority           : low
AllocateBy        : value
Detection Status  : Searching
LLDP Detect       : disabled
Configured Type   :
Value             : 17 W
Power Class       : 0
Over Current Cnt  : 0
MPS Absent Cnt   : 0
```

Power Denied Cnt	:0	Short Cnt	: 0
Voltage	: 0 V	Current	: 0 mA
Power	: 0 W		

Overview of port trunking

Port trunking allows you to assign up to eight physical links to one logical link (trunk) that functions as a single, higher-speed link providing dramatically increased bandwidth. This capability applies to connections between backbone devices as well as to connections in other network areas where traffic bottlenecks exist. A **trunk group** is a set of up to eight ports configured as members of the same port trunk. The ports in a trunk group do not have to be consecutive. For Example:

Figure 14: Conceptual Example: of port trunking



With full-duplex operation in a eight-port trunk group, trunking enables the following bandwidth capabilities:

Port connections and configuration

All port trunk links must be point-to-point connections between a switch and another switch, router, server, or workstation configured for port trunking. No intervening, non-trunking devices are allowed. It is important to note that ports on both ends of a port trunk group must have the same mode (speed and duplex) and flow control settings.



Link connections

The switch does not support port trunking through an intermediate, non-trunking device such as a hub, or using more than onemedia type in a port trunk group. Similarly, for proper trunk operation, all links in the same trunk group must have the samespeed, duplex, and flow control.

Port security restriction

Port security does not operate on a trunk group. If you configure port security on one or more ports that are later added to a trunk group, the switch resets the port security parameters for those ports to the factory-default configuration.



To avoid broadcast storms or loops in your network while configuring a trunk, first disable or disconnect all ports you want to add to or remove from the trunk. After you finish configuring the trunk, enable or re-connect the ports.

Port trunk features and operation

The switches covered in this guide offer these options for port trunking:

- LACP: IEEE 802.3ad—[Trunk group operation using LACP](#) on page 112
- Trunk: Non-Protocol—[Trunk group operation using the "trunk" option](#) on page 118

Up to 144 trunk groups are supported on the switches. The actual maximum depends on the number of ports available on the switch and the number of links in each trunk. (Using the link aggregation control protocol—LACP—option, you can include standby trunked ports in addition to the maximum of eight actively trunking ports.) The trunks do not have to be the same size; For example, 100 two-port trunks and 11 eight-port trunks are supported.



LACP requires full-duplex (FDx) links of the same media type (10/100Base-T, 100FX, and so on) and the same speed, and enforces speed and duplex conformance across a trunk group. For most installations, Hewlett Packard Enterprise Switch recommends that you leave the port Mode settings at `Auto` (the default). LACP also operates with `Auto-10`, `Auto-100`, and `Auto-1000` (if negotiation selects FDx), and `10FDx`, `100FDx`, and `1000FDx` settings. (The 10-gigabit ports available for some switch models allow only the `Auto` setting.)

Fault tolerance

If a link in a port trunk fails, the switch redistributes traffic originally destined for that link to the remaining links in the trunk. The trunk remains operable as long as there is at least one link in operation. If a link is restored, that link is automatically included in the traffic distribution again. The LACP option also offers a standby link capability, which enables you to keep links in reserve for service if one or more of the original active links fails. (See [Trunk group operation using LACP](#) on page 112.)

Trunk configuration methods

Dynamic LACP trunk

The switch automatically negotiates trunked links between LACP-configured ports on separate devices, and offers one dynamic trunk option: LACP. To configure the switch to initiate a dynamic LACP trunk with another device, use the `interface` command in the CLI to set the default LACP option to `active` on the ports you want to use for the trunk. For example, the following command sets ports C1 to C4 to `LACP active`:

```
switch(config) int c1-c4 lacp active
```

The preceding Example: works if the ports are not already operating in a trunk. To change the LACP option on ports already operating as a trunk, you must first remove them from the trunk. For example, if ports C1 to C4 are LACP-active and operating in a trunk with another device, you would do the following to change them to LACP-passive:

```
switch(config)# no int c1-c4 lacp
```

Removes the ports from the trunk.

```
switch(config)# int c1-c4 lacp passive
```

Configures LACP passive.

Static trunk

The switch uses the links you configure with the Port/Trunk Settings screen in the menu interface or the `trunk` command in the CLI to create a static port trunk. The switch offers two types of static trunks: LACP and Trunk.

Table 11: Trunk types used in static and dynamic trunk groups

Trunking method	LACP	Trunk
Dynamic	Yes	No
Static	Yes	Yes

The following table describes the trunking options for LACP and Trunk protocols.

Table 12: Trunk configuration protocols

Protocol	Trunking Options
LACP (802.3ad)	<p>Provides dynamic and static LACP trunking options.</p> <ul style="list-style-type: none"> • Dynamic LACP — Use the switch-negotiated dynamic LACP trunk when: <ul style="list-style-type: none"> ◦ The port on the other end of the trunk link is configured for Active or Passive LACP. ◦ You want fault-tolerance for high-availability applications. If you use an eight-link trunk, you can also configure one or more additional links to operate as standby links that will activate only if another active link goes down. • Static LACP — Use the manually configured static LACP trunk when: <ul style="list-style-type: none"> ◦ The port on the other end of the trunk link is configured for a static LACP trunk. ◦ You want to configure non-default spanning tree or IGMP parameters on an LACP trunk group. ◦ You want an LACP trunk group to operate in a VLAN other than the default VLAN and GVRP is disabled. (See VLANs and dynamic LACP on page 116.) ◦ You want to use a monitor port on the switch to monitor an LACP trunk. <p>For more information, see Trunk group operation using LACP on page 112.</p>
Trunk (non-protocol)	<p>Provides manually configured, static-only trunking to:</p> <ul style="list-style-type: none"> • Most HPE Switch and routing switches not running the 802.3ad LACP protocol. • Windows NT and HP-UX workstations and servers <p>Use the Trunk option when:</p> <ul style="list-style-type: none"> • The device to which you want to create a trunk link is using a non-802.3ad trunking protocol. • You are unsure which type of trunk to use, or the device to which you want to create a trunk link is using an unknown trunking protocol. • You want to use a monitor port on the switch to monitor traffic on a trunk. <p>See Trunk group operation using the "trunk" option on page 118.</p>

Table 13: General operating rules for port trunks

<p>Media:</p>	<p>For proper trunk operation, all ports on both ends of a trunk group must have the same media type and mode (speed and duplex). (For the switches, HPE Switch recommends leaving the port Mode setting at <code>Auto</code> or, in networks using Cat 3 cabling, <code>Auto-10</code>.)</p>
<p>Port Configuration:</p>	<p>The default port configuration is <code>Auto</code>, which enables a port to sense speed and negotiate duplex with an auto-enabled port on another device. HPE recommends that you use the <code>Auto</code> setting for all ports you plan to use for trunking. Otherwise, you must manually ensure that the mode setting for each port in a trunk is compatible with the other ports in the trunk. See: <u>Recommended port mode setting for LACP example</u></p> <p>All of the following operate on a per-port basis, regardless of trunk membership:</p> <ul style="list-style-type: none"> • Enable/Disable • Flow control (Flow Ctrl) <p>LACP is a full-duplex protocol. See <u>Trunk group operation using LACP</u> on page 112.</p>
<p>Trunk configuration:</p>	<p>All ports in the same trunk group must be the same trunk type (LACP or trunk). All LACP ports in the same trunk group must be either all static LACP or all dynamic LACP. A trunk appears as a single port labeled <code>Dyn1</code> (for an LACP dynamic trunk) or <code>Trk1</code> (for a static trunk of type LACP, Trunk) on various menu and CLI screens. For a listing of which screens show which trunk types, see <u>How the switch lists trunk data</u> on page 118. For spanning-tree or VLAN operation, configuration for all ports in a trunk is done at the trunk level. (You cannot separately configure individual ports within a trunk for spanning-tree or VLAN operation.)</p>
<p>Traffic distribution:</p>	<p>All of the switch trunk protocols use the SA/DA (source address/destination address) method of distributing traffic across the trunked links. See <u>Outbound traffic distribution across trunked links</u> on page 118.</p>

Table Continued

Spanning Tree:	<p>802.1D (STP) and 802.1w (RSTP) Spanning Tree operate as a global setting on the switch (with one instance of Spanning Tree per switch). 802.1s (MSTP) Spanning Tree operates on a per-instance basis (with multiple instances allowed per switch). For each Spanning Tree instance, you can adjust Spanning Tree parameters on a per-port basis. A static trunk of any type appears in the Spanning Tree configuration display, and you can configure Spanning Tree parameters for a static trunk in the same way that you would configure Spanning Tree parameters on a non-trunked port. (Note that the switch lists the trunk by name—such as Trk1—and does not list the individual ports in the trunk.) For example, if ports C1 and C2 are configured as a static trunk named Trk1, they are listed in the Spanning Tree display as Trk1 and do not appear as individual ports in the Spanning Tree displays. See <u>A port trunk in a Spanning Tree listing</u> on page 106. When Spanning Tree forwards on a trunk, all ports in the trunk will be forwarding. Conversely, when Spanning Tree blocks a trunk, all ports in the trunk are blocked.</p> <p>A dynamic LACP trunk operates only with the default Spanning Tree settings. Also, this type of trunk appears in the CLI <code>show spanning-tree</code> display, but not in the Spanning Tree Operation display of the Menu interface.</p> <p>If you remove a port from a static trunk, the port retains the same Spanning Tree settings that were configured for the trunk. In the below Example:, ports C1 and C2 are members of TRK1 and do not appear as individual ports in the port configuration part of the listing. See: <u>A port trunk in a Spanning Tree listing example</u></p>
IP multicast protocol (IGMP):	<p>A static trunk of any type appears in the IGMP configuration display, and you can configure IGMP for a static trunk in the same way that you would configure IGMP on a non-trunked port. (Note that the switch lists the trunk by name—such as Trk1—and does not list the individual ports in the trunk.) Also, creating a new trunk automatically places the trunk in IGMP Auto status if IGMP is enabled for the default VLAN. A dynamic LACP trunk operates only with the default IGMP settings and does not appear in the IGMP configuration display or <code>show ip igmp</code> listing.</p>
VLANs:	<p>Creating a new trunk automatically places the trunk in the DEFAULT_VLAN, regardless of whether the ports in the trunk were in another VLAN. Similarly, removing a port from a trunk group automatically places the port in the default VLAN. You can configure a static trunk in the same way that you configure a port for membership in any VLAN.</p> <p>For a dynamic LACP trunk to operate in a VLAN other than the default VLAN (DEFAULT_VLAN), GVRP must be enabled. See <u>Trunk group operation using LACP</u> on page 112.</p>
Port security:	<p>Trunk groups (and their individual ports) cannot be configured for port security, and the switch excludes trunked ports from the <code>show port-security</code> listing. If you configure non-default port security settings for a port, then subsequently try to place the port in a trunk, you see the following message and the command is not executed: <code>< port-list > Command cannot operate over a logical port.</code></p>
Monitor port:	<p>A trunk cannot be a monitor port. A monitor port can monitor a static trunk but cannot monitor a dynamic LACP trunk.</p>

Recommended port mode setting for LACP

```
switch(config)# show interfaces config
```

Port Settings

Port	Type	Enabled	Mode	Flow Ctrl	MDI
1	10/100TX	Yes	Auto	Enable	Auto
2	10/100TX	Yes	Auto	Enable	MDI

A port trunk in a Spanning Tree listing

Port	Type	Cost	Priority	State	Designated Bridge
C3	100/1000T	5	12B	Forwarding	0020c1-b27ac0
C4	100/1000T	5	12B	Forwarding	0060b0-889e00
C5	100/1000T	5	12B	Disabled	
C6	100/1000T	5	12B	Disabled	
Trk1		1	64	Forwarding	0001e7-a0ec00

Viewing and configuring a static trunk group (Menu)

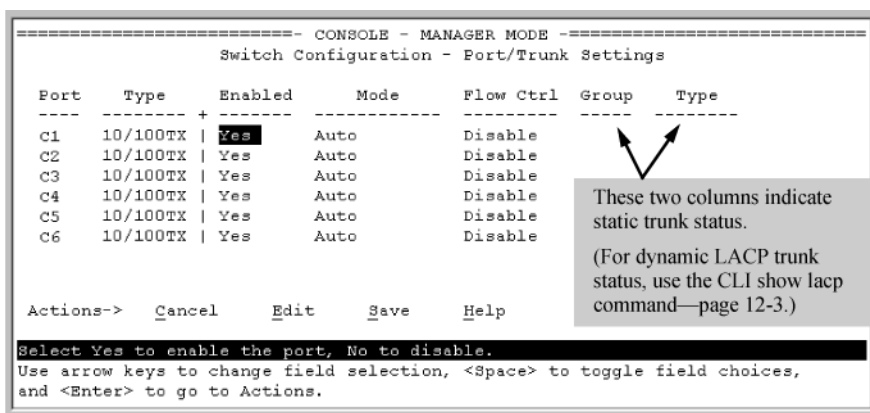
- ⓘ Configure port trunking **before** you connect the trunked links to another switch, routing switch, or server. Otherwise, a broadcast storm could occur. (If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until the trunk is configured. See "Enabling or Disabling Ports and Configuring Port Mode".)

This procedure uses the Port/Trunk Settings screen to configure a static port trunk group on the switch.

Procedure

1. Follow the procedures in the preceding IMPORTANT note.
2. From the Main Menu, select:
 2. Switch Configuration...
 2. Port/Trunk Settings
3. Press [E] (for Edit) and then use the arrow keys to access the port trunk parameters.

Figure 15: Example: of the menu screen for configuring a port trunk group



4. In the Group column, move the cursor to the port you want to configure.
5. Use the Space bar to choose a trunk group assignment (Trk1, Trk2, and so on) for the selected port.
 - a. For proper trunk operation, all ports in a trunk must have the same media type and mode (such as 10/100TX set to 100FDx, or 100FX set to 100FDx). The flow control settings must also be the same for all ports in a given trunk. To verify these settings, see "Viewing Port Status and Configuring Port Parameters".
 - b. You can configure the trunk group with up to eight ports per trunk. If multiple VLANs are configured, all ports within a trunk will be assigned to the same VLAN or set of VLANs. (With the 802.1Q VLAN capability built into the switch, more than one VLAN can be assigned to a trunk. See the "Static Virtual LANs (VLANs)" in the advanced traffic management guide for your switch.)

(To return a port to a non-trunk status, keep pressing the Space bar until a blank appears in the highlighted Group value for that port.)

Figure 16: Example: of the Configuration for a Two-Port Trunk Group

```

=====  CONSOLE - MANAGER MODE  =====
                Switch Configuration - Port/Trunk Settings
-----
Port   Type      Enabled  Mode    Flow Ctrl  Group  Type
-----
C1    10/100TX | Yes     Auto    Disable   -
C2    10/100TX | Yes     Auto    Disable   -
C3    10/100TX | Yes     Auto    Disable   -
C4    10/100TX | Yes     Auto    Disable   -
C5    10/100TX | Yes     Auto    Disable   Trk1  Trunk
C6    10/100TX | Yes     Auto    Disable   Trk1  Trunk

Actions->  _Cancel    _Edit     _Save     _Help

Select whether the port is part of a trunk or Mesh.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
  
```

6. Move the cursor to the Type column for the selected port and use the Space bar to select the trunk type:
 - a. LACP
 - b. Trunk (the default type if you do not specify a type)

All ports in the same trunk group on the same switch must have the same Type (LACP or Trunk).

7. When you are finished assigning ports to the trunk group, press **[Enter]**, then **[S]** (for Save) and return to the Main Menu. (It is not necessary to reboot the switch.)

During the Save process, traffic on the ports configured for trunking is delayed for several seconds. If the Spanning Tree Protocol is enabled, the delay may be up to 30 seconds.

8. Connect the trunked ports on the switch to the corresponding ports on the opposite device. If you previously disabled any of the trunked ports on the switch, enable them now. (See "Viewing Port Status and Configuring Port Parameters")

Check the Event Log ("Using the Event Log for Troubleshooting Switch Problems") to verify that the trunked ports are operating properly.

Viewing and configuring port trunk groups (CLI)

You can list the trunk type and group for all ports on the switch or for selected ports. You can also list LACP-only status information for LACP-configured ports.

Viewing static trunk type and group for all ports or for selected ports

Syntax:

```
show trunks [< port-list >]
```

Omitting the `<port-list>` parameter results in a static trunk data listing for all LAN ports in the switch.

Example:

In a switch where ports A4 and A5 belong to Trunk 1 and ports A7 and A8 belong to Trunk 2, you have the options shown in [Listing specific ports belonging to static trunks](#) on page 108 and [A show trunk listing without specifying ports](#) on page 108 for displaying port data for ports belonging to static trunks.

Using a port list specifies, for switch ports in a static trunk group, only the ports you want to view. In this case, the command specifies ports A5 through A7. However, because port A6 is not in a static trunk group, it does not appear in the resulting listing:

Listing specific ports belonging to static trunks

```
switch# show trunks e 5-7
```

Load Balancing

Port	Name	Type	Group	Type
5	Print-Server-Trunk	10/100TX	Trk1	Trunk
7		10/100TX	Trk2	Trunk

The `show trunks <port-list>` command in the above Example: includes a port list, and thus shows trunk group information only for specific ports that have membership in a static trunk. In [A show trunk listing without specifying ports](#) on page 108, the command does not include a port list, so the switch lists all ports having static trunk membership.

A show trunk listing without specifying ports

```
switch# show trunks
```

Load Balancing

Port	Name	Type	Group	Type
4	Print-Server-Trunk	10/100TX	Trk1	Trunk
5	Print-Server-Trunk	10/100TX	Trk1	Trunk
7		10/100TX	Trk2	Trunk
8		10/100TX	Trk2	Trunk

Viewing static LACP and dynamic LACP trunk data

Syntax:

```
show lacp
```

Lists data for only the LACP-configured ports.

Example:

Ports A1 and A2 have been previously configured for a static LACP trunk. (For more on the `Active` parameter, see table "[LACP port status data](#)".)

A show LACP listing

```
switch# show lacp
```

LACP

Port	LACP Enabled	Trunk Group	Port Status	Partner	LACP Status	Admin Key	Oper Key
A1	Active	Trkl	Up	Yes	Success	0	250
A2	Active	Trkl	Up	Yes	Success	0	250
A3	Active	A3	Down	No	Success	0	300
A4	Passive	A4	Down	No	Success	0	0
A5	Passive	A5	Down	No	Success	0	0
A6	Passive	A6	Down	No	Success	0	0

For a description of each of the above-listed data types, see table "[LACP port status data](#)".

Dynamic LACP Standby Links

Dynamic LACP trunking enables you to configure standby links for a trunk by including more than eight ports in a dynamic LACP trunk configuration. When eight ports (trunk links) are up, the remaining link(s) will be held in standby status. If a trunked link that is "Up" fails, it will be replaced by a standby link, which maintains your intended bandwidth for the trunk. (Refer to also the "Standby" entry under "Port Status" in "Table 4-5. LACP Port Status Data".) In the next Example:, ports A1 through A9 have been configured for the same LACP trunk. Notice that one of the links shows Standby status, while the remaining eight links are "Up".

A Dynamic LACP trunk with one standby link

```
switch# show lacp
```

Port	LACP						
	LACP Enabled	Trunk Group	Port Status	Partner	LACP Status	Admin Key	Oper Key
A1	Active	Dyn1	Up	Yes	Success	100	100
A2	Active	Dyn1	Up	Yes	Success	100	100
A3	Active	Dyn1	Up	Yes	Success	100	100
A4	Active	Dyn1	Up	Yes	Success	100	100
A5	Active	Dyn1	Up	Yes	Success	100	100
A6	Active	Dyn1	Up	Yes	Success	100	100
A7	Active	Dyn1	Up	Yes	Success	100	100
A8	Active	Dyn1	Up	Yes	Success	100	100
A9	Active	Dyn1	Standby	Yes	Success	100	100

Configuring a static trunk or static LACP trunk group

- ⓘ Configure port trunking **before** you connect the trunked links between switches. Otherwise, a broadcast storm could occur. (If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until the trunk is configured. See "Enabling or Disabling Ports and Configuring Port Mode".)

The "[Port trunk features and operation](#)" section describes the maximum number of trunk groups you can configure on the switch. An individual trunk can have up to eight links, with additional standby links if you're using LACP. You can configure trunk group types as follows:

Trunk Type	Trunk Group Membership	
	TrkX (Static)	DynX (Dynamic)
LACP	Yes	Yes
Trunk	Yes	No

The following examples show how to create different types of trunk groups.

Syntax:

```
trunk <port-list> <trk1 ... trk144> {<trunk | lacp>}
```

Configures the specified static trunk type.

Example:

This Example: uses ports C4 to C6 to create a non-protocol static trunk group with the group name `Trk2`.

```
switch(config)# trunk c4-c6 trk2 trunk
```

Removing ports from a static trunk group



Removing a port from a trunk can create a loop and cause a broadcast storm. When you remove a port from a trunk where spanning tree is not in use, HPE Switch recommends that you first disable the port or disconnect the link on that port.

Syntax:

```
no trunk <port-list>
```

Removes the specified ports from an existing trunk group.

Example:

To remove ports C4 and C5 from an existing trunk group:


```
switch(config)# no trunk c4-c5
```

Enabling a dynamic LACP trunk group

In the default port configuration, all ports on the switch are set to disabled. To enable the switch to automatically form a trunk group that is dynamic on both ends of the link, the ports on one end of a set of links must be LACP Active. The ports on the other end can be either LACP Active or LACP Passive. The `active` command enables the switch to automatically establish a (dynamic) LACP trunk group when the device on the other end of the link is configured for LACP Passive.

Example:

Figure 17: Criteria for automatically forming a dynamic LACP trunk



Syntax:

```
interface <port-list> lacp active
```

Configures *<port-list>* as LACP active. If the ports at the other end of the links on *<port-list>* are configured as LACP passive, this command enables a dynamic LACP trunk group on *<port-list>*.

Example:

This Example: uses ports C4 and C5 to enable a dynamic LACP trunk group.

```
switch(config)# interface c4-c5 lacp active
```

Removing ports from a dynamic LACP trunk group

To remove a port from dynamic LACP trunk operation, you must turn off LACP on the port. (On a port in an operating, dynamic LACP trunk, you cannot change between LACP *Active* and LACP *passive* without first removing LACP operation from the port.)



Unless spanning tree is running on your network, removing a port from a trunk can result in a loop. To help prevent a broadcast storm when you remove a port from a trunk where spanning tree is not in use, Hewlett Packard Enterprise recommends that you first disable the port or disconnect the link on that port.

Syntax:

```
no interface <port-list> lacp
```

Removes *<port-list>* from any dynamic LACP trunk and returns the ports in *<port-list>* to passive LACP.

Example:

Port C6 belongs to an operating, dynamic LACP trunk. To remove port C6 from the dynamic trunk and return it to passive LACP, do the following:

```
switch(config)# no interface c6 lacp
switch(config)# interface c6 lacp passive
```

In the above Example:, if the port on the other end of the link is configured for active LACP or static LACP, the trunked link will be re-established almost immediately.

Viewing existing port trunk groups (WebAgent)

While the WebAgent does not enable you to configure a port trunk group, it does provide a view of an existing trunk group.

To view any port trunk groups:

1. In the navigation pane, click **Interface**.
2. Click **Port Info/Config**. The trunk information for the port displays in the **Port Properties** box.

Trunk group operation using LACP

The switch can automatically configure a dynamic LACP trunk group, or you can manually configure a static LACP trunk group.



LACP requires full-duplex (FDx) links of the same media type (10/100Base-T, 100FX, and so on) and the same speed and enforces speed and duplex conformance across a trunk group. For most installations, HPE Switch recommends that you leave the port mode settings at `Auto` (the default). LACP also operates with `Auto-10`, `Auto-100`, and `Auto-1000` (if negotiation selects FDx), and `10FDx`, `100FDx`, and `1000FDx` settings.

LACP trunk status commands include:

Trunk display method	Static LACP trunk	Dynamic LACP trunk
CLI <code>show lacp</code> command	Included in listing.	Included in listing.
CLI <code>show trunk</code> command	Included in listing.	Not included.
Port/Trunk Settings screen in menu interface	Included in listing.	Not included

Thus, to display a listing of dynamic LACP trunk ports, you must use the `show lacp` command.

In most cases, trunks configured for LACP on the switches operate as described in the following table.

Table 14: LACP trunk types

LACP port trunk configuration	Operation
<p>Dynamic LACP</p>	<p>This option automatically establishes an 802.3ad-compliant trunk group, with LACP for the port Type parameter and DynX for the port Group name, where X is an automatically assigned value from 1 to 144, depending on how many dynamic and static trunks are currently on the switch. (The switch allows a maximum of 144 trunk groups in any combination of static and dynamic trunks.)</p> <p>Dynamic LACP trunks operate only in the default VLAN (unless GVRP is enabled and <code>Forbid</code> is used to prevent the trunked ports from joining the default VLAN). Thus, if an LACP dynamic port forms using ports that are not in the default VLAN, the trunk automatically moves to the default VLAN unless GVRP operation is configured to prevent this from occurring. In some cases, this can create a traffic loop in your network. For more information on this topic, see VLANs and dynamic LACP on page 116 .</p> <p>Under the following conditions, the switch automatically establishes a dynamic LACP port trunk group and assigns a port Group name:</p> <ul style="list-style-type: none"> • The ports on both ends of each link have compatible mode settings (speed and duplex). • The port on one end of each link must be configured for LACP Active and the port on the other end of the same link must be configured for either LACP Passive or LACP Active. For Example: <div data-bbox="607 974 1357 1201" data-label="Diagram"> <pre> graph LR subgraph Switch1 [Switch 1] P1[Port X: LACP Enable: Active] P2[Port Y: LACP Enable: Active] end subgraph Switch2 [Switch 2] P3[Port A: LACP Enable: Active] P4[Port B: LACP Enable: Passive] end P1 --- Active-to-Active P3 P2 --- Active-to-Passive P4 </pre> </div> <p>Either of the above link configurations allows a dynamic LACP trunk link.</p> <p>Backup Links: A maximum of eight operating links are allowed in the trunk, but, with dynamic LACP, you can configure one or more additional (backup) links that the switch automatically activates if a primary link fails. To configure a link as a standby for an existing eight-port dynamic LACP trunk, ensure that the ports in the standby link are configured as either active-to-active or active-to-passive between switches.</p> <p>Displaying dynamic LACP trunk data: To list the configuration and status for a dynamic LACP trunk, use the CLI <code>show lacp</code> command.</p> <p>The dynamic trunk is automatically created by the switch and is not listed in the static trunk listings available in the menu interface or in the CLI <code>show trunk</code> listing.</p>
<p>Static LACP</p>	<p>Provides a manually configured, static LACP trunk to accommodate these conditions:</p>

LACP port trunk configuration	Operation
	<ul style="list-style-type: none"> • The port on the other end of the trunk link is configured for a static LACP trunk. • You want to configure non-default Spanning Tree or IGMP parameters on an LACP trunk group. • You want an LACP trunk group to operate in a VLAN other than the default VLAN and GVRP is disabled. (See VLANS and dynamic LACP on page 116.) • You want to use a monitor port on the switch to monitor an LACP trunk. <p>The trunk operates if the trunk group on the opposite device is running one of the following trunking protocols:</p> <ul style="list-style-type: none"> • Active LACP • Passive LACP • Trunk <p>This option uses LACP for the port Type parameter and TrkX for the port Group parameter, where X is an automatically assigned value in a range corresponding to the maximum number of trunks the switch allows. (See Port trunk features and operation for the maximum number of trunk groups allowed on the switches.)</p> <p>Displaying static LACP trunk data : To list the configuration and status for a static LACP trunk, use the CLI <code>show lacp</code> command. To list a static LACP trunk with its assigned ports, use the CLI <code>show trunk</code> command or display the menu interface Port/Trunk Settings screen. Static LACP does not allow standby ports.</p>

Default port operation

In the default configuration, LACP is disabled for all ports. If LACP is not configured as Active on at least one end of a link, the port does not try to detect a trunk configuration and operates as a standard, untrunked port. The following table lists the elements of per-port LACP operation. To display this data for a switch, execute the following command in the CLI:

```
switch# show lacp
```

Table 15: LACP port status data

Status name	Meaning
Port Numb	Shows the physical port number for each port configured for LACP operation (C1, C2, C3 ...). Unlisted port numbers indicate that the missing ports that are assigned to a static trunk group are not configured for any trunking.
LACP Enabled	<p>Active: The port automatically sends LACP protocol packets.</p> <p>Passive: The port does not automatically send LACP protocol packets and responds only if it receives LACP protocol packets from the opposite device. A link having either two active LACP ports or one active port and one passive port can perform dynamic LACP trunking. A link having two passive LACP ports does not perform LACP trunking because both ports are waiting for an LACP protocol packet from the opposite device.</p> <p>In the default switch configuration, LACP is disabled for all ports.</p>

Table Continued

Status name	Meaning
Trunk Group	<p>TrkX: This port has been manually configured into a static LACP trunk.</p> <p>Trunk group same as port number: The port is configured for LACP, but is not a member of a port trunk.</p>
Port Status	<p>Up: The port has an active LACP link and is not blocked or in standby mode.</p> <p>Down: The port is enabled, but an LACP link is not established. This can indicate, For example, a port that is not connected to the network or a speed mismatch between a pair of linked ports.</p> <p>Disabled: The port cannot carry traffic.</p> <p>Blocked: LACP, Spanning Tree has blocked the port. (The port is not in LACP standby mode.) This may be caused by a (brief) trunk negotiation or a configuration error, such as differing port speeds on the same link or trying to connect the switch to more trunks than it can support. (See Trunk configuration protocols.)</p> <p>Some older devices are limited to four ports in a trunk. When eight LACP-enabled ports are connected to one of these older devices, four ports connect, but the other four ports are blocked.</p> <p>Standby: The port is configured for dynamic LACP trunking to another device, but the maximum number of ports for the dynamic trunk to that device has already been reached on either the switch or the other device. This port will remain in reserve, or "standby" unless LACP detects that another, active link in the trunk has become disabled, blocked, or down. In this case, LACP automatically assigns a standby port, if available, to replace the failed port.</p>
LACP Partner	<p>Yes: LACP is enabled on both ends of the link.</p> <p>No: LACP is enabled on the switch, but either LACP is not enabled or the link has not been detected on the opposite device.</p>
LACP Status	<p>Success: LACP is enabled on the port, detects and synchronizes with a device on the other end of the link, and can move traffic across the link.</p> <p>Failure: LACP is enabled on a port and detects a device on the other end of the link, but is not able to synchronize with this device, and therefore is not able to send LACP packets across the link. This can be caused, For example, by an intervening device on the link (such as a hub), a bad hardware connection, or if the LACP operation on the opposite device does not comply with the IEEE 802.3ad standard.</p>

LACP notes and restrictions

802.1X (Port-based access control) configured on a port

To maintain security, LACP is not allowed on ports configured for 802.1X authenticator operation. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port, and enables 802.1X on that port.

```
switch(config)# aaa port-access authenticator b1
LACP has been disabled on 802.1x port(s).
switch(config)#
```

The switch does not allow you to configure LACP on a port on which port access (802.1X) is enabled. For Example:

```
switch(config)# int b1 lacp passive
Error configuring port < port-number > : LACP and 802.1x cannot
```

```
be run together.  
switch(config)#
```

To restore LACP to the port, you must first remove the 802.1X configuration of the port and then re-enable LACP active or passive on the port.

Port security configured on a port

To maintain security, LACP is not allowed on ports configured for port security. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port, and enables port security on that port. For example:

```
switch(config)# port-security a17 learn-mode static address-  
limit 2 LACP has been disabled on secured port(s).  
switch(config)#
```

The switch does not allow you to configure LACP on a port on which port security is enabled. For example:

```
switch(config)# int a17 lacp passive  
Error configuring port A17: LACP and port security cannot be  
run together.  
switch(config)#
```

To restore LACP to the port, you must remove port security and re-enable LACP active or passive.

Changing trunking methods

To convert a trunk from static to dynamic, you must first eliminate the static trunk.

Static LACP trunks

When a port is configured for LACP (active or passive), but does not belong to an existing trunk group, you can add that port to a static trunk. Doing so disables dynamic LACP on that port, which means you must manually configure both ends of the trunk.

Dynamic LACP trunks

You can configure a port for LACP-active or LACP-passive, but on a dynamic LACP trunk you cannot configure the other options that you can on static trunks. If you want to manually configure a trunk, use the `trunk` command.

VLANs and dynamic LACP

A dynamic LACP trunk operates only in the default VLAN (unless you have enabled GVRP on the switch and use `Forbid` to prevent the ports from joining the default VLAN).

If you want to use LACP for a trunk on a non-default VLAN and GVRP is disabled, configure the trunk as a static trunk.

Blocked ports with older devices

Some older devices are limited to four ports in a trunk. When eight LACP-enabled ports are connected to one of these older devices, four ports connect, but the other four ports are blocked. The LACP status of the blocked ports is shown as "Failure."

If one of the other ports becomes disabled, a blocked port replaces it (Port Status becomes "Up"). When the other port becomes active again, the replacement port goes back to blocked (Port Status is "Blocked"). It can take a few seconds for the switch to discover the current status of the ports.


Blocked ports with LACP

```
switch(eth-B1-B8) # show lacp
```

LACP					
PORT NUMB	LACP ENABLED	TRUNK GROUP	PORT STATUS	LACP PARTNER	LACP STATUS
B1	Active	Dyn1	Up	Yes	Success
B2	Active	Dyn1	Up	Yes	Success
B3	Active	Dyn1	Up	Yes	Success
B4	Active	Dyn1	Up	Yes	Success
B5	Active	Dyn1	Blocked	Yes	Failure
B6	Active	Dyn1	Blocked	Yes	Failure
B7	Active	B7	Down	No	Success
B8	Active	B8	Down	No	Success

If there are ports that you do not want on the default VLAN, ensure that they cannot become dynamic LACP trunk members. Otherwise a traffic loop can unexpectedly occur. For Example:

Figure 18: A dynamic LACP trunk forming in a VLAN can cause a traffic loop



Easy control methods include either disabling LACP on the selected ports or configuring them to operate in static LACP trunks.

Spanning Tree and IGMP

If Spanning Tree, IGMP, or both are enabled in the switch, a dynamic LACP trunk operates only with the default settings for these features and does not appear in the port listings for these features.

Half-duplex, different port speeds, or both not allowed in LACP trunks

The ports on both sides of an LACP trunk must be configured for the same speed and for full-duplex (FDx). The 802.3ad LACP standard specifies a full-duplex (FDx) requirement for LACP trunking. (10-gigabit ports operate only at FDx.)

A port configured as LACP passive and not assigned to a port trunk can be configured to half-duplex (HDx). However, in any of the following cases, a port cannot be reconfigured to an HDx setting:

- If the port is a 10-gigabit port.
- If a port is set to LACP Active, you cannot configure it to HDx.
- If a port is already a member of a static or dynamic LACP trunk, you cannot configure it to HDx.
- If a port is already set to HDx, the switch does not allow you to configure it for a static or dynamic LACP trunk.

Dynamic/static LACP interoperation

A port configured for dynamic LACP can properly interoperate with a port configured for static (TrkX) LACP, but any ports configured as standby LACP links are ignored.

Trunk group operation using the "trunk" option

This method creates a trunk group that operates independently of specific trunking protocols and does not use a protocol exchange with the device on the other end of the trunk. With this choice, the switch simply uses the SA/DA method of distributing outbound traffic across the trunked ports without regard for how that traffic is handled by the device at the other end of the trunked links. Similarly, the switch handles incoming traffic from the trunked links as if it were from a trunked source.

When a trunk group is configured with the `trunk` option, the switch automatically sets the trunk to a priority of "4" for Spanning Tree operation (even if Spanning Tree is currently disabled). This appears in the running-config file as `spanning-tree Trkn priority 4`. Executing `write memory` after configuring the trunk places the same entry in the startup-config file.

Use the `trunk` option to establish a trunk group between a switch and another device, where the other device's trunking operation fails to operate properly with LACP trunking configured on the switches.

How the switch lists trunk data

Static trunk group	Appears in the menu interface and the output from the CLI <code>show trunk</code> and <code>show interfaces</code> commands.
Dynamic LACP trunk group	Appears in the output from the CLI <code>show lacp</code> command.

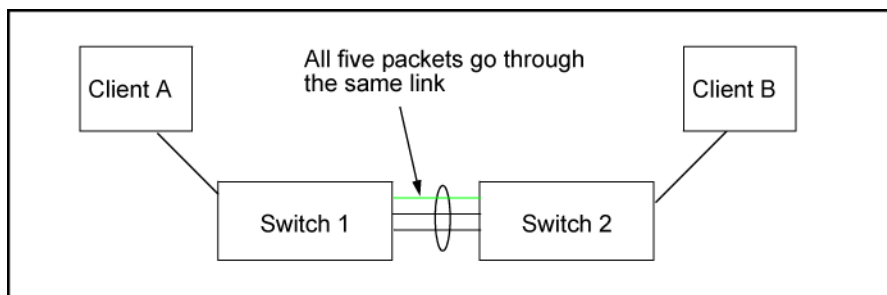
Interface option	Dynamic LACP trunk group	Static LACP trunk group	Static non-protocol
Menu interface	No	Yes	Yes
CLI <code>show trunk</code>	No	Yes	Yes
CLI <code>show interfaces</code>	No	Yes	Yes
CLI <code>show lacp</code>	Yes	Yes	No
CLI <code>show spanning-tree</code>	No	Yes	Yes
CLI <code>show igmp</code>	No	Yes	Yes
CLI <code>show config</code>	No	Yes	Yes

Outbound traffic distribution across trunked links

The two trunk group options (LACP and trunk) use SA/DA pairs for distributing outbound traffic over trunked links. That is, the switch sends traffic from the same source address to the same destination address through the same trunked link, and may also send traffic from the same source address to a different destination address through the same link or a different link, depending on the mapping of path assignments among the links in the trunk. Likewise, the switch distributes traffic for the same destination address but from different source addresses through links depending on the path assignment.

The load-balancing is done on a per-communication basis. Otherwise, traffic is transmitted across the same path as shown in the figure below. That is, if Client A attached to Switch 1 sends five packets of data to Server A attached to Switch 2, the same link is used to send all five packets. The SA/DA address pair for the traffic is the same. The packets are not evenly distributed across any other existing links between the two switches; they all take the same path.

Figure 19: Example: of single path traffic through a trunk



The actual distribution of the traffic through a trunk depends on a calculation using bits from the SA/DA. When an IP address is available, the calculation includes the last five bits of the IP source address and IP destination address; otherwise, the MAC addresses are used. The result of that process undergoes a mapping that determines which link the traffic goes through. If you have only two ports in a trunk, it is possible that all the traffic will be sent through one port even if the SA/DA pairs are different. The more ports you have in the trunk, the more likely it is that the traffic will be distributed among the links.

When a new port is added to the trunk, the switch begins sending traffic, either new traffic or existing traffic, through the new link. As links are added or deleted, the switch redistributes traffic across the trunk group. For example, in the figure below showing a three-port trunk, traffic could be assigned as shown in the following table.

Figure 20: Example: of port-trunked network




Table 16: Example: of link assignments in a trunk group (SA/DA distribution)

Source	Destination	Link
Node A	Node W	1
Node B	Node X	2
Node C	Node Y	3
Node D	Node Z	1
Node A	Node Y	2
Node B	Node W	3

Because the amount of traffic coming from or going to various nodes in a network can vary widely, it is possible for one link in a trunk group to be fully utilized while other links in the same trunk have unused bandwidth capacity, even if the assignments were evenly distributed across the links in a trunk.

VLAN-based rate-limiting

VLAN-based rate-limiting provides specific bandwidth for a specific VLAN for the inbound traffic on the VLAN. It allows the user to specify the maximum number of kilobits per second (kbps) a VLAN can receive. The specified VLAN drops all traffic that exceeds the configured rate. Unused bandwidth is not carried over from one interval to the next.

VLAN-based rate limiting and rate limiting per interface/port are supported together on the same port. If applying any of these rate-limit configurations results in blocking the inbound traffic on that port, further packets on that port are dropped. Any VLAN-based rate limiting applies to all traffic on that VLAN.

Global QoS priority policies and rate limit policies cannot be configured at the same time.

VLAN-based rate limiting is disabled by default.



NOTE

VLAN-based rate limiting does not work on broadcast traffic. The hardware design lowers the priority of broadcast traffic to avoid flooding the CPU.



NOTE

The HPE Switch supports data rate in multiples of 64 kbps. If the configured value is between 1-63, it is rounded off to 64. For all other non-zero values, it is rounded off to the nearest lower multiple of 64, and the rate-limit is applied

Setting limits for inbound traffic on a specific VLAN

```
switch(vlan-1)# rate-limit all in kbps <rate>
```

```
Usage: rate-limit all in kbps <rate>
       no rate-limit all in
```

```
all: Set limits for all traffic
in : Set limits for all inbound traffic
```

ICMP rate-limiting

In IP networks, ICMP messages are generated in response to either inquiries or requests from routing and diagnostic functions. These messages are directed to the applications originating the inquiries. In unusual situations, if the messages are generated rapidly with the intent of overloading network circuits, they can threaten network availability. This problem is visible in denial-of-service (DoS) attacks or other malicious behaviors where a worm or virus overloads the network with ICMP messages to an extent where no other traffic can get through. (ICMP messages themselves can also be misused as virus carriers). Such malicious misuses of ICMP can include a high number of ping packets that mimic a valid source IP address and an invalid destination IP address (spoofed pings), and a high number of response messages (such as Destination Unreachable error messages) generated by the network.

ICMP rate-limiting provides a method for limiting the amount of bandwidth that may be used for inbound ICMP traffic on a switch port. This feature allows users to restrict ICMP traffic to percentage levels that permit necessary ICMP functions, but throttle additional traffic that may be caused by worms or viruses (reducing their spread and effect). In addition, ICMP rate-limiting preserves inbound port bandwidth for non-ICMP traffic.



ICMP is necessary for routing, diagnostic, and error responses in an IP network. ICMP rate-limiting is primarily used for throttling worm or virus-like behavior and should normally be configured to allow one to five percent of available inbound bandwidth (at 10 Mbps or 100 Mbps speeds) or 100 to 10,000 kbps (1Gbps or 10 Gbps speeds) to be used for ICMP traffic. **This feature should not be used to remove all ICMP traffic from a network.**



ICMP rate-limiting does not throttle non-ICMP traffic. In cases where you want to throttle both ICMP traffic and all other inbound traffic on a given interface, you can separately configure both ICMP rate-limiting and all-traffic rate-limiting.

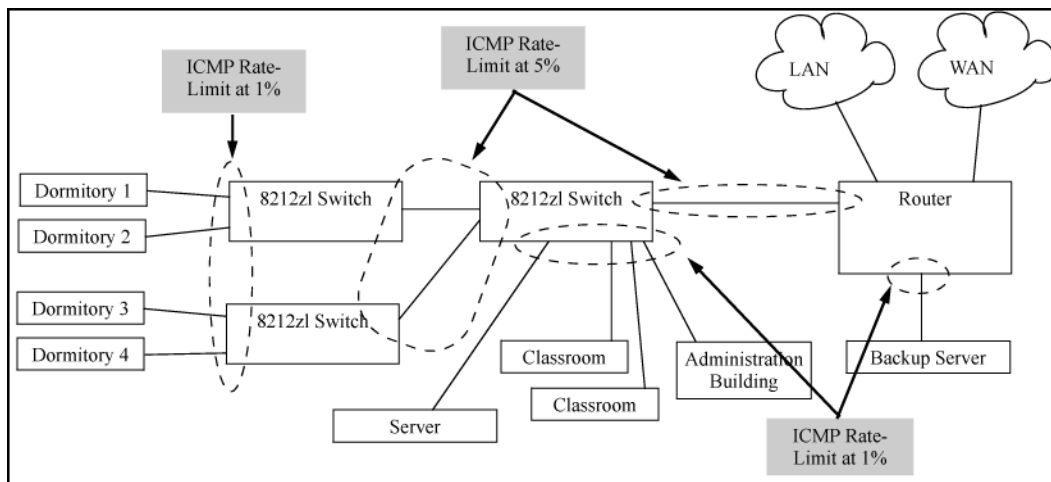
The all-traffic rate-limiting command (`rate-limit all`) and the ICMP rate-limiting command (`rate-limit icmp`) operate differently:

- All-traffic rate-limiting applies to both inbound and outbound traffic and can be specified either in terms of a percentage of total bandwidth or in terms of bits per second;
- ICMP rate-limiting applies only to inbound traffic and can be specified as only a percentage of total bandwidth.

Guidelines for configuring ICMP rate-limiting

Apply ICMP rate-limiting on all connected interfaces on the switch to effectively throttle excessive ICMP messaging from any source. **Figure 21: Example: of ICMP rate-limiting** on page 121 shows an Example: of how to configure this for a small to mid-sized campus though similar rate-limit thresholds are applicable to other network environments. On edge interfaces, where ICMP traffic should be minimal, a threshold of 1% of available bandwidth should be sufficient for most applications. On core interfaces, such as switch-to-switch and switch-to-router, a maximum threshold of 5% should be sufficient for normal ICMP traffic. ("Normal" ICMP traffic levels should be the maximums that occur when the network is rebooting.)

Figure 21: Example: of ICMP rate-limiting



Configuring ICMP rate-limiting

For detailed information about ICMP rate-limiting, see **ICMP rate-limiting** on page 120.

The `rate-limit icmp` command controls inbound usage of a port by setting a limit on the bandwidth available for inbound ICMP traffic.

Syntax:

```
[no] int <port-list> rate-limit icmp {< percent < 0-100 > | kbps < 0-10000000 > | [trap-clear>]}
```

Configures inbound ICMP traffic rate-limiting. You can configure a rate limit from either the global configuration level (as shown above) or from the interface context level. The `no` form of the command disables ICMP rate-limiting on the specified interfaces.

(Default: Disabled.)

<code>percent <1-100></code>	Values in this range allow ICMP traffic as a percentage of the bandwidth available on the interface.
<code>kbps <0-10000000></code>	Specifies the rate at which to forward traffic in kilobits per second.
<code>0</code>	Causes an interface to drop all incoming ICMP traffic and is not recommended. See the caution .
<code>trap-clear</code>	Clears existing ICMP rate limiting trap condition.

Note: ICMP rate-limiting is not supported on meshed ports. (Rate-limiting can reduce the efficiency of paths through a mesh domain).

Example:

Either of the following commands configures an inbound rate limit of 1% on ports A3 to A5, which are used as network edge ports:

```
switch(config) # int a3-a5 rate-limit icmp 1
switch(eth-A3-A5) # rate-limit icmp 1
```



When using kbps-mode ICMP rate-limiting, the rate-limiting only operates on the IP payload part of the ICMP packet (as required by metering RFC 2698). This means that effective metering is at a rate greater than the configured rate, with the disparity increasing as the packet size decreases (the packet to payload ratio is higher).

Also, in kbps mode, metering accuracy is limited at low values, For example, less than 45 Kbps. This is to allow metering to function well at higher media speeds such as 10 Gbps.

For information on using ICMP rate-limiting and all-traffic rate-limiting on the same interface, see [Using both ICMP rate-limiting and all-traffic rate-limiting on the same interface](#) on page 122.

Using both ICMP rate-limiting and all-traffic rate-limiting on the same interface

ICMP and all-traffic rate-limiting can be configured on the same interface. All-traffic rate-limiting applies to all inbound or outbound traffic (including ICMP traffic), while ICMP rate-limiting applies only to inbound ICMP traffic.



If the all-traffic load on an interface meets or exceeds the currently configured all-traffic inbound rate-limit while the ICMP traffic rate-limit on the same interface has not been reached, all excess traffic is dropped, including any inbound ICMP traffic above the all-traffic limit (regardless of whether the ICMP rate-limit has been reached).

Example:

Suppose:

- The all-traffic inbound rate-limit on port "X" is configured at 55% of the port's bandwidth.
- The ICMP traffic rate-limit on port "X" is configured at 2% of the port's bandwidth.

If at a given moment:

- Inbound ICMP traffic on port "X" is using 1% of the port's bandwidth, and
- Inbound traffic of all types on port "X" demands 61% of the ports's bandwidth,

all inbound traffic above 55% of the port's bandwidth, including any additional ICMP traffic, is dropped as long as all inbound traffic combined on the port demands 55% or more of the port's bandwidth.

Viewing the current ICMP rate-limit configuration

The `show rate-limit icmp` command displays the per-interface ICMP rate-limit configuration in the running-config file.

Syntax:

```
show rate-limit icmp [< port-list >]
```

Without `[port-list]`, this command lists the ICMP rate-limit configuration for all ports on the switch.

With `[port-list]`, this command lists the rate-limit configuration for the specified interfaces. This command operates the same way in any CLI context

If you want to view the rate-limiting configuration on ports 1–6:

Listing the rate-limit configuration

```
switch(config)# show rate-limit icmp 1-6

Inbound ICMP Rate Limit Maximum Percentage
```

Port	Mode	Rate Limit
1	Disabled	Disabled
2	kbps	100
3	%	5
4	%	1
5	%	1
6	Disabled	Disable

The `show running` command displays the currently applied setting for any interfaces in the switch configured for all traffic rate-limiting and ICMP rate-limiting.

The `show config` command displays this information for the configuration currently stored in the `startup-config` file. (Note that configuration changes performed with the CLI, but not followed by a `write mem` command, do not appear in the `startup-config` file.)

For more information on ICMP rate-limiting, see [Operating notes for ICMP rate-limiting](#) on page 123.

Operating notes for ICMP rate-limiting

ICMP rate-limiting operates on an interface (per-port) basis to allow, on average, the highest expected amount of legitimate, inbound ICMP traffic.

- **Interface support:** ICMP rate-limiting is available on all types of ports (other than trunk ports or mesh ports), and at all port speeds configurable for the switch.
- **Rate-limiting is not permitted on mesh ports:** Either type of rate-limiting (all-traffic or ICMP) can reduce the efficiency of paths through a mesh domain.
- **Rate-limiting is not supported on port trunks:** Neither all-traffic nor ICMP rate-limiting are supported on ports configured in a trunk group.
- **ICMP percentage-based rate-limits are calculated as a percentage of the negotiated link speed:** For example, if a 100 Mbps port negotiates a link to another switch at 100 Mbps and is ICMP rate-limit configured

at 5%, the inbound ICMP traffic flow through that port is limited to 5 Mbps. Similarly, if the same port negotiates a 10 Mbps link, it allows 0.5 Mbps of inbound traffic. If an interface experiences an inbound flow of ICMP traffic in excess of its configured limit, the switch generates a log message and an SNMP trap (if an SNMP trap receiver is configured).

- **ICMP rate-limiting is port-based:** ICMP rate-limiting reflects the available percentage of an interface's entire inbound bandwidth. The rate of inbound flow for traffic of a given priority and the rate of flow from an ICMP rate-limited interface to a particular queue of an outbound interface are not measures of the actual ICMP rate limit enforced on an interface.
- **Below-maximum rates:** ICMP rate-limiting operates on a per-interface basis, regardless of traffic priority. Configuring ICMP rate-limiting on an interface where other features affect inbound port queue behavior (such as flow control) can result in the interface not achieving its configured ICMP rate-limiting maximum. For example, in some situations with flow control configured on an ICMP rate-limited interface, there can be enough "back pressure" to hold high-priority inbound traffic from the upstream device or application to a rate that does not allow bandwidth for lower-priority ICMP traffic. In this case, the inbound traffic flow may not permit the forwarding of ICMP traffic into the switch fabric from the rate-limited interface. (This behavior is termed "head-of-line blocking" and is a well-known problem with flow-control.) In cases where both types of rate-limiting (`rate-limit all` and `rate-limit icmp`) are configured on the same interface, this situation is more likely to occur.

In another type of situation, an outbound interface can become oversubscribed by traffic received from multiple ICMP rate-limited interfaces. In this case, the actual rate for traffic on the rate-limited interfaces may be lower than configured because the total traffic load requested to the outbound interface exceeds the interface's bandwidth, and thus some requested traffic may be held off on inbound.

- **Monitoring (mirroring) ICMP rate-limited interfaces:** If monitoring is configured, packets dropped by ICMP rate-limiting on a monitored interface are still forwarded to the designated monitor port. (Monitoring shows what traffic is inbound on an interface, and is not affected by "drop" or "forward" decisions.)
- **Optimum rate-limiting operation:** Optimum rate-limiting occurs with 64-byte packet sizes. Traffic with larger packet sizes can result in performance somewhat below the configured inbound bandwidth. This is to ensure the strictest possible rate-limiting of all sizes of packets.
- **Outbound traffic flow:** Configuring ICMP rate-limiting on an interface does **not** control the rate of outbound traffic flow on the interface.

Notes on testing ICMP rate-limiting

ICMP rate-limiting is applied to the available bandwidth on an interface. If the total bandwidth requested by all ICMP traffic is less than the available, configured maximum rate, no ICMP rate-limit can be applied. That is, an interface must be receiving more inbound ICMP traffic than the configured bandwidth limit allows. If the interface is configured with both `rate-limit all` and `rate-limit icmp`, the ICMP limit can be met or exceeded only if the rate limit for all types of inbound traffic has not already been met or exceeded. Also, to test the ICMP limit you need to generate ICMP traffic that exceeds the configured ICMP rate limit. Using the recommended settings—1% for edge interfaces and 5% maximum for core interfaces—it is easy to generate sufficient traffic. However, if you are testing with higher maximums, you need to ensure that the ICMP traffic volume exceeds the configured maximum.

When testing ICMP rate-limiting where inbound ICMP traffic on a given interface has destinations on multiple outbound interfaces, the test results must be based on the received outbound ICMP traffic.

ICMP rate-limiting is not reflected in counters monitoring inbound traffic because inbound packets are counted before the ICMP rate-limiting drop action occurs.

ICMP rate-limiting trap and Event Log messages

If the switch detects a volume of inbound ICMP traffic on a port that exceeds the ICMP rate-limit configured for that port, it generates one SNMP trap and one informational Event Log message to notify the system operator of the condition. (The trap and Event Log message are sent within two minutes of when the event occurred on the port.) For Example:

```
I 06/30/05 11:15:42 RateLim: ICMP traffic exceeded configured limit on port A1
```

These trap and Event Log messages provide an advisory that inbound ICMP traffic on a given interface has exceeded the configured maximum. The additional ICMP traffic is dropped, but the excess condition may indicate an infected host (or other traffic threat or network problem) on that interface. The system operator should investigate the attached devices or network conditions further; the switch does not send more traps or Event Log messages for excess ICMP traffic on the affected port until the system operator resets the port's ICMP trap function.

The switch does not send more traps or Event Log messages for excess ICMP traffic on the affected port until the system operator resets the port's ICMP trap function. The reset can be done through SNMP from a network management station or through the CLI with the `trap-clear` command option.

Syntax:

```
interface <port-list> rate-limit icmp trap-clear
```

On a port configured with ICMP rate-limiting, this command resets the ICMP trap function, which allows the switch to generate a new SNMP trap and an Event Log message if ICMP traffic in excess of the configured limit is detected on the port.

Example:

An operator noticing an ICMP rate-limiting trap or Event Log message originating with port 1 on a switch would use the following command to reset the port to send a new message if the condition occurs again:

```
HP Switch(config)# interface 1 rate-limit icmp trap-clear
```

Determining the switch port number used in ICMP port reset commands

To enable excess ICMP traffic notification traps and Event Log messages, use the `setmib` command described on **ICMP rate-limiting trap and Event Log messages** on page 124. The port number included in the command corresponds to the internal number the switch maintains for the designated port and not the port's external identity.

To match the port's external slot/number to the internal port number, use the `walkmib ifDescr` command, as shown in the following example:

Matching internal port numbers to external port numbers

```
switch# walkmib ifDescr
ifDescr.1 = 1
ifDescr.2 = 2
ifDescr.3 = 3
ifDescr.4 = 4
ifDescr.5 = 5
ifDescr.6 = 6
ifDescr.7 = 7
ifDescr.8 = 8
ifDescr.9 = 9
ifDescr.10 = 10
ifDescr.11 = 11
ifDescr.12 = 12
ifDescr.13 = 13
ifDescr.14 = 14
ifDescr.15 = 15
ifDescr.16 = 16
ifDescr.17 = 17
ifDescr.18 = 18
ifDescr.19 = 19
ifDescr.20 = 20
ifDescr.21 = 21
```

```
ifDescr.22 = 22
ifDescr.23 = 23
ifDescr.24 = 24
ifDescr.210 = Trk1
ifDescr.211 = Trk2
ifDescr.330 = DEFAULT_VLAN
ifDescr.4425 = HP Switch software loopback interface
ifDescr.4426 = HP Switch software loopback interface
.
.
.
```

Configuring inbound rate-limiting for broadcast and multicast traffic

You can configure rate-limiting (throttling) of inbound broadcast and multicast traffic on the switch, which helps prevent the switch from being disrupted by traffic storms if they occur on the rate-limited port. The rate-limiting is implemented as a percentage of the total available bandwidth on the port.

The `rate-limit` command can be executed from the global or interface context, for Example:

```
switch(config)# interface 3 rate-limit bcast in percent 10
```

or

```
switch(config)# interface 3
switch(eth-3)# rate-limit bcast in percent 10
```

Syntax:

```
rate-limit {< bcast | mcast >} in percent < 0-100 >
```

Option

```
in percent <0-100>
```

Also supports configuring limit in *kbps*

```
[no] rate-limit {<bcast | [mcast >]} in
```

Enables rate-limiting and sets limits for the specified inbound broadcast or multicast traffic. Only the amount of traffic specified by the percent is forwarded.

Default: Disabled

If you want to set a limit of 50% on inbound broadcast traffic for port 3, you can first enter interface context for port 3 and then execute the `rate-limit` command, as shown in [Inbound broadcast rate-limiting of 50% on port 3](#) on page 126. Only 50% of the inbound broadcast traffic will be forwarded.

Inbound broadcast rate-limiting of 50% on port 3

```
switch(config)# int 3
switch(eth-3)# rate-limit bcast in percent 50

switch(eth-3)# show rate-limit bcast
Broadcast-Traffic Rate Limit Maximum %
```

Port	Inbound Limit	Mode	Radius Override
1	Disabled	Disabled	No-override
2	Disabled	Disabled	No-override
3	Disabled	%	No-override
4	Disabled	Disabled	No-override
5	Disabled	Disabled	No-override

If you rate-limit multicast traffic on the same port, the multicast limit is also in effect for that port, as shown in **Inbound multicast rate-limiting of 20% on port 3** on page 127. Only 20% of the multicast traffic will be forwarded.

Inbound multicast rate-limiting of 20% on port 3

```
switch(eth-3)# rate-limit mcast in percent 20
switch(eth-3)# show rate-limit mcast
```

Multicast-Traffic Rate Limit Maximum %

Port	Inbound Limit	Mode	Radius Override
1	Disabled	Disabled	No-override
2	Disabled	Disabled	No-override
3	20	%	No-override
4	Disabled	Disabled	No-override

To disable rate-limiting for a port enter the `no` form of the command, as shown in **Disabling inbound multicast rate-limiting for port 3** on page 127.

Disabling inbound multicast rate-limiting for port 3

```
switch(eth-3)# no rate-limit mcast in
```

```
switch(eth-3)# show rate-limit mcast
```

Multicast-Traffic Rate Limit Maximum %

Port	Inbound Limit	Mode	Radius Override
1	Disabled	Disabled	No-override
2	Disabled	Disabled	No-override
3	Disabled	Disabled	No-override
4	Disabled	Disabled	No-override

Operating Notes

The following information is displayed for each installed transceiver:

- Port number on which transceiver is installed.
- Type of transceiver.
- Product number — Includes revision letter, such as A, B, or C. If no revision letter follows a product number, this means that no revision is available for the transceiver.
- Part number — Allows you to determine the manufacturer for a specified transceiver and revision number.
- For a non-HPE switches installed transceiver (see **line 23 of "The show tech transceivers command" example**), no transceiver type, product number, or part information is displayed. In the Serial Number field, `non-operational` is displayed instead of a serial number.
- The following error messages may be displayed for a non-operational transceiver:
 - Unsupported Transceiver. (SelfTest Err#060)
 - This switch only supports revision B and above transceivers.

- Self test failure.
- Transceiver type not supported in this port.
- Transceiver type not supported in this software version.
- Not an HPE Switch Transceiver.

Jumbo frames

The maximum transmission unit(MTU) is the maximum size IP frame the switch can receive for Layer 2 frames inbound on a port. The switch drops any inbound frames larger than the MTU allowed on the port. Ports operating at a minimum of 1 Gbps can accept forward frames of up to 9220 bytes (including four bytes for a VLAN tag) when configured for jumbo traffic. You can enable inbound jumbo frames on a per-VLAN basis. That is, on a VLAN configured for jumbo traffic, all ports belonging to that VLAN and **operating** at a minimum of 1 Gbps allow inbound jumbo frames of up to 9220 bytes.

Operating rules

- **Required port speed:** This feature allows inbound and outbound jumbo frames on ports operating at a minimum of 1 Gbps.
- **GVRP operation:** A VLAN enabled for jumbo traffic cannot be used to create a dynamic VLAN. A port belonging to a statically configured, jumbo-enabled VLAN cannot join a dynamic VLAN.
- **Port adds and moves:** If you add a port to a VLAN that is already configured for jumbo traffic, the switch enables that port to receive jumbo traffic. If you remove a port from a jumbo-enabled VLAN, the switch disables jumbo traffic capability on the port only if the port is not currently a member of another jumbo-enabled VLAN. This same operation applies to port trunks.
- **Jumbo traffic sources:** A port belonging to a jumbo-enabled VLAN can receive inbound jumbo frames through any VLAN to which it belongs, including non-jumbo VLANs. For example, if VLAN 10 (without jumbos enabled) and VLAN 20 (with jumbos enabled) are both configured on a switch, and port 1 belongs to both VLANs, port 1 can receive jumbo traffic from devices on either VLAN. For a method to allow only some ports in a VLAN to receive jumbo traffic, see [Configuring a maximum frame size](#) on page 131.

Jumbo traffic-handling

- HPE does not recommend configuring a voice VLAN to accept jumbo frames. Voice VLAN frames are typically small, and allowing a voice VLAN to accept jumbo frame traffic can degrade the voice transmission performance.
- You can configure the default, primary, and/or (if configured) the management VLAN to accept jumbo frames on all ports belonging to the VLAN.
- When the switch applies the default MTU (1522-bytes including 4 bytes for the VLAN tag) to a VLAN, all ports in the VLAN can receive incoming frames of up to 1522 bytes. When the switch applies the jumbo MTU (9220 bytes including 4 bytes for the VLAN tag) to a VLAN, all ports in that VLAN can receive incoming frames of up to 9220 bytes. A port receiving frames exceeding the applicable MTU drops such frames, causing the switch to generate an Event Log message and increment the "Giant Rx" counter (displayed by `show interfaces <port-list>`).
- The switch allows flow control and jumbo frame capability to co-exist on a port.
- The default MTU is 1522 bytes (including 4 bytes for the VLAN tag). The jumbo MTU is 9220 bytes (including 4 bytes for the VLAN tag).
- When a port is not a member of any jumbo-enabled VLAN, it drops all jumbo traffic. If the port is receiving "excessive" inbound jumbo traffic, the port generates an Event Log message to notify you of this condition. This same condition also increments the switch's "Giant Rx" counter.
- If you do not want all ports in a given VLAN to accept jumbo frames, you can consider creating one or more jumbo VLANs with a membership comprising only the ports you want to receive jumbo traffic. Because a port belonging to one jumbo-enabled VLAN can receive jumbo frames through any VLAN to which it belongs, this method enables you to include both jumbo-enabled and non-jumbo ports within the same VLAN.

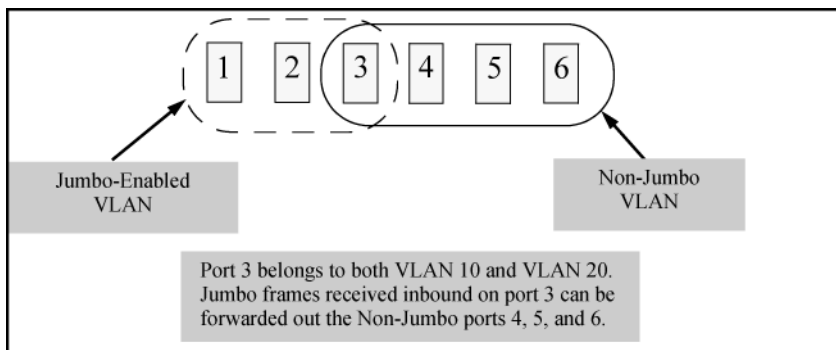
For example, suppose you want to allow inbound jumbo frames only on ports 6, 7, 12, and 13. However, these ports are spread across VLAN 100 and VLAN 200 and also share these VLANs with other ports you want excluded from jumbo traffic. A solution is to create a third VLAN with the sole purpose of enabling jumbo traffic on the desired ports, while leaving the other ports on the switch disabled for jumbo traffic. That is:

	VLAN 100	VLAN 200	VLAN 300
Ports	6-10	11-15	6, 7, 12, and 13
Jumbo-enabled?	No	No	Yes

If there are security concerns with grouping the ports as shown for VLAN 300, you can either use source-port filtering to block unwanted traffic paths or create separate jumbo VLANs, one for ports 6 and 7, and another for ports 12 and 13.

- **Outbound jumbo traffic.** Any port operating at 1 Gbps or higher can transmit outbound jumbo frames through any VLAN, regardless of the jumbo configuration. The VLAN is not required to be jumbo-enabled, and the port is not required to belong to any other, jumbo-enabled VLANs. This can occur in situations where a non-jumbo VLAN includes some ports that do not belong to another, jumbo-enabled VLAN and some ports that do belong to another, jumbo-enabled VLAN. In this case, ports capable of receiving jumbo frames can forward them to the ports in the VLAN that do not have jumbo capability, as shown in **Figure 22: Forwarding jumbo frames through non-jumbo ports** on page 129.

Figure 22: Forwarding jumbo frames through non-jumbo ports



Jumbo frames can also be forwarded out non-jumbo ports when the jumbo frames received inbound on a jumbo-enabled VLAN are routed to another, non-jumbo VLAN for outbound transmission on ports that have no memberships in other, jumbo-capable VLANs. Where either of the above scenarios is a possibility, the downstream device must be configured to accept the jumbo traffic. Otherwise, this traffic will be dropped by the downstream device.

Configuring jumbo frame operation

For detailed information about jumbo frames, see **Jumbo frames** on page 128.

Overview

1. Determine the VLAN membership of the ports or trunks through which you want the switch to accept inbound jumbo traffic. For operation with GVRP enabled, refer to the GVRP topic under “Operating Rules”, above.
2. Ensure that the ports through which you want the switch to receive jumbo frames are operating at least at gigabit speed. (Check the Mode field in the output for the `show interfaces brief <port-list>` command.)
3. Use the `jumbo` command to enable jumbo frames on one or more VLANs statically configured in the switch. (All ports belonging to a jumbo-enabled VLAN can receive jumbo frames.)
4. Execute `write memory` to save your configuration changes to the `startupconfig` file.

Viewing the current jumbo configuration

Syntax:

```
show vlans
```

Lists the static VLANs configured on the switch and includes a Jumbo column to indicate which VLANs are configured to support inbound jumbo traffic. All ports belonging to a jumbo-enabled VLAN can receive jumbo traffic. (For more information, see [Configuring a maximum frame size](#) on page 131.) See [Figure 23: Example: listing of static VLANs to show jumbo status per VLAN](#) on page 130.

Figure 23: Example: listing of static VLANs to show jumbo status per VLAN

```
HP Switch(config)# show vlans
Status and Counters - VLAN Information

Maximum VLANs to support : 256
Primary VLAN : DEFAULT_VLAN
Management VLAN :
```

VLAN ID	Name	Status	Voice	Jumbo
1	DEFAULT_VLAN	Port-based	No	Yes
5	VLAN5	Port-based	No	No
22	VLAN22	Port-based	No	No

Indicates which static VLANs are configured to enable jumbo frames.

Syntax:

```
show vlans ports <port-list>
```

Lists the static VLANs to which the specified ports belong, including the Jumbo column to indicate which VLANs are configured to support jumbo traffic.

Entering only one port in <port-list > results in a list of all VLANs to which that port belongs.

Entering multiple ports in <port-list > results in a superset list that includes the VLAN memberships of all ports in the list, even though the individual ports in the list may belong to different subsets of the complete VLAN listing.

Example:

If port 1 belongs to VLAN 1, port 2 belongs to VLAN 10, and port 3 belongs to VLAN 15, executing this command with a *port-list* of 1 - 3 results in a listing of all three VLANs, even though none of the ports belong to all three VLANs. (See [Figure 24: Listing the VLAN memberships for a range of ports](#) on page 130.)

Figure 24: Listing the VLAN memberships for a range of ports

```
HP Switch(config)# show vlans ports A1-A3
Status and Counters - VLAN Information - for ports A1-A3
```

VLAN ID	Name	Status	Voice	Jumbo
1	DEFAULT_VLAN	Port-based	No	Yes
10	VLAN10	Port-based	No	No
15	VLAN15	Port-based	No	No

Indicates which static VLANs are configured to enable jumbo frames.

Syntax:

```
show vlans <vid>
```

Shows port membership and jumbo configuration for the specified *vid* . (See **Figure 25: Example: of listing the port membership and jumbo status for a VLAN** on page 131.)

Figure 25: Example: of listing the port membership and jumbo status for a VLAN

```
HP Switch(config)# show vlan 100
Status and Counters - VLAN Information - VLAN 100
VLAN ID : 100
Name : VLAN100
Status : Port-based Voice : No
Jumbo : No
Port Information Mode      Unknown VLAN Status
-----
A1      Tagged    Learn    Up
A2      Tagged    Learn    Up
A3      Tagged    Learn    Up
A4      Tagged    Learn    Down
A5      Tagged    Learn    Up
```

Lists the ports belonging to VLAN 100 and whether the VLAN is enabled for jumbo frame traffic.

Enabling or disabling jumbo traffic on a VLAN

Syntax:

```
vlan <vid> jumbo
```

```
[no] vlan <vid> jumbo
```

Configures the specified VLAN to allow jumbo frames on all ports on the switch that belong to that VLAN. If the VLAN is not already configured on the switch, `vlan <vid> jumbo` also creates the VLAN.

A port belonging to one jumbo VLAN can receive jumbo frames through any other VLAN statically configured on the switch, regardless of whether the other VLAN is enabled for jumbo frames.

The `[no]` form of the command disables inbound jumbo traffic on all ports in the specified VLAN that do not also belong to another VLAN that is enabled for jumbo traffic. In a VLAN context, the command forms are `jumbo` and `no jumbo`.

(Default: Jumbos disabled on the specified VLAN.)

Configuring a maximum frame size

You can globally set a maximum frame size for jumbo frames that will support values from 1518 bytes to 9216 bytes for untagged frames.

Syntax:

```
jumbo max-frame-size <size>
```

Sets the maximum frame size for jumbo frames. The range is from 1518 bytes to 9216 bytes. (Default: 9216 bytes)



The `jumbo max-frame-size` is set on a GLOBAL level.
Default: 9216 bytes

Configuring IP MTU



The following feature is available on the switches covered in this guide. `jumbos` support is required for this feature. On switches that do not support this command, the IP MTU value is derived from the maximum frame size and is not configurable.

You can set the IP MTU globally by entering this command. The value of `max-frame-size` must be greater than or equal to 18 bytes more than the value selected for `ip-mtu`. For example, if `ip-mtu` is set to 8964, the `max-frame-size` is configured as 8982.

Syntax:

```
jumbo ip-mtu <size>
```

Globally sets the IP MTU size. Values range between 1500 and 9198 bytes. This value must be 18 bytes less than the value of `max-frame-size`.

(Default: 9198 bytes)

SNMP implementation

Jumbo maximum frame size

The maximum frame size for jumbos is supported with the following proprietary MIB object:

```
hpSwitchMaxFrameSize OBJECT-TYPE
```

This is the value of the global `max-frame-size` supported by the switch. The default value is set to 9216 bytes.

Jumbo IP MTU

The IP MTU for jumbos is supported with the following proprietary MIB object:

```
hpSwitchIpMTU OBJECT-TYPE
```

This is the value of the global jumbos IP MTU (or L3 MTU) supported by the switch. The default value is set to 9198 bytes (a value that is 18 bytes less than the largest possible maximum frame size of 9216 bytes). This object can be used only in switches that support `max-frame-size` and `ip-mtu` configuration.

Displaying the maximum frame size

Use the `show jumbos` command to display the globally configured untagged maximum frame size for the switch, as shown in the following Example:.

```
switch(config)# show jumbos
```

```
Jumbos Global Values
```

```
Configured : MaxFrameSize : 9216   Ip-MTU : 9198
In Use      : MaxFrameSize : 9216   Ip-MTU : 9198
```

For more information about frame size, see [Jumbo frames](#) on page 128.

Operating notes for maximum frame size

- When you set a maximum frame size for jumbo frames, it must be on a global level. You cannot use the `jumbo max-frame-size` command on a per-port or per-VLAN basis.
- The original way to configure jumbo frames remains the same, which is per-VLAN, but you cannot set a maximum frame size per-VLAN.
- Jumbo support must be enabled for a VLAN from the CLI or through SNMP.
- Setting the maximum frame size does not require a reboot.

- When you upgrade to a version of software that supports setting the maximum frame size from a version that did not, the `max-frame-size` value is set automatically to 9216 bytes.
- Configuring a jumbo maximum frame size on a VLAN allows frames up to `max-frame-size` even though other VLANs of which the port is a member are not enabled for jumbo support.

Troubleshooting

A VLAN is configured to allow jumbo frames, but one or more ports drops all inbound jumbo frames

The port may not be operating at a minimum of 1 Gbps on the other switches covered in this guide. Regardless of a port's configuration, if it is actually operating at a speed lower than 1 Gbps for the other switches, it drops inbound jumbo frames. For example, if a port is configured for `Auto` mode (`speed-duplex auto`), but has negotiated a 7 Mbps speed with the device at the other end of the link, the port cannot receive inbound jumbo frames. To determine the actual operating speed of one or more ports, view the `Mode` field in the output for the following command:

```
show interfaces brief <port-list>
```

A non-jumbo port is generating "Excessive undersize/giant frames" messages in the Event Log

The switches can transmit outbound jumbo traffic on any port, regardless of whether the port belongs to a jumbo VLAN. In this case, another port in the same VLAN on the switch may be jumbo-enabled through membership in a different, jumbo-enabled VLAN, and may be forwarding jumbo frames received on the jumbo VLAN to non-jumbo ports.

Overview

Detection of link-flap and taking action on the port is done via fault-finder command at 3 different sensitivity levels (low, medium and high). The configuration in fault-finder for link-flap is a global configuration affecting all ports on the switch/stack. To provide further granularity to link-flap detection and action which provides different link-flap detection and action configuration for each port rather than the same configuration for all ports on the switch/stack. The per-port configuration will supersede the global configuration for fault-finder link-flap.

A configurable option to re-enable ports disabled by link-flap after a waiting period is also been added. The waiting period time is expressed in unit of seconds in the range 0 to 604800. Maximum allowed waiting period is one week. Zero is the default value, meaning that the port will not be re-enabled automatically.



A very important point is the wording of “link-flap” itself – i.e. the word “link”. This condition should be at the link/port-level granular, allowing alerts and actions only on those certain links/ports where the functionality is needed.

Fault-finder link-flap

Syntax

In the config context:

```
[no] fault-finder link-flap [ethernet] PORT-LIST action warn | warn-and-disable SECONDS sensitivity low | medium | high
```

Description

Configures the link-flap on a port. The default value is `warn`.

Options

link-flap	Configure link-flap control.
warn	Log the event only.
warn-and-disable	Log the event and disable the port.
seconds	Re-enable the port after waiting for the specified number of seconds. The default value is 0, which indicates that the port will not be automatically enabled.
sensitivity	Indicate the sensitivity of the link-flap control threshold within a 10-second interval. <ul style="list-style-type: none"> • Low indicates 10 link-flaps. • Medium indicates 6 link-flaps. • High indicates 3 link-flaps.

Parameters

action	Configure the action taken when a fault is detected.
---------------	--

ethernet <i>PORT-LIST</i>	Enable link-flap control on a list of ports.
warn	Warn about faults found.
warn-and-disable	Warn and disable faulty component.
seconds	Configure the number of seconds for which the port remains disabled. A value of 0 means that the port will remain disabled until manually re-enabled.
sensitivity	Configure the fault sensitivity level.
low	Low sensitivity.
medium	Medium sensitivity
high	High sensitivity.

Subcommand Syntax

[no] fault-finder link-flap ethernet *PORT-LIST*

Description

To remove the current configuration of link-flap on a port

Usage

Enable a linkFault-Finder check and set parameters for it. These commands may be repeated to enable additional checks. The default sensitivity is medium and the default action is warn.

[no] fault-finder all | fault sensitivity low | medium | high action warn | warn-and-disable

[no] fault-finder broadcast-storm sensitivity low | medium | high action warn | warn-and-disable *SECONDS*

[no] fault-finder link-flap sensitivity low | medium | high action warn | warn-and-disable

[no] fault-finder link-flap *PORT-LIST* action warn | warn-and-disable *SECONDS* sensitivity low | medium | high

Configure ports for link-flap detection with high sensitivity

Configure ports A1 to A5 for link-flap detection with sensitivity of high (3 flaps over 10s) and to log and disable port for 65535s if the link-flap threshold is exceeded.

```
switch(config)# fault-finder link-flap ethernet A1-A5 action warn-and-disable
65535
sensitivity high
```

Configure ports for link-flap detection with medium sensitivity

Configure ports A8 for link-flap detection with sensitivity of medium (6 flaps over 10s) and to log and disable port if the link-flap threshold is exceeded. User will need to re-enable the port if disabled.

```
switch(config)# fault-finder link-flap ethernet A8 action warn-and-disable 0
sensitivity medium
```

Configure ports for link-flap detection with low sensitivity

Configure ports A22 for link-flap detection with sensitivity of low (10 flaps over 10s) and to log if the link-flap threshold is exceeded

```
switch(config)# fault-finder link-flap ethernet A22 action warn sensitivity low
```

Disable link-flap detection

Disable link-flap detection for port A5

```
switch(config)# no fault-finder link-flap ethernet A5
```

Show fault-finder link-flap

Syntax

```
show fault-finder link-flap ethernet PORT-LIST
```

Description

Display the link-flap control configuration.

Show fault-finder link-flap

```
switch# show fault-finder link-flap A1
```

Port	Link Flap	Port Status	Sensitivity	Action	Disable Timer	Disable Time Left
A1	Yes	Down	Low	warn-and-disable	65535	45303


```
switch# show fault-finder link-flap
```

Port	Link Flap	Port Status	Sensitivity	Action	Disable Timer	Disable Time Left
A1	Yes	Down	Low	warn-and-disable	65535	45303
A5	No	Up	None	None	-	-
A22	Yes	Down	Low	warn-and-disable	-	-
A23	Yes	Down	High	warn-and-disable	100	-



This example displays only the list of ports configured via the above per-port config commands, does not include the global configuration ports.

Event Log

Cause

Message	Cause
FFI: port <ID>- Excessive link state transitions.	Link-flap is detected by fault-finder per the sensitivity configured.
FFI: port <ID>- Excessive link state transitions.FFI: port <ID>-Port disabled by Fault-finder.FFI: port <ID>-Administrator action is required to re-enable.ports: Fault-finder (71) has disabled port <ID>.ports: port <ID> is now offline.vlan: VLAN<VLAN-ID> virtual LAN is disabled.	Link-flap is detected and the action is to disable the port with no disable timer.
FFI: port <ID>- Excessive link state transitions.FFI: port <ID>-Port disabled by Fault-finder.ports: Fault-finder(71) has disabled port <ID> for <SECONDS> seconds.ports: port <ID> is now off-line.vlan: VLAN<VLAN-ID> virtual LAN is disabled.	Link-flap is detected and the action is to disable the port with disable timer.
port <ID> timer (71) has expired.ports: port <ID> is now on-line.vlan: VLAN<VLAN-ID> virtual LAN is enabled.	The port is enabled when the disable timer expires.

Restrictions

- Per port configuration for options – link-flap only. Global settings for other options.
- No support for menu interface.
- No support for Web UI.
- No changes to PCM.
- No changes to IDM.
- No support for trunks.

Using SNMP tools to manage the switch

SNMP is a management protocol that allows an SNMP client application to retrieve device configuration and status information and to configure the device (**get** and **set**). You can manage the switch via SNMP from a network management station running an application such as PCM+. For more information on PCM+, see the HPE website at: <http://www.hpe.com/networking>.

From the **Products** menu, select **Network Management**. Then click on **PCM+ Network Management** under the **HPE Network Management** bar.

To implement SNMP management, the switch must have an IP address configured either manually or dynamically (using DHCP or Bootp). If multiple VLANs are configured, each VLAN interface should have its own IP address. For DHCP use with multiple VLANs, see section "The Primary VLAN" in the "Static Virtual LANs (VLANs)" of the advanced traffic management guide for your switch.



If you use the switch's Authorized IP Managers and Management VLAN features, ensure that the SNMP management station, the choice of switch port used for SNMP access to the switch, or both, are compatible with the access controls enforced by these features. Otherwise, SNMP access to the switch will be blocked.

For more information on Authorized IP Managers, see the access security guide for your switch. (The latest version of this guide is available on the HPE Networking website.) For information on the Management VLAN feature, see the section "The Secure Management VLAN" in the "Static Virtual LANs (VLANs)" chapter of the advanced traffic management guide for your switch.

SNMP management features

SNMP management features on the switch include:

- SNMP version 1, version 2c, or version 3 over IP
- Security via configuration of SNMP communities (**SNMPv3 communities** on page 145)
- Security via authentication and privacy for SNMPv3 access
- Event reporting via SNMP
 - Version 1 traps
 - RMON: groups 1, 2, 3, and 9
- PCM/PCM+
- Flow sampling using sFlow
- Standard MIBs, such as the Bridge MIB (RFC 1493), Ethernet MAU MIB (RFC 1515), and others.

The switch SNMP agent also uses certain variables that are included in an HPE proprietary MIB (management information base) file. If you are using HPE OpenView, you can ensure that it is using the latest version of the MIB file by downloading the file to the OpenView database. To do so, go to the HPE Networking website at: <http://www.hpe.com/networking>.

1. Type a model number of your switch (For example, 8212) or product number in the **Auto Search** text box.
2. Select an appropriate product from the drop down list.
3. Click the Display selected button.
4. From the options that appear, select Software downloads.
5. MIBs are available with switch software in the Other category.

Click on `software updates`, then MIBs.

SNMPv1 and v2c access to the switch

SNMP access requires an IP address and subnet mask configured on the switch. If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address.

Once an IP address is configured, the main steps for configuring SNMPv1 and v2c access management features are:

Procedure

1. Configure the appropriate SNMP communities. (See **SNMPv3 communities** on page 145.)
2. Configure the appropriate trap receivers.

In some networks, authorized IP manager addresses are not used. In this case, all management stations using the correct community name may access the switch with the View and Access levels that have been set for that community. If you want to restrict access to one or more specific nodes, you can use the switch's IP Authorized Manager feature. (See the access security guide for your switch.)



For PCM/PCM+ version 1.5 or earlier (or any TopTools version), deleting the "public" community disables some network management functions (such as traffic monitoring, SNMP trap generation, and threshold setting). If network management security is a concern, and you are using the above software versions, Hewlett Packard Enterprise recommends that you change the write access for the "public" community to "Restricted."

SNMPv3 access to the switch

SNMPv3 access requires an IP address and subnet mask configured on the switch. (See "IP Configuration" on page 8-2.) If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address. (See "DHCP/Bootp Operation".)

Once you have configured an IP address, the main steps for configuring SNMPv3 access management features are the following:

Procedure

1. Enable SNMPv3 for operation on the switch (see **Enabling SNMPv3** on page 140).
2. Configure the appropriate SNMP users (see **SNMPv3 users** on page 141).
3. Configure the appropriate SNMP communities (see **SNMPv3 communities** on page 145).
4. Configure the appropriate trap receivers (see **SNMP notifications** on page 148).

In some networks, authorized IP manager addresses are not used. In this case, all management stations using the correct User and community name may access the switch with the View and Access levels that have been set for that community. If you want to restrict access to one or more specific nodes, you can use the IP Authorized Manager feature for the switch. (See the access security guide for your switch.)

SNMP version 3 (SNMPv3) adds some new commands to the CLI for configuring SNMPv3 functions. To enable SNMMPv3 operation on the switch, use the `snmpv3 enable` command. An initial user entry will be generated with MD5 authentication and DES privacy.

You may (optionally) restrict access to only SNMPv3 agents by using the `snmpv3 only` command. To restrict write-access to only SNMPv3 agents, use the `snmpv3 restricted-access` command.



Restricting access to only version 3 messages will make the community named "public" inaccessible to network management applications (such as autodiscovery, traffic monitoring, SNMP trap generation, and threshold setting) from operating in the switch.

Enabling and disabling switch for access from SNMPv3 agents

This includes the creation of the initial user record.

Syntax:

```
[no] snmpv3 enable
```

Enabling or disabling restrictions to access from only SNMPv3 agents

When enabled, the switch rejects all non-SNMPv3 messages.

Syntax:

```
[no] snmpv3 only
```

Enabling or disabling restrictions from all non-SNMPv3 agents to read-only access

Syntax:

```
[no] snmpv3 restricted-access
```

Viewing the operating status of SNMPv3

Syntax:

```
show snmpv3 enable
```

Viewing status of message reception of non-SNMPv3 messages

Syntax:

```
show snmpv3 only
```

Viewing status of write messages of non-SNMPv3 messages

Syntax:

```
show snmpv3 restricted-access
```

Enabling SNMPv3

The `snmpv3 enable` command allows the switch to:

- Receive SNMPv3 messages.
- Configure initial users.
- Restrict non-version 3 messages to "read only" (optional).



Restricting access to only version 3 messages makes the community named "public" inaccessible to network management applications (such as autodiscovery, traffic monitoring, SNMP trap generation, and threshold setting) from running on the switch.

Example:

SNMP version 3 enable command

```
HP Switch(config)# snmpv3 enable
SHMPv3 Initialization process.
Creating user 'initial'
Authentication Protocol: MD5
Enter authentication password: *****
Privacy protocol is DES
Enter privacy password: *****

User 'initial' is created
Would you like to create a user that uses SHA? y
Enter user name: templateSHA
Authentication Protocol: SHA
Enter authentication password: *****
Privacy protocol is DES
Enter privacy password: *****

User creation is done. SHMPv3 is now functional.
Would you like to restrict SHMPv1 and SNMPv2c messages to have read only
access (you can set this later by the command 'snmp restrict-access'): n
```

Enable SNMPv3

Create initial user models for SNMPv3 Management Applications

Set restriction on non-SNMPv3 messages

SNMPv3 users



To create new users, most SNMPv3 management software requires an initial user record to clone. The initial user record can be downgraded and provided with fewer features, but not upgraded by adding new features. For this reason, Hewlett Packard Enterprise recommends that when you enable SNMPv3, you also create a second user with SHA authentication and DES privacy.

To use SNMPv3 on the switch, you must configure the users that will be assigned to different groups:

Procedure

1. Configure users in the User Table with the `snmpv3 user` command.

To view the list of configured users, enter the `show snmpv3 user` command (see [Adding users](#) on page 142)

2. Assign users to Security Groups based on their security model with the `snmpv3 group` command (see [Assigning users to groups \(CLI\)](#) on page 143).



If you add an SNMPv3 user without authentication, privacy, or both, to a group that requires either feature, the user will not be able to access the switch. Ensure that you add a user with the appropriate security level to an existing security group.

Adding users

To configure an SNMPv3 user, you must first add the user name to the list of known users with the `snmpv3 user` command, as shown in **Figure 26: Adding SNMPv3 users and displaying SNMPv3 configuration** on page 142.

Figure 26: Adding SNMPv3 users and displaying SNMPv3 configuration

```
HP Switch(config)# snmpv3 user NetworkAdmin
HP Switch(config)# snmpv3 user NetworkMgr auth md5 authpass priv privpass
HP Switch(config)# show snmpv3 user
```

Status and Counters - SNMP v3 Global Configuration Information

User Name	Auth. Protocol	Privacy Protocol
initial	MD5	CFB AES-128
NetworkAdmin	MD5	CBC-DES

Annotations:

- HP Switch(config)# snmpv3 user NetworkAdmin: Add user Network Admin with no authentication or privacy.
- HP Switch(config)# snmpv3 user NetworkMgr auth md5 authpass priv privpass: Add user Network Mgr with authentication and privacy.
- auth md5: MD5 authentication is enabled and the password is set to "authpass".
- priv privpass: Privacy is enabled and the password is set to "privpass".

SNMPv3 user commands

Syntax:

```
[no] snmpv3 user <user_name>
```

Adds or deletes a user entry for SNMPv3. Authorization and privacy are optional, but to use privacy, you must use authorization. When you delete a user, only the `user_name` is required.

```
[auth < {md5 | sha}> < auth_pass>]
```

With authorization, you can set either MD5 or SHA authentication. The authentication password `<auth_pass>` must be 6 to 32 characters and is mandatory when you configure authentication.

Default: None

```
[priv < {des | aes}> < priv_pass>]
```

With privacy, the switch supports DES (56-bit) and AES (128-bit) encryption. The privacy password `<priv_pass>` must be 6-32 characters in length and is mandatory when you configure privacy.

Default: DES



Only AES 128-bit and DES 56-bit encryption are supported as privacy protocols. Other non-standard encryption algorithms, such as AES-172, AES-256, and 3-DES are not supported.

Listing Users

To display the management stations configured to access the switch with SNMPv3 and view the authentication and privacy protocols that each station uses, enter the `show snmpv3 user` command.

Syntax:

```
show snmpv3 user
```

Display of the management stations configured on VLAN 1 on page 143 displays information about the management stations configured on VLAN 1 to access the switch.

Display of the management stations configured on VLAN 1

```
switch# configure terminal
switch(config)# vlan 1
switch(vlan-1)# show snmpv3 user
```

Status and Counters - SNMPv3 Global Configuration Information

User Name	Auth. Protocol	Privacy Protocol
initial	MD5	CFB AES-128
NetworkAdmin	MD5	CBC-DES

Assigning users to groups (CLI)

Next you must set the group access level for the user by assigning the user to a group. This is done with the `snmpv3 group` command, as shown in **Figure 27: Example: of assigning users to groups** on page 143. For more details on the MIBs access for a given group, see **Group access levels** on page 144.

Figure 27: Example: of assigning users to groups

```
Switch(config)# snmpv3 group operatornoauth user NetworkAdmin sec-model ver3
Switch(config)# snmpv3 group managerpriv user NetworkMgr sec-model ver3
Switch(config)# show snmpv3 group
```

Status and Counters - SNHP v3 Global Configuration Information

Security Name	Security Model	Group Name
CommunityManagerReadOnly	ver1	ComManagerR
CommunityManagerReadWrite	ver1	ComManagerRW
CommunityOperatorReadOnly	ver1	ComOperatorRW
CommunityOperatorReadWrite	ver1	ComOperatorRW
CommunityManagerReadOnly	ver2c	ComManagerR
CommunityManagerReadWrite	ver2c	ComManagerRW
CommunityOperatorReadOnly	ver2c	ComOperatorRW
CommunityOperatorReadWrite	ver2c	ComOperatorRW
NetworkMgr	ver3	ManagerPriv
NetworkAdmin	ver3	OperatorNoAuth

Syntax:

```
[no] snmpv3 group
```

Assigns or removes a user to a security group for access rights to the switch. To delete an entry, all of the following three parameters must be included in the command:

<code>group <group_name></code>	Identifies the group that has the privileges that will be assigned to the user. For more details, see Group access levels on page 144.
<code>user <user_name></code>	Identifies the user to be added to the access group. This must match the user name added with the <code>snmpv3 user</code> command.
<code>sec-model {<ver1 ver2c ver3></code>	Defines which security model to use for the added user. An SNMPv3 access group should use only the ver3 security model.

Group access levels

The switch supports eight predefined group access levels, shown in the following table. There are four levels for use by version 3 users and four are used for access by version 2c or version 1 management applications.

Table 17: Predefined group access levels

Group name	Group access type	Group read view	Group write view
managerpriv	Ver3 Must have Authentication and Privacy	ManagerReadView	ManagerWriteView
managerauth	Ver3 Must have Authentication	ManagerReadView	ManagerWriteView
operatorauth	Ver3 Must have Authentication	OperatorReadView	DiscoveryView
operatornoauth	Ver3 No Authentication	OperatorReadView	DiscoveryView
commanagerrw	Ver2c or Ver1	ManagerReadView	ManagerWriteView
commanagerr	Ver2c or Ver1	ManagerReadView	DiscoveryView
comoperatorrw	Ver2c or Ver1	OperatorReadView	OperatorReadView
comoperatorr	Ver2c or Ver1	OperatorReadView	DiscoveryView

Each view allows you to view or modify a different set of MIBs:

- **Manager Read View** – access to all managed objects
- **Manager Write View** – access to all managed objects except the following:
 - vacmContextTable
 - vacmAccessTable
 - vacmViewTreeFamilyTable
- **OperatorReadView** – no access to the following:
 - icfSecurityMIB
 - hpSwitchIpTftpMode
 - vacmContextTable
 - vacmAccessTable
 - vacmViewTreeFamilyTable
 - usmUserTable
 - snmpCommunityTable
- **Discovery View** – Access limited to samplingProbe MIB.



All access groups and views are predefined on the switch. There is no method to modify or add groups or views to those that are predefined on the switch.

SNMPv3 communities

SNMP communities are supported by the switch to allow management applications that use version 2c or version 1 to access the switch. The communities are mapped to Group Access Levels that are used for version 2c or version 1 support. This mapping happens automatically based on the communities access privileges, but special mappings can be added with the `snmpv3 community` command (see [Mapping SNMPv3 communities \(CLI\)](#) on page 145).

Mapping SNMPv3 communities (CLI)

SNMP communities are supported by the switch to allow management applications that use version 2c or version 1 to access the switch. For more details, see [SNMPv3 communities](#) on page 145.

Syntax:

[no] snmpv3 community

Maps or removes a mapping of a community name to a group access level. To remove a mapping you need to specify only the `index_name` parameter.

<code>index <index_name></code>	An index number or title for the mapping. The values of 1 to 5 are reserved and can not be mapped.
<code>name <community_name></code>	The community name that is being mapped to a group access level.
<code>sec-name <security_name></code>	The group level to which the community is being mapped.
<code>tag <tag_value></code>	This is used to specify which target address may have access by way of this index reference.

Example:

Figure 28: Assigning a community to a group access level on page 146 shows the assigning of the Operator community on MgrStation1 to the CommunityOperatorReadWrite group. Any other Operator has an access level of CommunityOperatorReadOnly.

Figure 28: Assigning a community to a group access level

```

Add mapping to allow write access for
Operator community on MgrStation1
HP Switch(config)# snmpv3 Community index 30 name Operator sec-name
CommunityManagerReadWrite tag MgrStation1
HP Switch(config)# show snmpv3 community
snmpCommunityTable [rfc2576]
Index Name          Community Name      Security Name
-----
1                   public             CommunityManagerReadWrite
2                   Operator           CommunityOperatorReadOnly
3                   Manager            CommunityManagerReadWrite
30                  Operator           CommunityManagerReadWrite
  
```

SNMP community features

Use SNMP communities to restrict access to the switch by SNMP management stations by adding, editing, or deleting SNMP communities. You can configure up to five SNMP communities, each with either an operator-level or a manager-level view and either restricted or unrestricted write access.

Using SNMP requires that the switch have an IP address and subnet mask compatible with your network.



For PCM/PCM+ version 1.5 or earlier (or any TopTools version), deleting the "public" community disables some network management functions (such as traffic monitoring, SNMP trap generation, and threshold setting). If network management security is a concern, and if you are using the above software versions, Hewlett Packard Enterprise recommends that you change the write access for the "public" community to "Restricted."

Viewing and configuring non-version-3 SNMP communities (Menu)

Procedure

1. From the Main Menu, select:
 2. **Switch Configuration...**
 6. **SNMP Community Names**

Figure 29: The SNMP Communities screen (default values)

Note: This screen gives an overview of the SNMP communities that are currently configured. All fields in this screen are read-only.

```

----- CONSOLE - MANAGER MODE -----
Switch Configuration - SNMP Communities
Community Name  MIB View  Write Access
-----
public          Manager   Unrestricted
  
```

Actions-> **Back** Add Edit Delete Help

Return to previous screen.
Use up/down arrow keys to change record selection, left/right arrow ke

2. Press **[A]** (for **Add**).

If you need information on the options in each field, press **[Enter]** to move the cursor to the Actions line, then select the Help option. When you are finished with Help, press **[E]** (for Edit) to return the cursor to the parameter fields.

3. Enter the name you want in the Community Name field, and use the Space bar to select the appropriate value in each of the other fields. (Use the **[Tab]** key to move from one field to the next.)
4. Press **[Enter]** , then **[S]** (for **Save**).

Listing community names and values (CLI)

This command lists the data for currently configured SNMP community names (along with trap receivers and the setting for authentication traps—see **SNMP notifications** on page 148).

Syntax:

```
show snmp-server [< community-string >]
```

Example:

Lists the data for all communities in a switch; that is, both the default "public" community name and another community named "blue-team."

Figure 30: Example: of the SNMP community listing with two communities

```
HP Switch# show snmp-server
SNMP Communities
-----
Community Name      MIB View  Write Access
-----
public              Manager   Unrestricted
blue-team           Operator  Restricted

Trap Receivers
Send Authentication Traps [No] : No

Address             Community  Events Sent in Trap
-----
```

To list the data for only one community, such as the "public" community, use the above command with the community name included. For Example:

```
switch# show snmp-server public
```

Configuring community names and values (CLI)

The `snmp-server` command enables you to add SNMP communities with either default or specific access attributes, and to delete specific communities.

Syntax:

```
[no] snmp-server community <community-name>
```

Configures a new community name.

- If you do not also specify `operator` or `manager`, the switch automatically assigns the community to the `operator` MIB view.
- If you do not specify `restricted` or `unrestricted`, the switch automatically assigns the community to `restricted` (read-only) access.

The `no` form uses only the `<community-name>` variable and deletes the named community from the switch.

<code>[operator manager]</code>	<p>Optionally assigns an access level.</p> <ul style="list-style-type: none"> At the <code>operator</code> level, the community can access all MIB objects except the CONFIG MIB. At the <code>manager</code> level, the community can access all MIB objects.
<code>[restricted unrestricted]</code>	<p>Optionally assigns MIB access type.</p> <ul style="list-style-type: none"> Assigning the <code>restricted</code> type allows the community to read MIB variables, but not to set them. Assigning the <code>unrestricted</code> type allows the community to read and set MIB variables.

Example:

To add the following communities:

Community	Access Level	Type of Access
red-team	manager (Access to all MIB objects.)	unrestricted (read/write)
blue-team	operator (Access to all MIB objects except the CONFIG MIB.)	restricted (read-only)

```
switch(config)# snmp-server community red-team
manager unrestricted
switch(config)# snmp-server community blue-team
operator restricted
```

To eliminate a previously configured community named "gold-team":

```
switch(config) # no snmp-server community gold-team
```

SNMP notifications

The switches:

- Fixed or "Well-Known" Traps: A switch automatically sends fixed traps (such as "coldStart", "warmStart", "linkDown", and "linkUp") to trap receivers using the public community name, which is the default. These traps can also be sent to non-public communities.
- SNMPv2c informs
- SNMP v3 notification process, including traps

This section describes how to configure a switch to send network security and link-change notifications to configured trap receivers.

Supported Notifications

By default, the following notifications are enabled on a switch:

- Manager password changes
- SNMP authentication failure
- Link-change traps: when the link on a port changes from up to down (linkDown) or down to up (linkUp)
- Port-security (web, MAC, or 802.1X) authentication failure

- Invalid password entered in a login attempt through a direct serial, Telnet, or SSH connection
- Inability to establish a connection with the RADIUS or TACACS+ authentication server
- DHCP snooping events
- ARP protection events

General steps for configuring SNMP notifications

Procedure

1. Determine the versions of SNMP notifications that you want to use in your network.

If you want to use SNMPv1 and SNMPv2c traps, you must also configure a trap receiver. See the following sections and follow the required configuration procedures:

- [SNMPv1 and SNMPv2c Traps](#) on page 149
- [Configuring an SNMP trap receiver \(CLI\)](#) on page 150
- [Enabling SNMPv2c informs \(CLI\)](#) on page 151

If you want to use SNMPv3 notifications (including traps), you must also configure an SNMPv3 management station. Follow the required configuration procedure in [Configuring SNMPv3 notifications \(CLI\)](#) on page 152.

2. To reconfigure any of the SNMP notifications that are enabled by default to be sent to a management station (trap receiver), see [Enabling Link-Change Traps \(CLI\)](#) on page 156.
3. (Optional) See the following sections to configure optional SNMP notification features and verify the current configuration:
 - [Configuring the source IP address for SNMP notifications \(CLI\)](#) on page 157
 - [Viewing SNMP notification configuration \(CLI\)](#) on page 159

SNMPv1 and SNMPv2c Traps

The switches support the following functionality from earlier SNMP versions (SNMPv1 and SNMPv2c):

- **Trap receivers:** A **trap receiver** is a management station to which the switch sends SNMP traps and (optionally) event log messages sent from the switch. From the CLI you can configure up to ten SNMP trap receivers to receive SNMP traps from the switch.
- **Fixed or "Well-Known" Traps:** A switch automatically sends fixed traps (such as "coldStart", "warmStart", "linkDown", and "linkUp") to trap receivers using the `public` community name. These traps cannot be redirected to other communities. If you change or delete the default `public` community name, these traps are not sent.
- **Thresholds:** A switch automatically sends all messages created when a system threshold is reached to the network management station that configured the threshold, regardless of the trap receiver configuration.

SNMP trap receivers

Use the `snmp-server host` command to configure a trap receiver that can receive SNMPv1 and SNMPv2c traps, and (optionally) Event Log messages. When you configure a trap receiver, you specify its community membership, management station IP address, and (optionally) the type of Event Log messages to be sent.

If you specify a community name that does not exist—that is, has not yet been configured on the switch—the switch still accepts the trap receiver assignment. However, no traps are sent to that trap receiver until the community to which it belongs has been configured on the switch.



To replace one community name with another for the same IP address, you must first enter the

```
no snmp-server host <community-name> {< ipv4-address | ipv6-address >}
```

command to delete the unwanted community name. Otherwise, if you add a new community name with an IP address that is already used with a different community name, two valid community name entries are created for the same management station.

If you do not specify the event level (`[none|all|not-info|critical|debug]`), the switch does not send Event Log messages as traps. However, "well-known" traps and threshold traps (if configured) are still sent.

Configuring an SNMP trap receiver (CLI)

Syntax:

```
snmp-server host {< ipv4-addr | ipv6-addr >} < community name >
```

Configures a destination network management station to receive SNMPv1/v2c traps and (optionally) Event Log messages sent as traps from the switch, using the specified community name and destination IPv4 or IPv6 address. You can specify up to ten trap receivers (network management stations). (The default community name is `public`.)

<pre>[[<none all not-info critical debug>]]</pre>	<p>(Optional) Configures the security level of the Event Log messages you want to send as traps to a trap receiver (see the following table).</p> <ul style="list-style-type: none"> • The type of Event Log message that you specify applies only to Event Log messages, not to threshold traps. • For each configured event level, the switch continues to send threshold traps to all network management stations that have the appropriate threshold level configured. • If you do not specify an event level, the switch uses the default value (<code>none</code>) and sends no Event Log messages as traps.
<pre>[<inform>]</pre>	<p>(Optional) Configures the switch to send SNMPv2 inform requests when certain events occur. For more information, see Enabling SNMPv2c informs (CLI).</p>

Table 18: Security levels for Event Log messages sent as traps

Security Level	Action
None (default)	Sends no Event Log messages.
All	Sends all Event Log messages.
Not-Info	Sends all Event Log messages that are not for information only.
Critical	Sends only Event Log messages for critical error conditions.
Debug	Sends only Event Log messages needed to troubleshoot network- and switch-level problems.

Example:

To configure a trap receiver in a community named "red-team" with an IP address of 10.28.227.130 to receive only "critical" event log messages, you can enter the following command:

```
switch(config)# snmp-server host 10.28.227.130 red-team critical
```

SNMPv2c informs

On a switch enabled for SNMPv2c, you can use the `snmp-server host inform` command (**Enabling SNMPv2c informs (CLI)** on page 151) to send inform requests when certain events occur. When an SNMP Manager receives an inform request, it can send an SNMP response back to the sending agent on the switch to let the agent know that the inform request reached its destination.

If the sending agent on the switch does not receive an SNMP response back from the SNMP Manager within the timeout period, the inform request may be resent, based on the retry count value.

When you enable SNMPv2c inform requests to be sent, you must specify the IP address and community name of the management station that will receive the inform notification.

Enabling SNMPv2c informs (CLI)

For information about enabling SNMPv2c informs, see **SNMPv2c informs** on page 151.

Syntax:

```
[no] snmp-server host {< ipv4-addr | ipv6-addr >} <community name> inform [retries < count >] [timeout < interval >]
```

Enables (or disables) the `inform` option for SNMPv2c on the switch and allows you to configure options for sending SNMP inform requests.

<code>retries</code>	Maximum number of times to resend an <code>inform</code> request if no SNMP response is received. (Default: 3)
<code>timeout</code>	Number of seconds to wait for an acknowledgement before resending the <code>inform</code> request. (Default: 15 seconds)



The `retries` and `timeout` values are not used to send trap requests.

To verify the configuration of SNMPv2c informs, enter the `show snmp-server` command, as shown in **Display of SNMPv2c inform configuration** on page 151 (note indication of inform Notify Type in bold below):

Display of SNMPv2c inform configuration

```
switch(config)# show snmp-server

SNMP Communities

Community Name      MIB View Write Access
-----
public              Manager Unrestricted

Trap Receivers

Link-Change Traps Enabled on Ports [All] : All
...
```

Address	Community	Events Sent	Notify Type	Retry	Timeout
15.28.333.456	guest	All	inform	3	15

Excluded MIBs

Snmpp Response Pdu Source-IP Information

Selection Policy : Default rfc1517

Trap Pdu Source-IP Information

Selection Policy : Configured IP

Ip Address : 10.10.10.10

Configuring SNMPv3 notifications (CLI)

The SNMPv3 notification process allows messages that are passed via SNMP between the switch and a network management station to be authenticated and encrypted.

Procedure

1. Enable SNMPv3 operation on the switch by entering the `snmpv3 enable` command (See "SNMP Version 3 Commands" on page N-7).

When SNMPv3 is enabled, the switch supports:

- Reception of SNMPv3 notification messages (traps and informs)
 - Configuration of initial users
 - (Optional) Restriction of non-SNMPv3 messages to "read only"
2. Configure SNMPv3 users by entering the `snmpv3 user` command (see [SNMPv3 users](#) on page 141). Each SNMPv3 user configuration is entered in the User Table.
 3. Assign SNMPv3 users to security groups according to their level of access privilege by entering the `snmpv3 group` command (see [Assigning users to groups \(CLI\)](#) on page 143).
 4. Define the name of an SNMPv3 notification configuration by entering the `snmpv3 notify` command.

Syntax:

```
[no] snmpv3 notify <notify_name> tagvalue <tag_name>
```

Associates the name of an SNMPv3 notification configuration with a tag name used (internally) in SNMPv3 commands. To delete a notification-to-tag mapping, enter `no snmpv3 notify notify_name`.

<code>notify <notify_name></code>	Specifies the name of an SNMPv3 notification configuration.
<code>tagvalue <tag_name></code>	Specifies the name of a tag value used in other SNMPv3 commands, such as <code>snmpv3 targetaddress params taglist tag_name</code> in Step 5.

5. Configure the target address of the SNMPv3 management station to which SNMPv3 informs and traps are sent by entering the `snmpv3 targetaddress` command.

Syntax:

```
[no] snmpv3 targetaddress {< ipv4-addr | ipv6-addr >} <name>
```

Configures the IPv4 or IPv6 address, name, and configuration filename of the SNMPv3 management station to which notification messages are sent.

<code>params <params_name></code>	Name of the SNMPv3 station's parameters file. The parameters filename configured with <code>params params_name</code> must match the <code>params params_name</code> value entered with the <code>snmpv3 params</code> command in Step 6.
<code>taglist <tag_name> [tag_name] ...</code>	Specifies the SNMPv3 notifications (identified by one or more <code>tag_name</code> values) to be sent to the IP address of the SNMPv3 management station. You can enter more than one <code>tag_name</code> value. Each <code>tag_name</code> value must be already associated with the name of an SNMPv3 notification configuration entered with the <code>snmpv3 notify</code> command in Step 4. Use a blank space to separate <code>tag_name</code> values. You can enter up to 103 characters in <code>tag_name</code> entries following the <code>taglist</code> keyword.
<code>[filter {<none debug all not-info critical>}]</code>	(Optional) Configures the type of messages sent to a management station. (Default: none.)
<code>[udp-port < port >]</code>	(Optional) Specifies the UDP port to use. (Default: 162.)
<code>[port-mask < mask >]</code>	(Optional) Specifies a range of UDP ports. (Default: 0.)
<code>[addr-mask < mask >]</code>	(Optional) Specifies a range of IP addresses as destinations for notification messages. (Default: 0.)
<code>[retries < value >]</code>	(Optional) Number of times a notification is retransmitted if no response is received. Range: 1-255. (Default: 3.)
<code>[timeout < value >]</code>	(Optional) Time (in millisecond increments) allowed to receive a response from the target before notification packets are retransmitted. Range: 0-2147483647. [Default: 1500 (15 seconds).]
<code>[max-msg-size < size >]</code>	(Optional) Maximum number of bytes supported in a notification message to the specified target. (Default: 1472)

6. Create a configuration record for the target address with the `snmpv3 params` command.

Syntax:

```
[no] snmpv3 params <params_name> user <user_name>
```

Applies the configuration parameters and IP address of an SNMPv3 management station (from the `params params_name` value configured with the `snmpv3 targetaddress` command in Step 5) to a specified SNMPv3 user (from the `user user_name` value configured with the `snmpv3 user` command in Step 2).

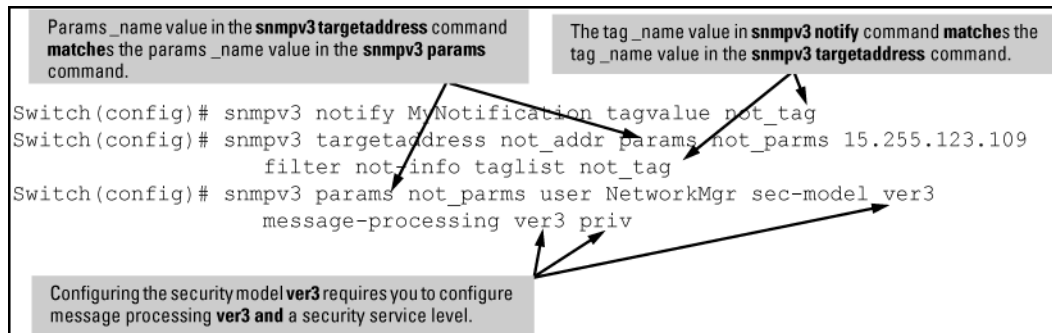
If you enter the `snmpv3 params user` command, you must also configure a security model (`sec-model`) and message processing algorithm (`msg-processing`).

<pre>{<sec-model [ver1 ver2c ver3>]}</pre>	<p>Configures the security model used for SNMPv3 notification messages sent to the management station configured with the <code>snmpv3 targetaddress</code> command in Step 5.</p> <p>If you configure the security model as <code>ver3</code>, you must also configure the message processing value as <code>ver3</code>.</p>
<pre>{msg-processing {<ver1 ver2c ver3>} [noaut auth priv]}</pre>	<p>Configures the algorithm used to process messages sent to the SNMPv3 target address.</p> <p>If you configure the message processing value as <code>ver3</code> and the security model as <code>ver3</code>, you must also configure a security services level (<code>noauth</code>, <code>auth</code>, or <code>priv</code>).</p>

Example:

An Example: of how to configure SNMPv3 notification is shown here:

Figure 31: Example: of an SNMPv3 notification configuration



Network security notifications

By default, a switch is enabled to send the SNMP notifications listed in **Supported Notifications** on page 148 when a network security event (For example, authentication failure) occurs. However, before security notifications can be sent, you must first configure one or more trap receivers or SNMPv3 management stations as described in:

- **Configuring an SNMP trap receiver (CLI)** on page 150
- **Configuring SNMPv3 notifications (CLI)** on page 152

You can manage the default configuration of the switch to disable and re-enable notifications to be sent for the following types of security events:

- ARP protection events
- Inability to establish a connection with the RADIUS or TACACS+ authentication server
- DHCP snooping events
- Dynamic IP Lockdown hardware resources consumed
- Link change notification
- Invalid password entered in a login attempt through a direct serial, Telnet, or SSH connection
- Manager password changes
- Port-security (web, MAC, or802.1X) authentication failure

- SNMP authentication failure
- Running configuration changes

Enabling or disabling notification/traps for network security failures and other security events (CLI)

For more information, see [Network security notifications](#) on page 154.

Syntax:

```
[no] snmp-server enable traps [snmp-auth | password-change-mgr | login-failure-mgr | port-security | auth-server-fail | dhcp-snooping | arp-protect | running-config-change]
```

Enables or disables sending one of the security notification types listed below to configured trap receivers. (Unless otherwise stated, all of the following notifications are enabled in the default configuration.)

The notification sends a trap:

arp-protect	If ARP packets are received with an invalid source or destination MAC address, an invalid IP address, or an invalid IP-to-MAC binding.
auth-server-fail	If the connection with a RADIUS or TACACS+ authentication server fails.
dhcp-snooping	If DHCP packets are received from an untrusted source or if DHCP packets contain an invalid IP-to-MAC binding.
dhcpv6-snooping	Set the traps for DHCPv6 snooping.
dyn-ip-lockdown	If the switch is out of hardware resources needed to program a dynamic IP lockdown rule
dyn-ipv6-lockdown	Enable traps for Dynamic IPv6 lockdown..
link-change <port-list>	When the link state on a port changes from up to down, or the reverse.
login-failure-mgr	For a failed login with a manager password.
password-change-mgr	When a manager password is reset.
mac-notify	Globally enables the generation of SNMP trap notifications upon MAC address table changes.
nd-snooping	Set the trap for nd snooping
port-security	For a failed authentication attempt through a web, MAC, or 801.X authentication session.
running-config-change	When changes to the running configuration file are made.

Table Continued

snmp-authentication [extended standard]	For a failed authentication attempt via SNMP.(Default: extended.)
Startup-config-change	Sends a trap when changes to the startup configuration file are made. See “Enabling SNMP Traps on Startup Configuration Changes” on page 6–34. (Default: Disabled)

To determine the specific cause of a security event, check the Event Log in the console interface to see why a trap was sent. For more information, see "Using the Event Log for Troubleshooting Switch Problems".

Viewing the current configuration for network security notifications (CLI)

Enter the `show snmp-server traps` command, as shown in [Display of configured network security notifications](#) on page 156. Note that command output is a subset of the information displayed with the `show snmp-server` command in [Display of SNMP notification configuration](#).

Display of configured network security notifications

```
switch(config)# show snmp-server traps

Trap Receivers

Link-Change Traps Enabled on Ports [All] : A1-A24

Traps Category                Current Status
-----
SNMP Authentication           : Extended
Password change                : Enabled
Login failures                 : Enabled
Port-Security                  : Enabled
Authorization Server Contact   : Enabled
DHCP Snooping                  : Enabled
Dynamic ARP Protection         : Enabled
Dynamic IP Lockdown            : Enabled

Address            Community  Events Sent  Notify Type  Retry  Timeout
-----
15.255.5.225      public     All          trap          3     15
2001:0db8:0000:0001
:0000:0000:0000:0121 user_1     All          trap          3     15

Excluded MIBs
```

Enabling Link-Change Traps (CLI)

By default, a switch is enabled to send a trap when the link state on a port changes from up to down (linkDown) or down to up (linkUp). To reconfigure the switch to send link-change traps to configured trap receivers, enter the `snmp-server enable traps link-change` command.

Syntax:

```
[no] snmp-server enable traps link-change <port-list> [all]
```

Enables or disables the switch to send a link-change trap to configured trap receivers when the link state on a port goes from up to down or down to up.

Enter `all` to enable or disable link-change traps on all ports on the switch.

Readable interface names in traps

The SNMP trap notification messages for linkup and linkdown events on an interface includes IfDesc and IfAlias var-bind information.

Source IP address for SNMP notifications

The switch uses an interface IP address as the source IP address in IP headers when sending SNMP notifications (traps and informs) or responses to SNMP requests.

For multi-netted interfaces, the source IP address is the IP address of the outbound interface of the SNMP reply, which may differ from the destination IP address in the IP header of the received request. For security reasons, it may be desirable to send an SNMP reply with the IP address of the destination interface (or a specified IP address) on which the corresponding SNMP request was received.

To configure the switch to use the source IP address on which an SNMP request was received in SNMP notification/traps and replies, enter the `snmp-server response-source` and `snmp-server trap-source` commands (**Configuring the source IP address for SNMP notifications (CLI)**).

Configuring the source IP address for SNMP notifications (CLI)

For more information, see [Source IP address for SNMP notifications](#) on page 157.

Syntax:

```
[no] snmp-server response-source [dst-ip-of-request | [ipv4-addr | ipv6-addr] | loopback <0-7>]
```

Specifies the source IP address of the SNMP response PDU. The default SNMP response PDU uses the IP address of the active interface from which the SNMP response was sent as the source IP address.

The `no` form of the command resets the switch to the default behavior (compliant with rfc-1517).

(Default: Interface IP address)

<code>dst-ip-of-request</code>	Destination IP address of the SNMP request PDU that is used as the source IP address in an SNMP response PDU.
<code>[ipv4-addr ipv6-addr]</code>	User-defined interface IP address that is used as the source IP address in an SNMP response PDU. Both IPv4 and IPv6 addresses are supported.
<code>loopback <0-7></code>	IP address configured for the specified loopback interface that is used as the source IP address in an SNMP response PDU. If multiple loopback IP addresses are configured, the lowest alphanumeric address is used.

To use the IP address of the destination interface on which an SNMP request was received as the source IP address in the IP header of SNMP traps and replies, enter the following command:

```
switch(config)# snmp-server response-source dst-ip-of-request
```

Syntax:

```
[no] snmp-server trap-source [ipv4-addr | loopback <0-7>]
```

Specifies the source IP address to be used for a trap PDU. To configure the switch to use a specified source IP address in generated trap PDUs, enter the `snmp-server trap-source` command.

The `no` form of the command resets the switch to the default behavior (compliant with rfc-1517).

(Default: Use the interface IP address in generated trap PDUs)

<code>ipv4-addr</code>	User-defined interface IPv4 address that is used as the source IP address in generated traps. IPv6 addresses are not supported.
<code>loopback <0-7></code>	P address configured for the specified loopback interface that is used as the source IP address in a generated trap PDU. If multiple loopback IP addresses are configured, the lowest alphanumeric address is used.



When you use the `snmp-server response-source` and `snmp-server trap-source` commands, note the following behavior:

- The `snmp-server response-source` and `snmp-server trap-source` commands configure the source IP address for IPv4 interfaces only.
- You must manually configure the `snmp-server response-source` value if you wish to change the default user-defined interface IP address that is used as the source IP address in SNMP traps (RFC 1517).
- The values configured with the `snmp-server response-source` and `snmp-server trap-source` commands are applied globally to all interfaces that are sending SNMP responses or SNMP trap PDUs.
- Only the source IP address field in the IP header of the SNMP response PDU can be changed.
- Only the source IP address field in the IP header and the SNMPv1 Agent Address field of the SNMP trap PDU can be changed.

Verifying the configuration of the interface IP address used as the source IP address in IP headers for SNMP replies and traps sent from the switch (CLI)

Enter the `show snmp-server` command to display the SNMP policy configuration, as shown in [Display of source IP address configuration](#) on page 158.

Display of source IP address configuration

```
switch(config)# show snmp-server

SNMP Communities

Community Name      MIB View Write Access
-----
public              Manager  Unrestricted

Trap Receivers
Link-Change Traps Enabled on Ports [All] : All

...

Excluded MIBs
Snmp Response Pdu Source-IP Information
Selection Policy : dstIpOfRequest 1

Trap Pdu Source-IP Information
Selection Policy : Configured IP
```

¹ `dstIpOfRequest`: The destination IP address of the interface on which an SNMP request is received is used as the source IP address in SNMP replies.

Viewing SNMP notification configuration (CLI)

Syntax:

```
show snmp-server
```

Displays the currently configured notification settings for versions SNMPv1 and SNMPv2c traps, including SNMP communities, trap receivers, link-change traps, and network security notifications.

Example:

In the following Example:, the `show snmp-server` command output shows that the switch has been configured to send SNMP traps and notifications to management stations that belong to the "public," "red-team," and "blue-team" communities.

Figure 32: Display of SNMP notification configuration

```
HP Switch(config)# show snmp-server

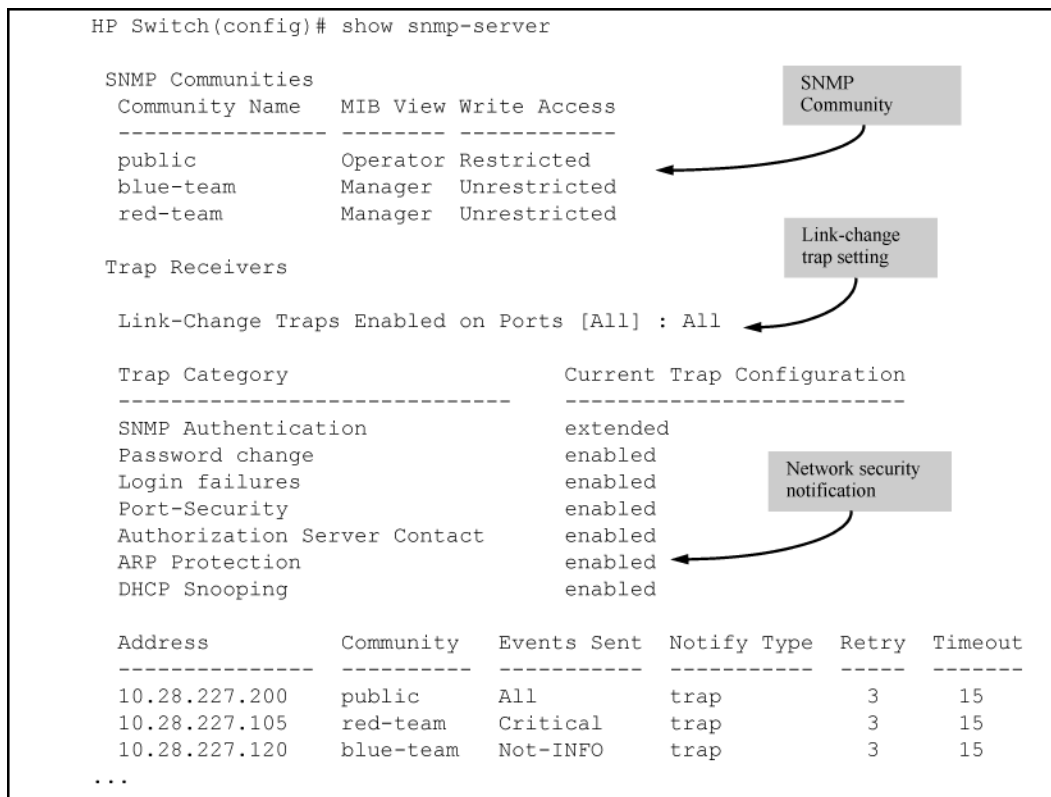
SNMP Communities
Community Name  MIB View Write Access
-----
public          Operator Restricted
blue-team       Manager Unrestricted
red-team        Manager Unrestricted

Trap Receivers

Link-Change Traps Enabled on Ports [All] : All

Trap Category                Current Trap Configuration
-----
SNMP Authentication          extended
Password change              enabled
Login failures               enabled
Port-Security                enabled
Authorization Server Contact  enabled
ARP Protection               enabled
DHCP Snooping               enabled

Address      Community  Events Sent  Notify Type  Retry  Timeout
-----
10.28.227.200 public     All          trap         3      15
10.28.227.105 red-team   Critical     trap         3      15
10.28.227.120 blue-team  Not-INFO    trap         3      15
...
```



Advanced management: RMON

The switch supports RMON (remote monitoring) on all connected network segments. This allows for troubleshooting and optimizing your network.

The following RMON groups are supported:

- Ethernet Statistics (except the numbers of packets of different frame sizes)
- Alarm
- History (of the supported Ethernet statistics)
- Event

The RMON agent automatically runs in the switch. Use the RMON management station on your network to enable or disable specific RMON traps and events. Note that you can access the Ethernet statistics, Alarm, and Event groups from the HPE Switch Manager network management software. For more information on PCM+, see the HPE Networking web site at <http://www.hpe.com/networking>.

From the Products menu, select Network Management. Then click on PCM+ Network Management under the Network Management bar.

CLI-configured sFlow with multiple instances

sFlow can also be configured via the CLI for up to three distinct sFlow instances: once enabled, an sFlow receiver/destination can be independently configured for full flow-sampling and counter-polling. CLI-configured sFlow instances may be saved to the startup configuration to persist across a switch reboot.

Configuring sFlow (CLI)

The following sFlow commands allow you to configure sFlow instances via the CLI. For more information, see **Advanced management: RMON** on page 159.

Syntax:

```
[no] sflow <receiver-instance> destination <ip-address> [< udp-port-num >]
```

Enables an sFlow receiver/destination. The receiver-instance number must be a 1, 2, or 3.

By default, the udp destination port number is 6343.

To disable an sFlow receiver/destination, enter `no sflow receiver-instance` .

Syntax:

```
sflow <receiver-instance> sampling <port-list> <sampling rate>
```

Once an sFlow receiver/destination has been enabled, this command enables flow sampling for that instance. The receiver-instance number is 1, 2, or 3, and the sampling rate is the allowable non-zero skipcount for the specified port or ports.

To disable flow-sampling for the specified port-list, repeat the above command with a sampling rate of 0.

Syntax:

```
sflow <receiver-instance> polling <port-list> <polling interval>
```

Once an sFlow receiver/destination has been enabled, this command enables counter polling for that instance. The receiver-instance number is 1, 2, or 3, and the polling interval may be set to an allowable non-zero value to enable polling on the specified port or ports.

To disable counter-polling for the specified port-list, repeat the above command with a polling interval of 0.



Under the multiple instance implementation, sFlow can be configured via the CLI or via SNMP. However, CLI-owned sFlow configurations cannot be modified via SNMP, whereas SNMP-owned instances can be disabled via the CLI using the `no sflow <receiver-instance>` command.

Viewing sFlow Configuration and Status (CLI)

The following sFlow commands allow you to display sFlow configuration and status via the CLI. **Viewing sFlow destination information** on page 161 is an Example: of `sflow agent` information.

Syntax:

```
show sflow agent
```


Displays sFlow agent information. The agent address is normally the IP address of the first VLAN configured.

The `show sflow agent` command displays read-only switch agent information. The version information shows the sFlow version, MIB support, and software versions; the agent address is typically the IP address of the first VLAN configured on the switch.

Viewing sflow agent information

```
switch# show sflow agent

Version          1.3;HP;XX.11.40
Agent Address    10.0.10.228
```

Syntax:

```
show sflow <receiver instance> destination
```

Displays information about the management station to which the sFlow sampling-polling data is sent.

The `show sflow instance destination` command includes information about the management-station's destination address, receiver port, and owner, as shown in [Viewing sFlow destination information](#) on page 161.

Viewing sFlow destination information

```
switch# show sflow 2 destination

Destination Instance      2
sflow                     Enabled
Datagrams Sent           221
Destination Address       10.0.10.41
Receiver Port             6343
Owner                     Administrator, CLI-owned, Instance 2
Timeout (seconds)        99995530
Max Datagram Size        1400
Datagram Version Support  5
```

Note the following details:

- **Destination Address** remains blank unless it has been configured.
- **Datagrams Sent** shows the number of datagrams sent by the switch agent to the management station since the switch agent was last enabled.
- **Timeout** displays the number of seconds remaining before the switch agent will automatically disable sFlow (this is set by the management station and decrements with time).
- **Max Datagram Size** shows the currently set value (typically a default value, but this can also be set by the management station).

Syntax:

```
show sflow <receiver instance> sampling-polling <port-list/range>
```

Displays status information about sFlow sampling and polling.

The `show sflow instance sampling-polling [port-list]` command displays information about sFlow sampling and polling on the switch, as shown in **Figure 33: Example: of viewing sFlow sampling and polling information** on page 162. You can specify a list or range of ports for which to view sampling information.

Figure 33: Example: of viewing sFlow sampling and polling information

```
HP Switch# show sflow 2 sampling-polling A1-A4
```

Number denotes the sampling/polling instance to which the receiver is coupled.

Port	Sampling			Dropped Samples	Polling	
	Enabled	Rate	Header		Enabled	Interval
A1	Yes (2)	40	128	1234567890	---	---
A2	---	---	---	0	Yes (1)	60
A3	No (1)	0	100	898703	No	30
A4	Yes (3)	50	128	0	No (3)	0



The sampling and polling instances (noted in parentheses) coupled to a specific receiver instance are assigned dynamically, and so the instance numbers may not always match. The key thing to note is whether sampling or polling is enabled on a port, and the sampling rates or polling intervals for the receiver instance configured on each port.

Configuring UDLD Verify before forwarding

When an UDLD enabled port transitions to link-up, the port will begin with a UDLD blocking state. UDLD will probe via protocol packet exchange to determine the bidirectional state of the link. Until UDLD has completed the probe, all data traffic will be blocked. If the link is found to be bidirectional, UDLD will unblock the port for data traffic to pass. Once UDLD unblocks the port, other protocols will see the port as up and data traffic can be safely forwarded.

The default mode of a switch is “forward first then verify”. Enabling UDLD link-up will default to “forward first then verify”. To change the mode to “verify then forward”, you need to configure using the commands found in section 6.72.



Link-UP data traffic will resumed after probing the link partner completes. All other protocols running will see the port as down.

UDLD time delay

UDLD protocol informs the link partner simultaneously as it detects a state change from unidirectional to bidirectional traffic. Additional packet exchanges will be carried out by UDLD in addition to the existing UDLD exchanges whenever state changes from unidirectional to bidirectional.

Table 19: Peer state transition timings

Interval Time	Interval 1	Interval 1 + delta	Interval 2	Interval 3
	5 sec	5+(<5) sec*	10 sec	15 sec
With triggered updates	State = blockedPeer State = blocked	Inform PeerState = unblockedPeer State = unblocked	Regular UDLD TX	Regular UDLD TX

Table Continued

Interval Time	Interval 1	Interval 1 + delta	Interval 2	Interval 3
Without triggered updates	State = blockedPeer State = blocked	State = unblockedPeer State = blocked	Inform PeerState = unblockedPeer State = unblocked	Regular UDLD TX
*delta is the time when the unblock event occurs on local side				

Restrictions

- There is no support available when configuring this mode from the web and menu interface.
- There are no new packet types are introduced with UDLD.
- There are no new UDLD timers being introduced.

UDLD configuration commands

Syntax:

```
HP Switch(config)# link-keepalive mode [verify-then-forward | forward-then-verify]
```

This command configures the link-keepalive mode.

Link-keepalive provides two modes of operation; `verify-then-forward` and `forward-then-verify`.

When using the `verify-then-forward` mode, the port is in a blocking state until the link configured for UDLD establishes bidirectional communication. When using the `forward-then-verify` mode, the port forwards the data then verifies the status of the link-in state.

When a unidirectional state is detected, the port is moved to a blocked state.

When a bidirectional state is detected, the data is forwarded without interruption.

Syntax:

```
HP Switch(config)# link-keepalive mode verify-then-forward
```

Keeps the port in a logically blocked state until the link configured for UDLD has been successfully established in bi-directional communication.

Syntax:

```
HP Switch(config)# link-keepalive mode forward-then-verify
```

Forwards the data then verifies the status of the link. If a unidirectional state is detected, the port is then moved to a blocked state.

Syntax:

```
HP Switch(config)# link-keepalive interval <deciseconds>
```

Configure the interval for `link-keepalive`. The `link-keepalive` interval is the time between sending two UDLD packets. The time interval is entered in deciseconds (1/10 sec). The default keepalive interval is 50 deciseconds.

Example:

A value of 10 is 1 sec., 11 is 1.1 sec.

Syntax:

```
HP Switch(config)# link-keepalive retries <number>
```

Maximum number of sending attempts for UDLD packets before declaring the link as faulty.

Default keepalive attempt is 4.

Show commands

Syntax:

```
switch(config)# show link-keepalive
```

Sample output:

```
Total link-keepalive enabled ports: 8
Keepalive Retries : 4
Keepalive Interval: 5 sec
Keepalive Mode : verify-then-forward
Physical Keepalive Adjacent UDLD
```

Port	Enabled	Status	Status	Switch	VLAN
1	Yes	down	off-line	000000-000000	untagged
2	Yes	down	off-line	000000-000000	untagged
3	Yes	down	off-line	000000-000000	untagged
4	Yes	down	off-line	000000-000000	untagged
5	Yes	down	off-line	000000-000000	untagged
6	Yes	down	off-line	000000-000000	untagged
7	Yes	down	off-line	000000-000000	untagged
8	Yes	down	off-line	000000-000000	untagged

RMON generated when user changes UDLD mode

RMON events are generated when UDLD is configured. The first time you configure the mode, the UDLD states will be re-initialized. An event log entry is initiated to include the reason for the initial UDLD blocking state during link up.

Example:

```
UDLD mode [verify-then-forward | forward-then-verify] is configured
```

Severity: - Info.

LLDP

To standardize device discovery on all HPE switches, LLDP is implemented while offering limited read-only support for CDP, as documented in this manual. For the latest information on your switch model, consult the Release Notes (available on the HPE Networking website). If LLDP has not yet been implemented (or if you are running an older version of software), consult a previous version of the *Management and Configuration Guide* for device discovery details.

LLDP (Link Layer Discovery Protocol): provides a standards-based method for enabling the switches covered in this guide to advertise themselves to adjacent devices and to learn about adjacent LLDP devices.

LLDP-MED (LLDP Media Endpoint Discovery): Provides an extension to LLDP and is designed to support VoIP deployments.



LLDP-MED is an extension for LLDP, and the switch requires that LLDP be enabled as a prerequisite to LLDP-MED operation.

An SNMP utility can progressively discover LLDP devices in a network by:

Procedure

1. Reading a given device's Neighbors table (in the Management Information Base, or MIB) to learn about other, neighboring LLDP devices.
2. Using the information learned in step 1 to find and read the neighbor devices' Neighbors tables to learn about additional devices, and so on.

Also, by using `show` commands to access the switch's neighbor database for information collected by an individual switch, system administrators can learn about other devices connected to the switch, including device type (capability) and some configuration information. In VoIP deployments using LLDP-MED on the switches, additional support unique to VoIP applications is also available. See [LLDP-MED \(media-endpoint-discovery\)](#) on page 179.

General LLDP operation

An LLDP packet contains data about the transmitting switch and port. The switch advertises itself to adjacent (neighbor) devices by transmitting LLDP data packets out all ports on which outbound LLDP is enabled and by reading LLDP advertisements from neighbor devices on ports that are inbound LLDP-enabled. (LLDP is a one-way protocol and does not include any acknowledgement mechanism.) An LLDP-enabled port receiving LLDP packets inbound from neighbor devices stores the packet data in a Neighbor database (MIB).

LLDP-MED

This capability is an extension to LLDP and is available on the switches. See [LLDP-MED \(media-endpoint-discovery\)](#) on page 179.

Packet boundaries in a network topology

- Where multiple LLDP devices are directly connected, an outbound LLDP packet travels only to the next LLDP device. An LLDP-capable device does not forward LLDP packets to any other devices, regardless of whether they are LLDP-enabled.
- An intervening hub or repeater forwards the LLDP packets it receives in the same manner as any other multicast packets it receives. Thus, two LLDP switches joined by a hub or repeater handle LLDP traffic in the same way that they would if directly connected.
- Any intervening 802.1D device or Layer-3 device that is either LLDP-unaware or has disabled LLDP operation drops the packet.

LLDP operation configuration options

In the default configuration, LLDP is enabled and in both transmit and receive mode on all active ports. The LLDP configuration includes global settings, which apply to all active ports on the switch, and per-port settings, which affect only the operation of the specified ports.

The commands in the LLDP sections affect both LLDP and LLDP-MED operation. For information on operation and configuration unique to LLDP-MED, see [LLDP-MED \(media-endpoint-discovery\)](#) on page 179.

Enable or disable LLDP on the switch

In the default configuration, LLDP is globally enabled on the switch. To prevent transmission or receipt of LLDP traffic, you can disable LLDP operation.

Enable or disable LLDP-MED

In the default configuration for the switches, LLDP-MED is enabled by default. (Requires that LLDP is also enabled.) For more information, see [LLDP-MED \(media-endpoint-discovery\)](#) on page 179.

Change the frequency of LLDP packet transmission to neighbor devices

On a global basis, you can increase or decrease the frequency of outbound LLDP advertisements.

Change the Time-To-Live for LLDP packets sent to neighbors

On a global basis, you can increase or decrease the time that the information in an LLDP packet outbound from the switch will be maintained in a neighbor LLDP device.

Transmit and receive mode

With LLDP enabled, the switch periodically transmits an LLDP advertisement (packet) out each active port enabled for outbound LLDP transmissions and receives LLDP advertisements on each active port enabled to receive LLDP traffic ([Configuring per-port transmit and receive modes \(CLI\)](#) on page 175). Per-port configuration options include four modes:

- Transmit and receive (`tx_rx`): This is the default setting on all ports. It enables a given port to both transmit and receive LLDP packets and to store the data from received (inbound) LLDP packets in the switch's MIB.
- Transmit only (`txonly`): This setting enables a port to transmit LLDP packets that can be read by LLDP neighbors. However, the port drops inbound LLDP packets from LLDP neighbors without reading them. This prevents the switch from learning about LLDP neighbors on that port.
- Receive only (`rxonly`): This setting enables a port to receive and read LLDP packets from LLDP neighbors and to store the packet data in the switch's MIB. However, the port does not transmit outbound LLDP packets. This prevents LLDP neighbors from learning about the switch through that port.
- Disable (`disable`): This setting disables LLDP packet transmissions and reception on a port. In this state, the switch does not use the port for either learning about LLDP neighbors or informing LLDP neighbors of its presence.

SNMP notification

You can enable the switch to send a notification to any configured SNMP trap receiver(s) when the switch detects a remote LLDP data change on an LLDP-enabled port ([Configuring SNMP notification support](#) on page 174).

Per-port (outbound) data options

The following table lists the information the switch can include in the per-port, outbound LLDP packets it generates. In the default configuration, all outbound LLDP packets include this information in the TLVs transmitted to neighbor devices. However, you can configure LLDP advertisements on a per-port basis to omit some of this information ([Configuring a remote management address for outbound LLDP advertisements \(CLI\)](#) on page 175).

Table 20: Data available for basic LLDP advertisements

Data type	Configuration options	Default	Description
Time-to-Live	1	120 Seconds	The length of time an LLDP neighbor retains the advertised data before discarding it.
Chassis Type ^{2,3}	N/A	Always Enabled	Indicates the type of identifier used for Chassis ID.
Chassis ID ³	N/A	Always Enabled	Uses base MAC address of the switch.
Port Type ^{4,3}	N/A	Always Enabled	Uses "Local," meaning assigned locally by LLDP.
Port Id ³	N/A	Always Enabled	Uses port number of the physical port. This is an internal number reflecting the reserved slot/port position in the chassis. For more information on this numbering scheme, see the appendix "MAC Address Management".
Remote Management Address			
Type ^{3,5}	N/A	Always Enabled	Shows the network address type.
Address ⁵	Default or Configured	Uses a default address selection method unless an optional address is configured. See Remote management address on page 168.	
System Name ³	Enable/Disable	Enabled	Uses the switch's assigned name.
System Description ³	Enable/Disable	Enabled	Includes switch model name and running software version, and ROM version.
Port Description ³	Enable/Disable	Enabled	Uses the physical port identifier.

Table Continued

Data type	Configuration options	Default	Description
System capabilities supported ^{3,6}	Enable/Disable	Enabled	Identifies the switch's primary capabilities (bridge, router).
System capabilities enabled ^{3,6}	Enable/Disable	Enabled	Identifies the primary switch functions that are enabled, such as routing.

¹ The Packet Time-to-Live value is included in LLDP data packets.

² Subelement of the Chassis ID TLV.

³ Populated with data captured internally by the switch. For more on these data types, refer to the IEEE P802.1AB Standard.

⁴ Subelement of the Port ID TLV.

⁵ Subelement of the Remote-Management-Address TLV.

⁶ Subelement of the System Capability TLV.

Remote management address

The switch always includes an IP address in its LLDP advertisements. This can be either an address selected by a default process or an address configured for inclusion in advertisements. See [IP address advertisements](#) on page 169.

Debug logging

You can enable LLDP debug logging to a configured debug destination (Syslog server, a terminal device, or both) by executing the `debug lldp` command. (For more information on Debug and Syslog, see the "Troubleshooting" appendix in this guide.) Note that the switch's Event Log does not record usual LLDP update messages.

Options for reading LLDP information collected by the switch

You can extract LLDP information from the switch to identify adjacent LLDP devices. Options include:

- Using the switch's `show lldp info` command options to display data collected on adjacent LLDP devices—as well as the local data the switch is transmitting to adjacent LLDP devices ([Displaying the global LLDP, port admin, and SNMP notification status \(CLI\)](#) on page 169).
- Using an SNMP application that is designed to query the Neighbors MIB for LLDP data to use in device discovery and topology mapping.
- Using the `walkmib` command to display a listing of the LLDP MIB objects

LLDP and LLDP-MED standards compatibility

The operation covered by this section is compatible with these standards:

- IEEE P802.1AB
- RFC 2922 (PTOPO, or Physical Topology MIB)
- RFC 2737 (Entity MIB)
- RFC 2863 (Interfaces MIB)
- ANSI/TIA-1057/D6 (LLDP-MED; refer to [LLDP-MED \(media-endpoint-discovery\)](#) on page 179.)

LLDP operating rules

For additional information specific to LLDP-MED operation, see [LLDP-MED \(media-endpoint-discovery\)](#) on page 179.

Port trunking

LLDP manages trunked ports individually. That is, trunked ports are configured individually for LLDP operation, in the same manner as non-trunked ports. Also, LLDP sends separate advertisements on each port in a trunk, and not on a per-trunk basis. Similarly, LLDP data received through trunked ports is stored individually, per-port.

IP address advertisements

In the default operation, if a port belongs to only one static VLAN, the port advertises the lowest-order IP address configured on that VLAN. If a port belongs to multiple VLANs, the port advertises the lowest-order IP address configured on the VLAN with the lowest VID. If the qualifying VLAN does not have an IP address, the port advertises 127.0.0.1 as its IP address. For example, if the port is a member of the default VLAN (VID=1), and there is an IP address configured for the default VLAN, the port advertises this IP address. In the default operation, the IP address that LLDP uses can be an address acquired by DHCP or Bootp.

You can override the default operation by configuring the port to advertise any IP address that is manually configured on the switch, even if the port does not belong to the VLAN configured with the selected IP address (**Configuring a remote management address for outbound LLDP advertisements (CLI)** on page 175). (Note that LLDP cannot be configured through the CLI to advertise an addresses acquired through DHCP or Bootp. However, as mentioned above, in the default LLDP configuration, if the lowest-order IP address on the VLAN with the lowest VID for a given port is a DHCP or Bootp address, the switch includes this address in its LLDP advertisements unless another address is configured for advertisements on that port.) Also, although LLDP allows configuring multiple remote management addresses on a port, only the lowest-order address configured on the port will be included in outbound advertisements. Attempting to use the CLI to configure LLDP with an IP address that is either not configured on a VLAN or has been acquired by DHCP or Bootp results in the following error message.

```
xxx.xxx.xxx.xxx: This IP address is not configured or is a DHCP address.
```

Spanning-tree blocking

Spanning tree does not prevent LLDP packet transmission or receipt on STP-blocked links.

802.1X blocking

Ports blocked by 802.1X operation do not allow transmission or receipt of LLDP packets.

Configuring LLDP operation

Displaying the global LLDP, port admin, and SNMP notification status (CLI)

In the default configuration, LLDP is enabled and in both transmit and receive mode on all active ports. The LLDP configuration includes global settings that apply to all active ports on the switch, and per-port settings that affect only the operation of the specified ports.

The commands in this section affect both LLDP and LLDP-MED operation. for information on operation and configuration unique to LLDP-MED, refer to “LLDP-MED (Media-Endpoint-Discovery)”.

Syntax:

```
show lldp config
```

Displays the LLDP global configuration, LLDP port status, and SNMP notification status. For information on port admin status, see **Configuring per-port transmit and receive modes (CLI)** on page 175.

`show lldp config` produces the following display when the switch is in the default LLDP configuration:

Viewing the general LLDP configuration

```
switch(config)# show lldp config
```

LLDP Global Configuration

```
LLDP Enabled [Yes] : Yes
LLDP Transmit Interval [30] : 30
LLDP Hold time Multiplier [4] : 4
LLDP Delay Interval [2] : 2
LLDP Reinit Interval [2] : 2
LLDP Notification Interval [5] : 5
LLDP Fast Start Count [5] : 5
```

LLDP Port Configuration

Port	AdminStatus	NotificationEnabled	Med Topology Trap Enabled
A1	Tx_Rx	False	False
A2	Tx_Rx	False	False
A3	Tx_Rx	False	False
A4	Tx_Rx	False	False
A5	Tx_Rx	False	False
A6	Tx_Rx	False	False
A7	Tx_Rx	False	False
A8	Tx_Rx	False	False



The values displayed in the LLDP column correspond to the `lldp refresh-interval` command

Viewing port configuration details (CLI)

Syntax:

```
show lldp config <port-list>
```

Displays the LLDP port-specific configuration for all ports in *<port-list>*, including which optional TLVs and any non-default IP address that are included in the port's outbound advertisements.

For information on the notification setting, see [Configuring SNMP notification support](#) on page 174. For information on the other configurable settings displayed by this command, see [Configuring per-port transmit and receive modes \(CLI\)](#) on page 175.

Figure 34: Per-port configuration display

```
HP Switch(config)# show lldp config 1

LLDP Port Configuration Detail

Port : 1
AdminStatus [Tx_Rx] : Tx_Rx
NotificationEnabled [False] : False
Med Topology Trap Enabled [False] : False

TLVS Advertised:
* port_descr
* system_name
* system_descr
* system_cap

[*capabilities]
|* network_policy|
|* location_id  |
|* poe          |
|* _ _ _ _ _ _ |
[*macphy_config]

IpAddress Advertised:
```

These fields appear when medtlvenable is enabled on the switch, which is the default setting.

This field appears when dot3tlvenable is enabled on the switch, which is the default setting.

The blank IpAddress field indicates that the default IP address will be advertised from this port.

Configuring Global LLDP Packet Controls

The commands in this section configure the aspects of LLDP operation that apply the same to all ports in the switch.

LLDP operation on the switch

Enabling LLDP operation (the default) causes the switch to:

- Use active, LLDP-enabled ports to transmit LLDP packets describing itself to neighbor devices.
- Add entries to its neighbors table based on data read from incoming LLDP advertisements.

Enabling or disabling LLDP operation on the switch (CLI)

For more information, see [LLDP operation on the switch](#) on page 171.

Syntax:

```
[no] lldp run
```

Enables or disables LLDP operation on the switch.

The `no` form of the command, regardless of individual LLDP port configurations, prevents the switch from transmitting outbound LLDP advertisements and causes the switch to drop all LLDP advertisements received from other devices.

The switch preserves the current LLDP configuration when LLDP is disabled. After LLDP is disabled, the information in the LLDP neighbors database remains until it times-out.

(Default: Enabled)

Disabling LLDP

```
switch(config)# no lldp run
```

Changing the packet transmission interval (CLI)

This interval controls how often active ports retransmit advertisements to their neighbors.

Syntax:

```
lldp refresh-interval <5-32768>
```

Changes the interval between consecutive transmissions of LLDP advertisements on any given port.

(Default: 30 seconds)



The `refresh-interval` must be greater than or equal to (4 x `delay-interval`). (The default `delay-interval` is 2). For example, with the default `delay-interval`, the lowest `refresh-interval` you can use is 8 seconds (4 x 2=8). Thus, if you want a `refresh-interval` of 5 seconds, you must first change the `delay-interval` to 1 (that is, 4 x 1 5). If you want to change the `delay-interval`, use the `setmib` command.

Time-to-Live for transmitted advertisements

The Time-to-Live value (in seconds) for all LLDP advertisements transmitted from a switch is controlled by the switch that generates the advertisement and determines how long an LLDP neighbor retains the advertised data before discarding it. The Time-to-Live value is the result of multiplying the `refresh-interval` by the `holdtime-multiplier`.

Changing the time-to-live for transmitted advertisements (CLI)

For more information, see [Time-to-Live for transmitted advertisements](#) on page 172.

Syntax:

```
lldp holdtime-multiplier <2-10>
```

Changes the multiplier an LLDP switch uses to calculate the Time-to-Live for the LLDP advertisements it generates and transmits to LLDP neighbors. When the Time-to-Live for a given advertisement expires, the advertised data is deleted from the neighbor switch's MIB.

(Default: 4; Range 2–10)

Example:

If the `refresh-interval` on the switch is 15 seconds and the `holdtime-multiplier` is at the default, the Time-to-Live for advertisements transmitted from the switch is 60 seconds (4 x 15).

To reduce the Time-to-Live, you could lower the `holdtime-interval` to 2, which would result in a Time-to-Live of 30 seconds.

```
switch(config)# lldp holdtime-multiplier 2
```

Delay interval between advertisements generated by value or status changes to the LLDP MIB

The switch uses a **delay-interval** setting to delay transmitting successive advertisements resulting from these LLDP MIB changes. If a switch is subject to frequent changes to its LLDP MIB, lengthening this interval can

reduce the frequency of successive advertisements. You can change the delay-interval by using either an SNMP network management application or the CLI `setmib` command.

Changing the delay interval between advertisements generated by value or status changes to the LLDP MIB (CLI)

Syntax:

```
setmib lldpTxDelay.0 -i <1-8192>
```

Uses `setmib` to change the minimum time (delay-interval) any LLDP port will delay advertising successive LLDP advertisements because of a change in LLDP MIB content.

(Default: 2; Range 1–8192)

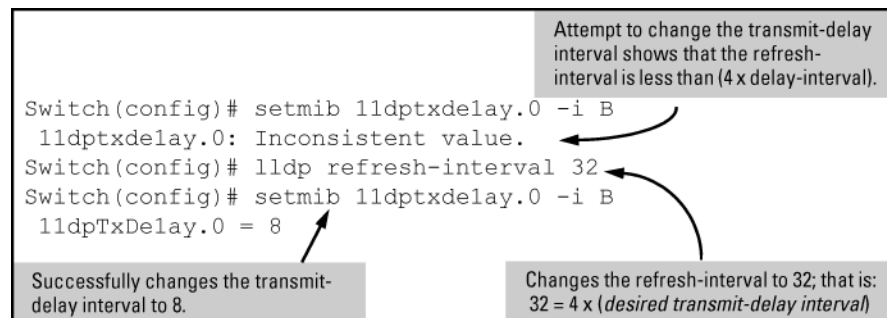


The LLDP refresh-interval (transmit interval) must be greater than or equal to (4 x delay-interval). The switch does not allow increasing the delay interval to a value that conflicts with this relationship. That is, the switch displays `Inconsistent value` if (4 x delay-interval) exceeds the current transmit interval, and the command fails. Depending on the current refresh-interval setting, it may be necessary to increase the refresh-interval before using this command to increase the delay-interval.

Example:

To change the delay-interval from 2 seconds to 8 seconds when the refresh-interval is at the default 30 seconds, you must first set the refresh-interval to a minimum of 32 seconds (32 = 4 x 8). (See [Figure 35: Changing the transmit-delay interval](#) on page 173.)

Figure 35: *Changing the transmit-delay interval*



Reinitialization delay interval

In the default configuration, a port receiving a `disable` command followed immediately by a `txonly`, `rxonly`, or `tx_rx` command delays reinitializing for two seconds, during which LLDP operation remains disabled. If an active port is subjected to frequent toggling between the LLDP disabled and enabled states, LLDP advertisements are more frequently transmitted to the neighbor device. Also, the neighbor table in the adjacent device changes more frequently as it deletes, then replaces LLDP data for the affected port which, in turn, generates SNMP traps (if trap receivers and SNMP notification are configured). All of this can unnecessarily increase network traffic. Extending the reinitialization-delay interval delays the ability of the port to reinitialize and generate LLDP traffic following an LLDP disable/enable cycle.

Changing the reinitialization delay interval (CLI)

Syntax:

```
setmib lldpReinitDelay.0 -i <1-10>
```

Uses `setmib` to change the minimum time (reinitialization delay interval) an LLDP port will wait before reinitializing after receiving an LLDP disable command followed closely by a `txonly` or `tx_rx` command. The delay interval commences with execution of the `lldp admin-status port-list disable` command.

(Default: 2 seconds; Range 1–10 seconds)

Example:

The following command changes the reinitialization delay interval to five seconds:

```
switch(config)# setmib lldpreinitdelay.0 -i 5
```

Configuring SNMP notification support

You can enable SNMP trap notification of LLDP data changes detected on advertisements received from neighbor devices, and control the interval between successive notifications of data changes on the same neighbor.

Enabling LLDP data change notification for SNMP trap receivers (CLI)

Syntax:

```
[no] lldp enable-notification <port-list>
```

Enables or disables each port in *port-list* for sending notification to configured SNMP trap receivers if an LLDP data change is detected in an advertisement received on the port from an LLDP neighbor.

(Default: Disabled)

For information on configuring trap receivers in the switch, see [SNMP notifications](#) on page 148.

Example:

This command enables SNMP notification on ports 1 - 5:

```
switch(config)# lldp enable-notification 1-5
```

Changing the minimum interval for successive data change notifications for the same neighbor

If LLDP trap notification is enabled on a port, a rapid succession of changes in LLDP information received in advertisements from one or more neighbors can generate a high number of traps. To reduce this effect, you can globally change the interval between successive notifications of neighbor data change.

Syntax:

```
setmib lldpnotificationinterval.0 -i <1-3600>
```

Globally changes the interval between successive traps generated by the switch. If multiple traps are generated in the specified interval, only the first trap is sent. The remaining traps are suppressed. (A network management application can periodically check the switch MIB to detect any missed change notification traps. See IEEE P802.1AB or later for more information.)

(Default: 5 seconds)

Example:

The following command limits change notification traps from a particular switch to one per minute.

```
switch(config)# setmib lldpnotificationinterval.0 -i 60 lldpNotificationInterval.0=60
```

Configuring per-port transmit and receive modes (CLI)

Syntax:

```
lldp admin-status <port-list> {<txonly | rxonly | tx_rx | disable>}
```

With LLDP enabled on the switch in the default configuration, each port is configured to transmit and receive LLDP packets. These options enable you to control which ports participate in LLDP traffic and whether the participating ports allow LLDP traffic in only one direction or in both directions.

<code>txonly</code>	Configures the specified ports to transmit LLDP packets, but block inbound LLDP packets from neighbor devices.
<code>rxonly</code>	Configures the specified ports to receive LLDP packets from neighbors, but block outbound packets to neighbors.
<code>tx_rx</code>	Configures the specified ports to both transmit and receive LLDP packets. (This is the default setting.)
<code>disable</code>	Disables LLDP packet transmit and receive on the specified ports.

Basic LLDP per-port advertisement content

In the default LLDP configuration, outbound advertisements from each port on the switch include both mandatory and optional data.

Mandatory Data

An active LLDP port on the switch always includes the mandatory data in its outbound advertisements. LLDP collects the mandatory data, and, except for the Remote Management Address, you cannot use LLDP commands to configure the actual data.

- Chassis Type (TLV subelement)
- Chassis ID (TLV)
- Port Type (TLV subelement)
- Port ID (TLV)
- Remote Management Address (TLV; actual IP address is a subelement that can be a default address or a configured address)

Configuring a remote management address for outbound LLDP advertisements (CLI)

This is an optional command you can use to include a specific IP address in the outbound LLDP advertisements for specific ports. For more information, see **Basic LLDP per-port advertisement content** on page 175.

Syntax:

```
[no] lldp config <port-list> ipAddrEnable <ip-address>
```

Replaces the default IP address for the port with an IP address you specify. This can be any IP address configured in a static VLAN on the switch, even if the port does not belong to the VLAN configured with the selected IP address.

The `no` form of the command deletes the specified IP address.

If there are no IP addresses configured as management addresses, the IP address selection method returns to the default operation.

Default: The port advertises the IP address of the lowest-numbered VLAN (VID) to which it belongs. If there is no IP address configured on the VLANs to which the port belongs, and if the port is not configured to advertise an IP address from any other (static) VLAN on the switch, the port advertises an address of 127.0.0.1.)



This command does not accept either IP addresses acquired through DHCP or Bootp, or IP addresses that are not configured in a static VLAN on the switch.

Example:

If port 3 belongs to a subnetted VLAN that includes an IP address of 10.10.10.100 and you want port 3 to use this secondary address in LLDP advertisements, you need to execute the following command:

```
switch(config)# lldp config 3 ipAddrEnable 10.10.10.100
```

Syntax:

[no] lldp config <port-list> basicTlvEnable <TLV-Type>

port_descr	For outbound LLDP advertisements, this TLV includes an alphanumeric string describing the port.(Default: Enabled)
system_name	For outbound LLDP advertisements, this TLV includes an alphanumeric string showing the assigned name of the system.(Default: Enabled)
system_descr	For outbound LLDP advertisements, this TLV includes an alphanumeric string describing the full name and version identification for the hardware type, software version, and networking application of the system.(Default: Enabled)
system_cap	For outbound advertisements, this TLV includes a bitmask of supported system capabilities (device functions). Also includes information on whether the capabilities are enabled. (Default: Enabled)

Example:

If you want to exclude the system name TLV from the outbound LLDP advertisements for all ports on a switch, use this command:

```
switch(config)# no lldp config 1-24 basicTlvEnable system_name
```

If you later decide to reinstate the system name TLV on ports 1-5, use this command:

```
switch(config)# lldp config 1-5 basicTlvEnable system_name
```

Optional Data

You can configure an individual port or group of ports to exclude one or more of the following data types from outbound LLDP advertisements.

- Port description (TLV)
- System name (TLV)
- System description (TLV)
- System capabilities (TLV)
 - System capabilities Supported (TLV subelement)
 - System capabilities Enabled (TLV subelement)
- Port speed and duplex (TLV subelement)

Optional data types, when enabled, are populated with data internal to the switch; that is, you cannot use LLDP commands to configure their actual content.

Support for port speed and duplex advertisements

This feature is optional for LLDP operation, but is **required** for LLDP-MED operation.

Port speed and duplex advertisements are supported on the switches to inform an LLDP endpoint and the switch port of each other's port speed and duplex configuration and capabilities. Configuration mismatches between a switch port and an LLDP endpoint can result in excessive collisions and voice quality degradation. LLDP enables discovery of such mismatches by supporting SNMP access to the switch MIB for comparing the current switch port and endpoint settings. (Changing a current device configuration to eliminate a mismatch requires intervention by the system operator.)

An SNMP network management application can be used to compare the port speed and duplex data configured in the switch and advertised by the LLDP endpoint. You can also use the CLI to display this information. For more information on using the CLI to display port speed and duplex information, see [Viewing the current port speed and duplex configuration on a switch port](#) on page 190.

Configuring support for port speed and duplex advertisements (CLI)

For more information, see [Support for port speed and duplex advertisements](#) on page 177.

Syntax:

```
[no] lldp config <port-list> dot3TlvEnable macphy_config
```

Options

<code>macphy_config</code>	MAC Physical Config TLV
<code>poepplus_config</code>	Power Via MDI Config TLV
<code>eee_config</code>	EEE Config TLV

For outbound advertisements, this TLV includes the (local) switch port's current speed and duplex settings, the range of speed and duplex settings the port supports, and the method required for reconfiguring the speed and duplex settings on the device (autonegotiation during link initialization, or manual configuration).

Using SNMP to compare local and remote information can help in locating configuration mismatches.

(Default: Enabled)



For LLDP operation, this TLV is optional. For LLDP-MED operation, this TLV is mandatory.

Port VLAN ID TLV support on LLDP

The `port-vlan-id` option enables advertisement of the port VLAN ID TLV as part of the regularly advertised TLVs. This allows discovery of a mismatch in the configured native VLAN ID between LLDP peers. The information is visible using `show` commands and is logged to the Syslog server.

Configuring the VLAN ID TLV

This TLV advertisement is enabled by default. To enable or disable the TLV, use this command. For more information, see [Port VLAN ID TLV support on LLDP](#) [Port VLAN ID TLV support on LLDP](#) on page 177.

Syntax:

```
[no] lldp config <port-list> dot1TlvEnable port-vlan-id
```

Enables the VLAN ID TLV advertisement.

The `no` form of the command disables the TLV advertisement.

Default: Enabled.

Options

<code>port-vlan-id</code>	Specifies the 802.1 TLV list to advertise.
<code>vlan-name</code>	Specifies that the VLAN name TLV is to be advertised.

Enabling the VLAN ID TLV

```
HP Switch(config)# lldp config a1 dot1TlvEnable port-vlan-id
```

Viewing the TLVs advertised

The `show` commands display the configuration of the TLVs. The command `show lldp config` lists the TLVs advertised for each port, as shown in the following examples.

Displaying the TLVs for a port

```
switch(config)# show lldp config a1

LLDP Port Configuration Detail

Port      : A1
AdminStatus [Tx_Rx] : Tx_Rx
NotificationEnabled [False] : False
Med Topology Trap Enabled [False] : False

TLVS Advertised:
* port_descr
* system_name
* system_descr
* system_cap

* capabilities
* network_policy
* location_id
* poe

* macphy_config

* port_vlan_id 1

IpAddress Advertised:
```

```
:
```

- ¹The VLAN ID TLV is being advertised.

Local device LLDP information

```
switch(config)# show lldp config info local-device a1
```

```
LLDP Port Configuration Information Detail
```

```
Port      : A1
PortType  : local
PortId    : 1
PortDesc  : A1
```

```
Port VLAN ID : 1 1
```

- ¹The information that LLDP used in its advertisement.

Remote device LLDP information

```
switch(config)# show lldp info remote-device a1
```

```
LLDP Remote Device Information Detail
```

```
Local Port      : A1
ChassisType     : mac-address
ChassisId       : 00 16 35 22 ca 40
PortType        : local
PortID          : 1
SysName         : esp-dback
System Descr    : HP J8693A Switch 3500yl-48G, revision XX.13.03, ROM...
PortDescr      : A1
```

```
System Capabilities Supported : bridge, router
System Capabilities Enabled   : bridge, router
```

```
Port VLAN ID : 200
```

```
Remote Management Address
Type      : ipv4
Address   : 192.168.1.1
```

SNMP support

The LLDP-EXT-DOT1-MIB has the corresponding MIB variables for the Port VLAN ID TLV. The TLV advertisement can be enabled or disabled using the MIB object `lldpXdot1ConfigPortVlanTxEnable` in the `lldpXdot1ConfigPortVlanTable`.

The port VLAN ID TLV local information can be obtained from the MIB object `lldpXdot1LocPortVlanId` in the local information table `lldpXdot1LocTable`.

The port VLAN ID TLV information about all the connected peer devices can be obtained from the MIB object `lldpXdot1RemPortVlanId` in the remote information table `lldpXdot1RemTable`.

LLDP-MED (media-endpoint-discovery)

LLDP-MED (ANSI/TIA-1057/D6) extends the LLDP (IEEE 802.1AB) industry standard to support advanced features on the network edge for Voice Over IP (VoIP) endpoint devices with specialized capabilities and LLDP-MED standards-based functionality. LLDP-MED in the switches uses the standard LLDP commands described


earlier in this section, with some extensions, and also introduces new commands unique to LLDP-MED operation. The `show` commands described elsewhere in this section are applicable to both LLDP and LLDP-MED operation. LLDP-MED benefits include:

- Plug-and-play provisioning for MED-capable, VoIP endpoint devices
- Simplified, vendor-independent management enabling different IP telephony systems to interoperate on one network
- Automatic deployment of convergence network policies (voice VLANs, Layer 2/CoS priority, and Layer 3/QoS priority)
- Configurable endpoint location data to support the Emergency Call Service (ECS) (such as Enhanced 911 service, 999, 112)
- Detailed VoIP endpoint data inventory readable via SNMP from the switch
- Power over Ethernet (PoE) status and troubleshooting support via SNMP
- support for IP telephony network troubleshooting of call quality issues via SNMP

This section describes how to configure and use LLDP-MED features in the switches to support VoIP network edge devices (media endpoint devices) such as:

- IP phones
- Voice/media gateways
- Media servers
- IP communications controllers
- Other VoIP devices or servers

Figure 36: Example: of LLDP-MED network elements



LLDP-MED endpoint support

LLDP-MED interoperates with directly connected IP telephony (endpoint) clients having these features and services:

- Autonegotiate speed and duplex configuration with the switch
- Use the following network policy elements configured on the client port
 - Voice VLAN ID
 - 802.1p (Layer 2) QoS
 - Diffserv codepoint (DSCP) (Layer 3) QoS
- Discover and advertise device location data learned from the switch
- Support ECS (such as E911, 999, and 112)
- Advertise device information for the device data inventory collected by the switch, including:

<ul style="list-style-type: none"> ◦ Hardware revision ◦ Firmware revision 	<ul style="list-style-type: none"> ◦ Software revision ◦ Serial number 	<ul style="list-style-type: none"> ◦ Manufacturer name ◦ Model name 	<ul style="list-style-type: none"> ◦ Asset ID
--	--	---	--

- Provide information on network connectivity capabilities (For example, a multi-port VoIP phone with Layer 2 switch capability)
- Support the fast-start capability



LLDP-MED is intended for use with VoIP endpoints and is not designed to support links between network infrastructure devices, such as switch-to-switch or switch-to-router links.

LLDP-MED endpoint device classes

LLDP-MED endpoint devices are, by definition, located at the network edge and communicate using the LLDP-MED framework. Any LLDP-MED endpoint device belongs to one of the following three classes:

- Class 1 (generic endpoint devices): These devices offer the basic LLDP discovery services, network policy advertisement (VLAN ID, Layer 2/802.1p priority, and Layer 3/DSCP priority), and PoE management. This class includes such devices as IP call controllers and communication-related servers.
- Class 2 (media endpoint devices): These devices offer all Class 1 features plus media-streaming capability, and include such devices as voice/media gateways, conference bridges, and media servers.
- Class 3 (communication devices): These devices are typically IP phones or end-user devices that otherwise support IP media and offer all Class 1 and Class 2 features, plus location identification and emergency 911 capability, Layer 2 switch support, and device information management.

LLDP-MED operational support

The switches offer two configurable TLVs supporting MED-specific capabilities:

- medTlvEnable (for per-port enabling or disabling of LLDP-MED operation)
- medPortLocation (for configuring per-port location or emergency call data)



LLDP-MED operation also requires the port speed and duplex TLV (dot3TlvEnable), which is enabled in the default configuration.

Topology change notifications provide one method for monitoring system activity. However, because SNMP normally employs UDP, which does not guarantee datagram delivery, topology change notification should not be relied upon as the sole method for monitoring critical endpoint device connectivity.

LLDP-MED fast start control

Syntax:

```
lldp fast-start-count <1-10>
```

An LLDP-MED device connecting to a switch port may use the data contained in the MED TLVs from the switch to configure itself. However, the `lldp refresh-interval` setting (default: 30 seconds) for transmitting advertisements can cause an unacceptable delay in MED device configuration.

To support rapid LLDP-MED device configuration, the `lldp fast-start-count` command temporarily overrides the `refresh-interval` setting for the `fast-start-count` advertisement interval. This results in the port initially advertising LLDP-MED at a faster rate for a limited time. Thus, when the switch detects a new LLDP-MED device on a port, it transmits one LLDP-MED advertisement per second out the port for the duration of the `fast-start-count` interval. In most cases, the default setting should provide an adequate `fast-start-count` interval.

(Default: 5 seconds)



This global command applies only to ports on which a new LLDP-MED device is detected. It does not override the `refresh-interval` setting on ports where non-MED devices are detected.

Advertising device capability, network policy, PoE status and location data

The `medTlvEnable` option on the switch is enabled in the default configuration and supports the following LLDP-MED TLVs:

- LLDP-MED capabilities: This TLV enables the switch to determine:
 - Whether a connected endpoint device supports LLDP-MED
 - Which specific LLDP-MED TLVs the endpoint supports
 - The device class (1, 2, or 3) for the connected endpoint

This TLV also enables an LLDP-MED endpoint to discover what LLDP-MED TLVs the switch port currently supports.

- Network policy operating on the port to which the endpoint is connected (VLAN, Layer 2 QoS, Layer 3 QoS)
- PoE (MED Power-over-Ethernet)
- Physical location data (see [Configuring location data for LLDP-MED devices](#) on page 185)



LLDP-MED operation requires the `macphy_config` TLV subelement (enabled by default) that is optional for IEEE 802.1AB LLDP operation. For more information, see the `dot3TlvEnable macphy_config` command ([Configuring support for port speed and duplex advertisements \(CLI\)](#) on page 177).

Network policy advertisements

Network policy advertisements are intended for real-time voice and video applications, and include these TLV subelements:

- Layer 2 (802.1p) QoS
- Layer 3 DSCP (diffserv code point) QoS
- Voice VLAN ID (VID)

VLAN operating rules

These rules affect advertisements of VLANs in network policy TLVs:

- The VLAN ID TLV subelement applies only to a VLAN configured for voice operation (`vlan <vid> voice`).
- If there are multiple voice VLANs configured on a port, LLDP-MED advertises the voice VLAN having the lowest VID.
- The voice VLAN port membership configured on the switch can be tagged or untagged. However, if the LLDP-MED endpoint expects a tagged membership when the switch port is configured for untagged, or the reverse, a configuration mismatch results. (Typically, the endpoint expects the switch port to have a tagged voice VLAN membership.)
- If a given port does not belong to a voice VLAN, the switch does not advertise the VLAN ID TLV through this port.

Policy elements

These policy elements may be statically configured on the switch or dynamically imposed during an authenticated session on the switch using a RADIUS server and 802.1X or MAC authentication. (Web authentication does not apply to VoIP telephones and other telecommunications devices that are not capable of accessing the switch through a Web browser.) The QoS and voice VLAN policy elements can be statically configured with the following CLI commands:

```
vlan <vid> voice
```

```
vlan <vid> {<tagged | untagged> <port-list>}
int <port-list> qos priority <0-7>
vlan <vid> qos dscp <codepoint>
```



A codepoint must have an 802.1p priority before you can configure it for use in prioritizing packets by VLAN-ID. If a codepoint you want to use shows **No Override** in the **Priority** column of the DSCP policy table (display with `show qos-dscp map`, then use `qos-dscp map <codepoint> priority <0-7>` to configure a priority before proceeding. For more information on this topic, see the "Quality of Service (QoS): Managing Bandwidth More Effectively" in the advanced traffic management guide for your switch.

Enabling or Disabling medTlvEnable

In the default LLDP-MED configuration, the TLVs controlled by `medTlvEnable` are enabled. For more information, see **Advertising device capability, network policy, PoE status and location data** on page 182.

Syntax:

```
[no] lldp config <port-list> medTlvEnable <medTlv>
```

Enables or disables advertisement of the following TLVs on the specified ports:

- Device capability TLV
- Configured network policy TLV
- Configured location data TLV (see **Configuring location data for LLDP-MED devices** on page 185.)
- Current PoE status TLV

(Default: All of the above TLVs are enabled.)

Helps to locate configuration mismatches by allowing use of an SNMP application to compare the LLDP-MED configuration on a port with the LLDP-MED TLVs advertised by a neighbor connected to that port.

<p>capabilities</p>	<p>This TLV enables the switch to determine:</p> <ul style="list-style-type: none"> • Which LLDP-MED TLVs a connected endpoint can discover • The device class (1, 2, or 3) for the connected endpoint <p>This TLV also enables an LLDP-MED endpoint to discover what LLDP-MED TLVs the switch port currently supports.(Default: enabled)</p> <p>This TLV cannot be disabled unless the <code>network_policy</code>, <code>poe</code>, and <code>location_id</code> TLVs are already disabled.</p>
<p>network_policy</p>	<p>This TLV enables the switch port to advertise its configured network policies (voice VLAN, Layer 2 QoS, Layer 3 QoS), and allows LLDP-MED endpoint devices to autoconfigure the voice network policy advertised by the switch. This also enables the use of SNMP applications to troubleshoot statically configured endpoint network policy mismatches.(Default: Enabled)</p> <p>Network policy is advertised only for ports that are configured as members of the voice VLAN. If the port belongs to more than one voice VLAN, the voice VLAN with the lowest-numbered VID is selected as the VLAN for voice traffic. Also, this TLV cannot be enabled unless the <code>capability</code> TLV is already enabled.</p> <p>For more information, see Network policy advertisements on page 182.</p>

Table Continued

location_id	<p>This TLV enables the switch port to advertise its configured location data (if any). For more information on configuring location data, see Configuring location data for LLDP-MED devices on page 185.(Default: Enabled)</p> <p>When disabled, this TLV cannot be enabled unless the capability TLV is already enabled.</p>
poe	<p>This TLV enables the switch port to advertise its current PoE state and to read the PoE requirements advertised by the LLDP-MED endpoint device connected to the port.(Default: Enabled)</p> <p>When disabled, this TLV cannot be enabled unless the <code>capability</code> TLV is already enabled.</p> <p>For more on this topic, see PoE advertisements on page 184.</p>

PoE advertisements

These advertisements inform an LLDP-MED endpoint of the power (PoE) configuration on switch ports. Similar advertisements from an LLDP-MED endpoint inform the switch of the endpoint's power needs and provide information that can be used to identify power priority mismatches.

PoE TLVs include the following power data:

- **Power type:** indicates whether the device is a power-sourcing entity (PSE) or a PD. Ports on the J8702A PoE zl module are PSE devices. A MED-capable VoIP telephone is a PD.
- **Power source:** indicates the source of power in use by the device. Power sources for PDs include PSE, local (internal), and PSE/local. The switches advertise Unknown.
- **Power priority:** indicates the power priority configured on the switch (PSE) port or the power priority configured on the MED-capable endpoint.
- **Power value:** indicates the total power in watts that a switch port (PSE) can deliver at a particular time, or the total power in watts that the MED endpoint (PD) requires to operate.

Viewing PoE advertisements

To display the current power data for an LLDP-MED device connected to a port, use the following command:

```
show lldp info remote-device <port-list>
```

For more information on this command, see page A-60.

To display the current PoE configuration on the switch, use the following commands:

```
show power brief <port-list>
```

```
show power <port-list>
```

For more information on PoE configuration and operation, see **[Power Over Ethernet \(PoE/PoE+\) Operation](#)**.

Location data for LLDP-MED devices

You can configure a switch port to advertise location data for the switch itself, the physical wall-jack location of the endpoint (recommended), or the location of a DHCP server supporting the switch, endpoint, or both. You also have the option of configuring these different address types:

- **Civic address:** physical address data such as city, street number, and building information
- **ELIN (Emergency Location Identification Number):** an emergency number typically assigned to MLTS (Multiline Telephone System) Operators in North America
- **Coordinate-based location:** attitude, longitude, and altitude information (Requires configuration via an SNMP application.)

Configuring location data for LLDP-MED devices

Syntax:

```
[no] lldp config <port-list> medPortLocation <Address-Type>
```

Configures location of emergency call data the switch advertises per port in the `location_id` TLV. This TLV is for use by LLDP-MED endpoints employing location-based applications.



The switch allows one `medPortLocation` entry per port (without regard to type). Configuring a new `medPortLocation` entry of any type on a port replaces any previously configured entry on that port.

```
civic-addr <COUNTRY-STR> <WHAT> <CA-TYPE> <CA-VALUE> ... [< CA-TYPE > < CA-VALUE >] ... [< CA-TYPE > < CA-VALUE >]
```

Enables configuration of a physical address on a switch port and allows up to 75 characters of address information.

COUNTRY-STR	A two-character country code, as defined by ISO 3166. Some examples include <code>FR</code> (France), <code>DE</code> (Germany), and <code>IN</code> (India). This field is required in a <code>civic-addr</code> command. (For a complete list of country codes, visit http://www.iso.org .)
WHAT	A single-digit number specifying the type of device to which the location data applies: 0 : Location of DHCP server 1 : Location of switch 2 : Location of LLDP-MED endpoint (recommended application) This field is required in a <code>civic-addr</code> command.

Table Continued

Type/Value Pairs (CA-TYPE and CA-VALUE)

A series of data pairs, each composed of a location data "type" specifier and the corresponding location data for that type. That is, the first value in a pair is expected to be the civic address "type" number (CA-TYPE), and the second value in a pair is expected to be the corresponding civic address data (CA-VALUE).

For example, if the CA-TYPE for "city name" is "3," the type/value pair to define the city of Paris is "3 Paris."

Multiple type/value pairs can be entered in any order, although Hewlett Packard Enterprise recommends that multiple pairs be entered in ascending order of the CA-TYPE.

When an emergency call is placed from a properly configured class 3 endpoint device to an appropriate PSAP, the country code, device type, and type/value pairs configured on the switch port are included in the transmission. The "type" specifiers are used by the PSAP to identify and organize the location data components in an understandable format for response personnel to interpret.

A `civic-addr` command requires a minimum of one type/value pair, but typically includes multiple type/value pairs as needed to configure a complete set of data describing a given location.

CA-TYPE: This is the first entry in a type/value pair and is a number defining the type of data contained in the second entry in the type/value pair (CA-VALUE). Some examples of CA-TYPE specifiers include:

- 3=city
- 6=street (name)
- 25=building name

(Range: 0 - 255)For a sample listing of CA-TYPE specifiers, see **Some location codes used in CA-TYPE fields.**

CA-VALUE: This is the second entry in a type/value pair and is an alphanumeric string containing the location information corresponding to the immediately preceding CA-TYPE entry.

Strings are delimited by either blank spaces, single quotes (' ... '), or double quotes ("... ").

Each string should represent a specific data type in a set of unique type/value pairs comprising the description of a location, and each string must be preceded by a CA-TYPE number identifying the type of data in the string.

A switch port allows one instance of any given CA-TYPE. For example, if a type/value pair of 6 Atlantic (to

Table Continued

specify "Atlantic" as a street name) is configured on port A5 and later another type/value pair of 6 Pacific is configured on the same port, Pacific replaces Atlantic in the civic address location configured for port A5.

`elin-addr <emergency-number>`

This feature is intended for use in ECS applications to support class 3 LLDP-MED VoIP telephones connected to a switch in an MLTS infrastructure.

An ELIN is a valid NANP format telephone number assigned to MLTS operators in North America by the appropriate authority. The ELIN is used to route emergency (E911) calls to a PSAP.

(Range: 1-15 numeric characters)

Configuring coordinate-based locations

Latitude, longitude, and altitude data can be configured per switch port using an SNMP management application. For more information, see the documentation provided with the application. A further source of information on this topic is RFC 3825-Dynamic host configuration protocol option for coordinate-based location configuration information.



Endpoint use of data from a medPortLocation TLV sent by the switch is device-dependent. See the documentation provided with the endpoint device.

Table 21: *Some location codes used in CA-TYPE fields*

Location element	Code ¹	Location element	Code
national subdivision	1	street number	19
regional subdivision	2	additional location data	22
city or township	3	unit or apartment	26
city subdivision	4	floor	27
street	6	room number	28
street suffix	18		

¹ The code assignments in this table are examples from a work-in-progress (the internet draft titled "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information draft-ietf-geopriv-dhcp-civil-06" dated May 30, 2005.) For the actual codes to use, contact the PSAP or other authority responsible for specifying the civic addressing data standard for your network.

Example:

Suppose a system operator wants to configure the following information as the civic address for a telephone connected to her company's network through port A2 of a switch at the following location:

CA-type	CA-type	CA-VALUE
national subdivision	1	CA

Table Continued

city	3	Widgitville
street	6	Main
street number	19	1433
unit	26	Suite 4-N
floor	27	4
room number	28	N4-3

The following example shows the commands for configuring and displaying the above data.

A civic address configuration

```
switch(config)# lldp config 2 medportlocation civic-addr US 2 1 CA 3
Widgitville 6 Main 19 1433 26 Suite_4-N 27 4 28 N4-3
```

```
switch(config)# show lldp config 2
LLDP Port Configuration Detail
Port : A2
AdminStatus [Tx_Rx] : Tx_Rx
NotificationEnabled [False] : False
Med Topology Trap Enabled [False] : False
Country Name      : US
What              : 2
Ca-Type           : 1
Ca-Length         : 2
Ca-Value          : CA
Ca-Type           : 3
Ca-Length         : 11
Ca-Value          : Widgitville
Ca-Type           : 6
Ca-Length         : 4
Ca-Value          : Main
Ca-Type           : 19
Ca-Length         : 4
Ca-Value          : 1433
Ca-Type           : 26
Ca-Length         : 9
Ca-Value          : Suite_4-N
Ca-Type           : 27
Ca-Length         : 1
Ca-Value          : 4
Ca-Type           : 28
Ca-Length         : 4
Ca-Value          : N4-3
```

Viewing switch information available for outbound advertisements

Syntax:

```
show lldp info local-device [port-list]
```

Without the [*port-list*] option, displays the global switch information and the per-port information currently available for populating outbound LLDP advertisements.

With the `[port-list]` option, displays only the following port-specific information that is currently available for outbound LLDP advertisements on the specified ports:

- PortType
- PortId
- PortDesc



This command displays the information available on the switch. Use the `lldp config <port-list>` command to change the selection of information that is included in actual outbound advertisements. In the default LLDP configuration, all information displayed by this command is transmitted in outbound advertisements.

In the default configuration, the switch information currently available for outbound LLDP advertisements appears similar to the display in the following example.

Displaying the global and per-port information available for outbound advertisements

```
switch(config)# show lldp info local-device
```

```
LLDP Local Device Information
```

```
Chassis Type : mac-address
Chassis Id : 00 23 47 4b 68 DD
System Name : HP Switch1
System Description : HP J9091A Switch 3500y1, revision XX.15.06...
System Capabilities Supported:bridge
System Capabilities Enabled:bridge
```

```
Management Address 1
```

```
  Type:ipv4
```

```
  Address:
```

```
LLDP Port Information
```

Port	PortType	PortId	PortDesc
1	local	1	1
2	local	2	2
3	local	3	3
4	local	4	4
5	local	5	5

- 1

The Management Address field displays only the LLDP-configurable IP addresses on the switch. (Only manually-configured IP addresses are LLDP-configurable.) If the switch has only an IP address from a DHCP or Bootp server, then the Management Address field is empty (because there are no LLDP-configurable IP addresses available).

The default per-port information content for ports 1 and 2

```
switch(config)# show lldp info local 1-2
```

```
LLDP Local Port Information Detail
```

```
Port      : 1
PortType  : local
PortId    : 1
PortDesc  : 1
```

```
-----  
Port      : 2  
PortType  : local  
PortId    : 2  
PortDesc  : 2
```

Displaying the current port speed and duplex configuration on a switch port

You can compare port speed and duplex information for a switch port and a connected LLDP-MED endpoint for configuration mismatches by using an SNMP application. You can also use the switch CLI to display this information, if necessary. The `show interfaces brief <port-list>` and `show lldp info remote-device [port-list]` (**Remote power information** on page 94) commands provide methods for displaying speed and duplex information for switch ports. For information on displaying the currently configured port speed and duplex on an LLDP-MED endpoint, see **Viewing the current port speed and duplex configuration on a switch port** on page 190.

Viewing the current port speed and duplex configuration on a switch port

Syntax:

```
show interfaces brief <port-list>
```

Includes port speed and duplex configuration in the `Mode` column of the resulting display.

Viewing advertisements currently in the neighbors MIB

Syntax:

```
show lldp info remote-device [port-list]
```

Without the `[port-list]` option, provides a global list of the individual devices it has detected by reading LLDP advertisements. Discovered devices are listed by the inbound port on which they were discovered.

Multiple devices listed for a single port indicates that such devices are connected to the switch through a hub.

Discovering the same device on multiple ports indicates that the remote device may be connected to the switch in one of the following ways:

- Through different VLANS using separate links. (This applies to switches that use the same MAC address for all configured VLANs.)
- Through different links in the same trunk.
- Through different links using the same VLAN. (In this case, spanning-tree should be invoked to prevent a network topology loop. Note that LLDP packets travel on links that spanning-tree blocks for other traffic types.)

With the `[port-list]` option, provides a listing of the LLDP data that the switch has detected in advertisements received on the specified ports.

For descriptions of the various types of information displayed by these commands, see **Data available for basic LLDP advertisements**.

A global listing of discovered devices

```
switch(config)# show lldp info remote
```

```
LLDP Remote Devices Information
```

```
LocalPort | ChassisId                               PortId PortDescr SysName  
----- + -----
```

1	00 11 85 35 3b 80	6	6	HP Switch
2	00 11 85 cf 66 60	8	8	HP Switch

An LLLDP-MED listing of an advertisement received from an LLDP-MED (VoIP telephone) source

```
switch(config)# show lldp info remote-device 1
```

```
LLDP Remote Device Information Detail

Local Port      : A2
ChassisType    : network-address
ChassisId      : 0f ff 7a 5c
PortType       : mac-address
PortId        : 08 00 0f 14 de f2
SysName        : HP Switch
System Descr   : HP Switch, revision xx.15.06.0000x
PortDescr     : LAN Port

System Capabilities Supported : bridge, telephone
System Capabilities Enabled   : bridge, telephone

Remote Management Address

MED Information Detail 1
  EndpointClass      :Class3
  Media Policy Vlan id :10
  Media Policy Priority :7
  Media Policy Dscp   :44
  Media Policy Tagged :False
  Poe Device Type     :PD
  Power Requested     :47
  Power Source        :Unknown
  Power Priority       :High
```

- ¹Indicates the policy configured on the telephone. A configuration mismatch occurs if the supporting port is configured differently.

Displaying LLDP statistics

LLDP statistics are available on both a global and a per-port levels. Rebooting the switch resets the LLDP statistics counters to zero. Disabling the transmit and/or receive capability on a port "freezes" the related port counters at their current values.

Viewing LLDP statistics

For more information, see [Displaying LLDP statistics](#) on page 191.

Syntax:

```
show lldp stats [port-list]
```

The **global LLDP** statistics command displays an overview of neighbor detection activity on the switch, plus data on the number of frames sent, received, and discarded per-port.

The **per-port LLDP** statistics command enhances the list of per-port statistics provided by the global statistics command with some additional per-port LLDP statistics.

Global LLDP Counters:

Neighbor Entries List Last Updated	The elapsed time since a neighbor was last added or deleted.
New Neighbor Entries Count	The total of new LLDP neighbors detected since the last switch reboot. Disconnecting, and then reconnecting a neighbor increments this counter.
Neighbor Entries Deleted Count	The number of neighbor deletions from the MIB for AgeOut Count and forced drops for all ports. For example, if the admin status for port on a neighbor device changes from <code>tx_rx</code> or <code>txonly</code> to <code>disabled</code> or <code>rxonly</code> , the neighbor device sends a "shutdown" packet out the port and ceases transmitting LLDP frames out that port. The device receiving the shutdown packet deletes all information about the neighbor received on the applicable inbound port and increments the counter.
Neighbor Entries Dropped Count	The number of valid LLDP neighbors the switch detected, but could not add. This can occur, For example, when a new neighbor is detected when the switch is already supporting the maximum number of neighbors. See <u>Neighbor maximum</u> on page 193.
Neighbor Entries AgeOut Count	The number of LLDP neighbors dropped on all ports because of Time-to-Live expiring.

Per-Port LLDP Counters:

NumFramesRecv	The total number of valid, inbound LLDP advertisements received from any neighbors on <code>port-list</code> . Where multiple neighbors are connected to a port through a hub, this value is the total number of LLDP advertisements received from all sources.
NumFramesSent	The total number of LLDP advertisements sent from <code>port-list</code> .
NumFramesDiscarded	The total number of inbound LLDP advertisements discarded by <code>port-list</code> . This can occur, For example, when a new neighbor is detected on the port, but the switch is already supporting the maximum number of neighbors. See <u>Neighbor maximum</u> on page 193. This can also be an indication of advertisement formatting problems in the neighbor device.
Frames Invalid	The total number of invalid LLDP advertisements received on the port. An invalid advertisement can be caused by header formatting problems in the neighbor device.
TLVs Unrecognized	The total number of LLDP TLVs received on a port with a type value in the reserved range. This can be caused by a basic management TLV from a later LLDP version than the one currently running on the switch.
TLVs Discarded	The total number of LLDP TLVs discarded for any reason. In this case, the advertisement carrying the TLV may be accepted, but the individual TLV is not usable.
Neighbor Ageouts	The number of LLDP neighbors dropped on the port because of Time-to-Live expiring.

Examples:

A global LLDP statistics display

```
switch(config)# show lldp stats
```

LLDP Device Statistics

```
Neighbor Entries List Last Updated : 2 hours
New Neighbor Entries Count : 20
Neighbor Entries Deleted Count : 20
Neighbor Entries Dropped Count : 0
Neighbor Entries AgeOut Count : 20
```

LLDP Port Statistics

Port	NumFramesRecvd	NumFramesSent	NumFramesDiscarded
A1	97317	97843	0
A2	21	12	0
A3	0	0	0
A4	446	252	0
A5	0	0	0
A6	0	0	0
A7	0	0	0
A8	0	0	0

A per-port LLDP statistics display

```
switch(config)# show lldp stats 1
```

LLDP Port Statistics Detail

```
PortName : 1
Frames Discarded : 0
Frames Invalid : 0
Frames Received : 7309
Frames Sent : 7231
TLVs Unrecognized : 0
TLVs Discarded : 0
Neighbor Ageouts : 0
```

LLDP Operating Notes

Neighbor maximum

The neighbors table in the switch supports as many neighbors as there are ports on the switch. The switch can support multiple neighbors connected through a hub on a given port, but if the switch neighbor maximum is reached, advertisements from additional neighbors on the same or other ports will not be stored in the neighbors table unless some existing neighbors time-out or are removed.

LLDP packet forwarding

An 802.1D-compliant switch does not forward LLDP packets, regardless of whether LLDP is globally enabled or disabled on the switch.

One IP address advertisement per port

LLDP advertises only one IP address per port, even if multiple IP addresses are configured by `lldp config port-list ipAddrEnable` on a given port.

802.1Q VLAN Information

LLDP packets do not include 802.1Q header information and are always handled as untagged packets.

Effect of 802.1X Operation

If 802.1X port security is enabled on a port, and a connected device is not authorized, LLDP packets are not transmitted or received on that port. Any neighbor data stored in the neighbor MIB for that port prior to the unauthorized device connection remains in the MIB until it ages out. If an unauthorized device later becomes authorized, LLDP transmit and receive operation resumes.

Neighbor data can remain in the neighbor database after the neighbor is disconnected

After disconnecting a neighbor LLDP device from the switch, the neighbor can continue to appear in the switch's neighbor database for an extended period if the neighbor's `holdtime-multiplier` is high; especially if the `refresh-interval` is large. See [Changing the time-to-live for transmitted advertisements \(CLI\)](#) on page 172.

Mandatory TLVs

All mandatory TLVs required for LLDP operation are also mandatory for LLDP-MED operation.

LLDP and CDP data management

This section describes points to note regarding LLDP and CDP (Cisco Discovery Protocol) data received by the switch from other devices. LLDP operation includes both transmitting LLDP packets to neighbor devices and reading LLDP packets received from neighbor devices. CDP operation is limited to reading incoming CDP packets from neighbor devices. (HPE switches do not generate CDP packets.)

Incoming CDP and LLDP packets tagged for VLAN 1 are processed even if VLAN 1 does not contain any ports. VLAN 1 must be present, but it is typically present as the default VLAN for the switch.



The switch may pick up CDP and LLDP multicast packets from VLAN 1 even when CDP- and /or LLDP-enabled ports are not members of VLAN 1.

LLDP and CDP neighbor data

With both LLDP and (read-only) CDP enabled on a switch port, the port can read both LLDP and CDP advertisements, and stores the data from both types of advertisements in its neighbor database. (The switch **stores** only CDP data that has a corresponding field in the LLDP neighbor database.) The neighbor database itself can be read by either LLDP or CDP methods or by using the `show lldp` commands. Take note of the following rules and conditions:

- If the switch receives both LLDP and CDP advertisements on the same port from the same neighbor, the switch stores this information as two separate entries if the advertisements have different chassis ID and port ID information.
- If the chassis and port ID information are the same, the switch stores this information as a single entry. That is, LLDP data overwrites the corresponding CDP data in the neighbor database if the chassis and port ID information in the LLDP and CDP advertisements received from the same device is the same.
- Data read from a CDP packet does not support some LLDP fields, such as "System Descr," "SystemCapSupported," and "ChassisType." For such fields, LLDP assigns relevant default values. Also:
 - The LLDP "System Descr" field maps to CDP's "Version" and "Platform" fields.
 - The switch assigns "ChassisType" and "PortType" fields as "local" for both the LLDP and the CDP advertisements it receives.
 - Both LLDP and CDP support the "System Capability" TLV. However, LLDP differentiates between what a device is capable of supporting and what it is actually supporting, and separates the two types of information into subelements of the System Capability TLV. CDP has only a single field for this data. Thus,

when CDP System Capability data is mapped to LLDP, the same value appears in both LLDP System Capability fields.

- System Name and Port Descr are not communicated by CDP, and thus are not included in the switch's Neighbors database.



Because HPE switches do not generate CDP packets, they are not represented in the CDP data collected by any neighbor devices running CDP.

A switch with CDP disabled forwards the CDP packets it receives from other devices, but does not store the CDP information from these packets in its own MIB.

LLDP data transmission/collection and CDP data collection are both enabled in the switch's default configuration. In this state, an SNMP network management application designed to discover devices running either CDP or LLDP can retrieve neighbor information from the switch regardless of whether LLDP or CDP is used to collect the device-specific information.

Protocol state	Packet generation	Inbound data management	Inbound packet forwarding
CDP Enabled ¹	N/A	Store inbound CDP data.	No forwarding of inbound CDP packets.
CDP Disabled	N/A	No storage of CDP data from neighbor devices.	Floods inbound CDP packets from connected devices to outbound ports.
LLDP Enabled ¹	Generates and transmits LLDP packets out all ports on the switch.	Store inbound LLDP data.	No forwarding of inbound LLDP packets.
LLDP Disabled	No packet generation.	No storage of LLDP data from neighbor devices.	No forwarding of inbound LLDP packets.

¹ Both CDP data collection and LLDP transmit/receive are enabled in the default configuration. If a switch receives CDP packets and LLDP packets from the same neighbor device on the same port, it stores and displays the two types of information separately if the chassis and port ID information in the two types of advertisements is different. In this case, if you want to use only one type of data from a neighbor sending both types, disable the unwanted protocol on either the neighbor device or on the switch. However, if the chassis and port ID information in the two types of advertisements is the same, the LLDP information overwrites the CDP data for the same neighbor device on the same port.

CDP operation and commands

By default the switches have CDP enabled on each port. This is a read-only capability, meaning that the switch can receive and store information about adjacent CDP devices but does not generate CDP packets.

When a CDP-enabled switch receives a CDP packet from another CDP device, it enters that device's data in the CDP Neighbors table, along with the port number where the data was received—and does not forward the packet. The switch also periodically purges the table of any entries that have expired. (The hold time for any data entry in the switch's CDP Neighbors table is configured in the device transmitting the CDP packet and cannot be controlled in the switch receiving the packet.) A switch reviews the list of CDP neighbor entries every three seconds and purges any expired entries.



For details on how to use an SNMP utility to retrieve information from the switch's CDP Neighbors table maintained in the switch's MIB, see the documentation provided with the particular SNMP utility.

Viewing the current CDP configuration of the switch

CDP is shown as enabled/disabled both globally on the switch and on a per-port basis.

Syntax:

```
show cdp
```

Lists the global and per-port CDP configuration of the switch.

The following example shows the default CDP configuration.

Default CDP configuration

```
switch(config)# show cdp

Global CDP information

  Enable CDP [Yes] : Yes (Receive Only)

Port CDP
-----
1      enabled
2      enabled
3      enabled
.      .
.      .
.      .
```

Viewing the current CDP neighbors table of the switch

Devices are listed by the port on which they were detected.

Syntax:

```
show cdp neighbors
```

Lists the neighboring CDP devices the switch detects, with a subset of the information collected from the device's CDP packet.

<code>[[e] port-numb [detail]]</code>	Lists the CDP device connected to the specified port. (Allows only one port at a time.) Using <code>detail</code> provides a longer list of details on the CDP device the switch detects on the specified port.
<code>[detail [[e] port-numb]]</code>	Provides a list of the details for all of the CDP devices the switch detects. Using <code>port-num</code> produces a list of details for the selected port.

The following example displays the CDP devices that the switch has detected by receiving their CDP packets.

CDP neighbors table listing

```
switch(config)# show cdp neighbors

CDP neighbors information
```

Port	Device ID	Platform	Capability
1	Accounting (0030c1-7fcc40)	J4812A HP Switch. . .	S
2	Research1-1 (0060b0-889e43)	J4121A HP Switch. . .	S
4	Support (0060b0_761a45)	J4121A HP Switch. . .	S
7	Marketing (0030c5_33dc59)	J4313A HP Switch. . .	S
12	Mgmt NIC (099a05-09df9b)	NIC Model X666	H
12	Mgmt NIC (099a05-09df11)	NIC Model X666	H

Enabling and Disabling CDP Operation

Enabling CDP operation (the default) on the switch causes the switch to add entries to its CDP Neighbors table for any CDP packets it receives from other neighboring CDP devices.

Disabling CDP operation clears the switch's CDP Neighbors table and causes the switch to drop inbound CDP packets from other devices without entering the data in the CDP Neighbors table.

Syntax:

```
[no] cdp run
```

Enables or disables CDP read-only operation on the switch.

(Default: Enabled)

Example:

To disable CDP read-only on the switch:

```
switch(config)# no cdp run
```

When CDP is disabled:

- `show cdp neighbors`
displays an empty CDP Neighbors table
- `show cdp`
displays Global CDP information
Enable CDP [Yes]: No

Enabling or disabling CDP operation on individual ports

In the factory-default configuration, the switch has all ports enabled to receive CDP packets. Disabling CDP on a port causes it to drop inbound CDP packets without recording their data in the CDP Neighbors table.

Syntax:

```
[no] cdp enable {< [e] port-list >}
```

Example:

To disable CDP on port A1:

```
switch(config)# no cdp enable a1
```

Filtering CDP information

In some environments it is desirable to be able to configure a switch to handle CDP packets by filtering out the MAC address learns from untagged VLAN traffic from IP phones. This means that normal protocol processing occurs for the packets, but the addresses associated with these packets is not learned or reported by the software

address management components. This enhancement also filters out the MAC address learns from LLDP and 802.1x EAPOL packets on untagged VLANs.

The feature is configured per-port.

Configuring the switch to filter untagged traffic

Enter this command to configure the switch not to learn CDP, LLDP, or EAPOL traffic for a set of interfaces.

Syntax:

```
[no] ignore-untagged-mac <port-list>
```

Prevents MAC addresses from being learned on the specified ports when the VLAN is untagged and the destination MAC address is one of the following:

- 01000C-CCCCC (CDP)
- 0180c2- 00000e (LLDP)
- 0180c2-000003 (EAPOL)

Configuring the switch to ignore packet MAC address learns for an untagged VLAN

```
switch(config) ignore-untagged-mac 1-2
```

Displaying the configuration

Enter the `show running-config` command to display information about the configuration.

Configuration showing interfaces to ignore packet MAC address learns

```
switch(config) show running-config
```

```
Running configuration:
```

```
; J9627 Configuration Editor; Created on release XX.15.XX  
; Ver #03:03.1f.ef:f0
```

```
hostname "HP Switch"  
interface 1  
    ignore-untagged-mac  
    exit  
interface 2  
    ignore-untagged-mac  
    exit  
.  
.  
.  
vlan 1  
    name "DEFAULT_VLAN"  
    untagged 1-24  
    ip address dhcp-bootp  
    exit  
.  
.  
.
```

Filtering PVID mismatch log messages

This enhancement filters out PVID mismatch log messages on a per-port basis. PVID mismatches are logged when there is a difference in the PVID advertised by a neighboring switch and the PVID of the switch port which receives the LLDP advertisement. Logging is an LLDP feature that allows detection of possible vlan leakage between adjacent switches. However, if these events are logged too frequently, they can overwhelm the log buffer and push relevant logging data out of log memory, making it difficult to troubleshoot another issue.

Logging is disabled and enabled with the support of CLI commands.

This enhancement also includes displaying the Mac-Address in the PVID mismatch log message when the port ID is Mac-Address instead of displaying garbage characters in the peer device port ID field.

Use the following command to disable the logging of the PVID mismatch log messages:

Syntax:

```
logging filter [filter-name][sub filter id] <regularexpression> deny
```

Regular-expression The regular expression should match the message which is to be filtered.

Syntax:

```
logging filter [filter-name] enable
```

Generic header ID in configuration file

Introduction

Auto deployment relies on DHCP options and the current DHCP auto-configuration function. Auto deployment is platform independent, avoiding the J-number validation of the downloaded configuration file when downloaded using DHCP option 66/67. The downloaded configuration file has an `IGNORE` tag immediately after the J-number in its header.

An option to add an `add-ignore-tag` to an existing copy command will insert an `ignore` tag into the configuration header. This insertion happens while transferring the configurations, (`startup configuration files` and `running configuration files`) from the switch to a configuration file setup on a remote server. The process uses TFTP/SFTP or can be accomplished with a serially connected workstation using XMODEM.

Add-Ignore-Tag option

The `add-ignore-tag` option is used in conjunction with the `copy` command to transfer the `startup configuration` or `running configuration` files from the switch to a remote server with `IGNORE` tag inserted into it.

The `IGNORE` tag is inserted into the first line of the configuration file directly after the J-number.

Configuration file

```
; J9782A IGNORE Configuration Editor; Created on release #YB.15.14.0000x
; Ver #04:63.ff.37.27:88
hostname "HP-2530-24"
snmp-server community "public" unrestricted
vlan 1
name "DEFAULT_VLAN"
no untagged 2,20-25
```



The J-number validation is ignored only when configuration file that contains the `IGNORE` tag is downloaded to a switch via DHCP option 66/67. When a `configuration` file containing the `IGNORE` tag is downloaded to a switch using CLI, SNMP or WebUI, the downloaded `configuration` file is only accepted if the J-number in it matches the J-number on the switch.

There is no change to the current switch configuration when executing the `copy` command with the `add-ignore-tag` option. The `IGNORE` tag is only added to the `configuration` file being exported to the external server. The `configuration` file stored on an external server is then downloaded to the switch using DHCP option 66 during bootup. If the `IGNORE` tag is available in the downloaded `configuration` file then the switch will avoid the J-number validation of the `configuration` file. The downloaded `configuration` file will then go through a line by line validation. Once the `configuration` file passes this validation, it gets updated in the flash. Once the `configuration` file has been updated, the switch will reboot automatically.



The J-number in the downloaded `configuration` file is replaced with that of the switch. The `IGNORE` tag is removed from the downloaded `configuration` file before updating it to flash. The `show running-configuration` command will not display the `IGNORE` tag but displays the switch's J-number as part of the output.

Copy with add-ignore-tag

```
switch(config)# copy startup-config tftp <ip-addr> <filename> add-ignore-tag
switch(config)# copy running-config tftp <ip-addr> <filename> add-ignore-tag
switch(config)# copy startup-config sftp <ip-addr> <filename> add-ignore-tag
switch(config)# copy running-config sftp <ip-addr> <filename> add-ignore-tag
switch(config)# copy startup-config xmodem add-ignore-tag
switch(config)# copy running-config xmodem add-ignore-tag
```

Configuration commands for the add-ignore-tag option

Configuration files can be transferred to the switch from a server using the following `copy` commands:

- `copy tftp`
- `copy xmodem`
- `copy sftp`

Copy commands

```
copy tftp < startup-config | running-config > < ip-address > < remote-file >[ pc |
unix ]
copy xmodem startup-config < pc | unix >
copy sftp < startup-config | running-config > < ip-address > < remote-file >
```

Configuration files that are downloaded using the `copy` commands as described in the **example** will be accepted by the switch if they pass J-number validations and line by line validations after download. The downloaded configuration file will be discarded by the switch if the validations fail. If the validations fail, the switch will work with its previous configuration.

Show logging commands for the add-ignore-tag option

The `show logging` command is used to locate errors during a configuration validation process. The event log catalogs entries with the ID#00158 and updates for each invalid entry found in the `configuration` file.

Show logging

```
-- Reverse event Log listing: Events Since Boot ----  
W 01/07/14 00:29:31 00158 update: line 13. Module command missing for port or  
invalid port: 36  
I 01/07/14 00:29:30 00131 tftp: Transfer completed  
I 01/07/14 00:29:29 00090 dhcp: Trying to download Config File (using TFTP)  
received in DHCP from 192.168.1.1
```



Downloading manually edited configuration file is not encouraged.

Exclusions

The `IGNORE` tag is not an available option when using external SCP, SFTP or TFTP clients such as PuTTY, OpenSSH, WinSCP and SSH Secure Shell to transfer configuration files out of the switch.

The Captive Portal feature allows the support of the ClearPass Policy Manager (CPPM) into the ArubaOS-Switch product line. The switch provides configuration to allow you to enable or disable the Captive Portal feature. By default, Captive Portal is disabled to avoid impacting existing installations as this feature is mutually exclusive with the following web-based authentication mechanisms: Web Authentication, EWA, MAFR, and BYOD Redirect.

Captive Portal is user-based, rather than port or VLAN-based, therefore the configuration is on a switch global basis. ArubaOS-Switch supports the following authentication types on the switch with RADIUS for Captive Portal:

- Media Access Control (MAC)
- 802.1X

Once you enable Captive Portal, the redirect functionality is triggered only if a redirect URL attribute is provided as part of the RADIUS Access-Accept response from an authentication request of type 802.1X or MAC. The redirect enables the client to self-register or directly login with valid credentials via the CPPM. Upon subsequent re-authentication, it provides access to the network per the CPPM configured policies that are communicated via the RADIUS attributes.

The redirect feature offers:

- Client self-registration
- Client direct login with valid credentials via CPPM Captive Portal
- On-boarding
- Ability to quarantine devices to remedy their status

Requirements

- HTTPS support requires a certificate to be configured on the switch with a usage type of `all` or `captive-portal`.
- If you are running HPE 5400 Series v2 modules, you must turn off the compatibility mode with the following command:

```
switch(config)# no allow-v1-modules
```

This will ensure that the switch will only power up with the v2 modules.

Best Practices

- Use the Port Bounce VSA via a CoA message, instead of the Disconnect message, to cause the second RADIUS authentication to occur during the Captive Portal exchange. This is the more reliable method for forcing a re-DHCP for the client.
- Configure Captive Portal such that the first `ACCESS_ACCEPT` returns a rate limit VSA to reduce the risk of DoS attacks. This configuration enables rate limiting for the HTTP/HTTPS ACL for traffic sent to CPPM.
- Do not use the keyword `cpy` in any other `NAS-Filter-Rules`. The keyword `cpy` in the enforcement profile attributes is specific to CPPM use. It is only supported with the `deny` attribute. If you configure the `cpy` keyword to `permit`, no ACL will be applied.

Limitations

- Captive Portal will not work with RADIUS configured on a loopback port or on the Out-of-Band Management (OOBM) port.
- Captive Portal is supported in CPPM versions 6.5.5 and later. However, by manually modifying the RADIUS dictionary files, any CPPM version 6.5.* can be used.
- Captive Portal does not support v1 modules, and will not work unless compatibility mode is turned off.
- Captive Portal does not support IPv6.
- Simultaneous Captive Portal client connections: maximum of 512
- Captive Portal does not support web proxy. The permit CPPM ACLs and the steal ACLs only use port 80 and 443. Non-standard ports for HTTP and HTTPS are not supported.
- Captive Portal is mutually exclusive with the following web-based authentication mechanisms: Web Authentication, EWA, MAFR, and BYOD.
- URL-string limitation of 253 characters.

Features

High Availability

Captive Portal includes support for High Availability (HA). The Captive Portal configurations (such as enablement, authenticated clients, and redirect URLs) are replicated to standby or other members.

If the feature is enabled and a failover occurs, clients in the process of onboarding are still redirected to Captive Portal, and authenticated clients continue to have the same access to the network.

Clients that are in the process of authenticating via MAC or 802.1X authentication will not be replicated to the standby. Replication of client data is only done when MAC or 802.1X authentication has resulted in a successful authentication.

Load balancing and redundancy

The following options are available to create load balancing and provide redundancy for CPPM:

- Virtual IP use for a CPPM server cluster
- CPPM servers configured in the switch RADIUS server group
- External load balancer

Captive Portal when disabled

By default, Captive Portal is disabled. If the Captive Portal feature is disabled and the switch receives a redirect URL attribute from the RADIUS server as part of the Access-Accept, it will view the redirect as an error. The authentication success will be overridden, the session will be flushed, and the switch will send the Accounting Start and Accounting Stop messages to indicate the client is no longer authenticated.

The Captive Portal feature may be disabled while there are in flight authentication requests. These are authentication sessions that have not finished the final authentication with the switch. The switch flushes all sessions with a redirect URL associated with them when Captive Portal is disabled.

Fully authenticated sessions are not impacted when Captive Portal is disabled. If CPPM deems these sessions to be invalid, a RADIUS Disconnect can be sent to flush all these sessions.

Disabling Captive Portal

To disable Captive Portal, enter one of the following:

```
switch(config)# aaa authentication captive-portal disable
```

```
switch(config)# no aaa authentication captive-portal enable
```

Configuring Captive Portal on CPPM

Procedure

1. Import the HP RADIUS dictionary
2. Create enforcement profiles
3. Create a ClearPass guest self-registration
4. Configure the login delay

Import the HP RADIUS dictionary

For CPPM versions 6.5.*, you must update the HP RADIUS dictionary. To import the dictionary in CPPM, follow these steps:

Procedure

1. Go to **Administration** -> **Dictionaries** -> **RADIUS** and click **Import**.
2. Select the XML HP RADIUS Dictionary from your Hard Drive.
3. Click **Import**.

Create enforcement profiles



Create the HPE Bounce Host-Port profile and the Guest Login profile only if they do not already exist.

For the HPE Bounce Host-Port profile, configure Captive Portal so that the RADIUS CoA message that includes the Port Bounce VSA is sent to force the second RADIUS re-authentication after the user registers their device and makes it known.

Procedure

1. In CPPM, go to **Configuration** -> **Enforcement** -> **Profiles**
2. Click **Add**.
3. Enter the Profile Name: **HPE Bounce Host-Port**
4. Enter the Description: **Custom-defined profile to bounce host port (HPE)**.
5. Select the type **RADIUS_CoA**.
6. Select the action **CoA**.
7. Add all of the attributes required for a CoA message, and specify the port bounce duration (valid values are between 0 and 60). This is the amount of time in seconds the port will be held in the down state. The recommended setting is 12 seconds.

Summary		Profile	Attributes
Profile:			
Name:	HPE Bounce Host-Port		
Description:	Custom-defined profile to bounce host port (HPE)		
Type:	RADIUS_CoA		
Action:	CoA		
Device Group List:	-		
Attributes:			
Type	Name		Value
1. Radius:IETF	User-Name	=	%{Radius:IETF:User-Name}
2. Radius:IETF	Calling-Station-Id	=	%{Radius:IETF:Calling-Station-Id}
3. Radius:IETF	NAS-Port	=	%{Radius:IETF:NAS-Port}
4. Radius:IETF	NAS-IP-Address	=	%{Radius:IETF:NAS-IP-Address}
5. Radius:IETF	Event-Timestamp	=	%{Radius:IETF:Event-Timestamp}
6. Radius:HPE	HPE-Port-Bounce-Host	=	12

- Repeat **Step 2** to **Step 6** to configure the Guest Login profile that will be sent as part of the first RADIUS Access-Accept and enforce the redirect to the Captive Portal on CPPM. For this profile, select **RADIUS** as the type and **Accept** as the action.
- Add all of the NAS-Filter-Rule attributes specified below, replacing the IP address in the first two NAS-Filter-Rule attributes with your CPPM address. Add the HPE-Captive-Portal-URL attribute to specify the redirect URL, replacing the IP address with your CPPM address. This will cause the client to be redirected to the Captive Portal on CPPM. You can add other attributes, such as a VLAN to isolate onboarding clients, or a rate limit to help prevent DoS attacks.



The HPE-Captive-Portal-URL value must be a URL normalized string. The scheme and host must be in lower case, for example `http://www.example.com/`.

Summary		Profile	Attributes
Profile:			
Name:	HPE Wired Guest Login		
Description:			
Type:	RADIUS		
Action:	Accept		
Device Group List:	-		
Attributes:			
Type	Name		Value
1. Radius:IETF	Tunnel-Type	=	VLAN (13)
2. Radius:IETF	Tunnel-Medium-Type	=	IEEE-802 (6)
3. Radius:IETF	Tunnel-Private-Group-Id	=	100
4. Radius:HPE	HPE-Captive-Portal-URL	=	http://10.73.4.136/guest/aruba_guest.php
5. Radius:IETF	NAS-Filter-Rule	=	permit in tcp from any to 10.73.4.136 80
6. Radius:IETF	NAS-Filter-Rule	=	permit in tcp from any to 10.73.4.136 443
7. Radius:IETF	NAS-Filter-Rule	=	deny in tcp from any to any 80 cpy
8. Radius:IETF	NAS-Filter-Rule	=	deny in tcp from any to any 443 cpy
9. Radius:IETF	NAS-Filter-Rule	=	permit in udp from any to any 53
10. Radius:IETF	NAS-Filter-Rule	=	permit in udp from any to any 67

Create a ClearPass guest self-registration

Procedure

- From the Customize Guest Registration window, select **Server-initiated** as the Login Method.
- Optionally, under Security Hash, select the level of checking to apply to the redirect URL.

Customize Guest Registration	
Login Options controlling logging in for self-registered guests.	
Enabled:	Enable guest login to a Network Access Server
* Vendor Settings:	Aruba Networks Select a predefined group of settings suitable for standard network configurations.
Login Method:	Server-initiated -- Change of authorization (RFC 3576) sent to controller Select how the user's network login will be handled. Server-initiated logins require the user's MAC address to be available, usually from the captive portal redirection process.
Security Hash:	Do not check -- login will always be permitted Select the level of checking to apply to URL parameters passed to the web login page. Use this option to detect when URL parameters have been modified by the user, for example their MAC address.
Default Destination Options for controlling the destination clients will redirect to after login.	
* Default URL:	<input type="text"/> Enter the default URL to redirect clients. Please ensure you prepend "http://" for any external domain.
Override Destination:	<input type="checkbox"/> Force default destination for all clients If selected, the client's default destination will be overridden regardless of its value.
<input type="button" value="Save Changes"/> <input type="button" value="Save and Continue"/>	

Configure the login delay

Enter the **Login Delay** value. The value must be greater than the `HPE-Port-Bounce-Host` attribute. In this example, we set the login delay value to 20 seconds.

Automatic Login

Options controlling automatically logging in from the receipt form.

* Login Delay:	<input type="text" value="20"/> seconds The time in seconds to delay while displaying the login message.
----------------	---

Social Logins

Optionally present guests with various social login options.

Social Login:	<input type="checkbox"/> Enable login with social network credentials
---------------	---

Configuring the switch

Once you have configured Captive Portal, you can configure the switch. To configure the switch, you must first configure the switch as a RADIUS client, then configure the ports that will be used for Captive Portal, as follows:

Procedure

1. Configure the switch as a RADIUS client. In this example, the CPPM IP address is `10.73.4.136` and `secret` is the secret key shared with the RADIUS server:
 - a. `switch(config)# radius-server host 10.73.4.136 key "secret"`
 - b. `switch(config)# radius-server host 10.73.4.136 dyn-authorization`
 - c. `switch(config)# radius-server host 10.73.4.136 time-window 0`



Make sure to set your time-window to 0. See **Event Timestamp not working**.

2. Configure the ports that will be used for Captive Portal. In this example, the commands enable ports `B3-B5` for MAC Authentication:

- a. `switch(config)# aaa authentication port-access chap-radius`
- b. `switch(config)# aaa port-access mac-based B3-B5`
3. If you configured the Security Hash to Deny login on validation error in **Create a ClearPass guest self-registration** on page 205, configure the URL key.
See **Configure the URL key** on page 207.
4. Configure the certificate. See **Configuring a certificate for Captive Portal usage** on page 207
5. Enable Captive portal:

```
switch(config)# aaa authentication captive-portal enable
```

By default, Captive Portal is disabled. Once enabled, you are redirected to the URL supplied via the HPE-Captive-Portal-URL VSA. Captive Portal is enabled on a global/switch wide basis.

Configure the URL key

You can optionally configure a URL hash key to provide some security for the Captive Portal exchange with CPPM. The key is a shared secret between CPPM and the switch. When configured, the switch generates a HMAC-SHA1 hash of the entire redirect URL, and appends the hash to the URL to be sent to CPPM as part of the HTTP redirect. If CPPM is configured to check the hash, it will generate the hash of the URL using its version of the URL hash key and compare against the value provided by the switch. The action taken by CPPM upon a match or mismatch is determined by what is configured on CPPM.

CPPM provides the following options:

- Do not check - login will always be permitted
- Deny login on validation error - login will not be permitted

The URL hash key is globally configured and will be used for all redirects to Captive Portal. This key is not configured on a per CPPM or RADIUS server basis. If the key is not specified, the hash is not added to the URL. The URL hash key is an ASCII string with a maximum length of 64 characters.

The URL key supports the FIPS certification feature `encrypt-credentials` and can optionally be encrypted for more robust security. This option is only available when the global `encrypt-credentials` is enabled.

To configure a plain text captive-portal URL key:

```
switch(config)# aaa authentication captive-portal url-hash-key plaintext <KEY>
```

To configure an encrypted captive-portal URL key when `encrypt-credentials` is enabled:

```
switch(config)# aaa authentication captive-portal url-hash-key encrypted  
<ENCRYPTED-KEY>
```

To clear a captive-portal URL key:

```
switch(config)# no aaa authentication captive-portal url-hash-key
```

Configuring a certificate for Captive Portal usage

HTTPS support requires the use of a certificate. If a certificate for Captive Portal does not exist, the certificate designated for all use is used instead.

To create a certificate signing request for Captive Portal, enter:

```
switch(config)# crypto pki create-csr certificate-name <cert-name> usage captive-portal
```

To create a self-signed certificate for Captive Portal, enter:

```
switch(config)# crypto pki enroll-self-signed certificate-name
```

Display Captive Portal configuration

To display the Captive Portal configuration settings, enter the `show captive-portal` command:

```
switch(config)# show captive-portal
```

```
Captive Portal Configuration
Redirection Enabled      : Yes
URL Hash Key Configured : No
```

Show certificate information

To view the certificate information, enter:

```
switch(config)# show crypto pki local-certificate
```

Name	Usage	Expiration	Parent / Profile
-----	-----	-----	-----
cp	Captive Portal	2016/08/14	default

Troubleshooting

Event Timestamp not working

Symptom

The client gets a credentials request on the web browser even though the valid credentials were already provided, or the client is not redirected to the Captive Portal.

Cause

- ClearPass 6.5.x does not support the sending of `Event Timestamp` in automated workflows (manual via Access Tracker works).
- The switch will reject CoA requests when the time on CPPM is ahead of the switch time by even a second.

Action

Procedure

1. Set the time-window security feature in PVOS to 0:


```
radius-server host<CLEARPASS-IP> time-window 0
```

Cannot enable Captive Portal

Symptom

When running the `aaa authentication captive-portal enable` command, getting the following error message:

```
Captive portal cannot be enabled when BYOD redirect, MAC authentication failure redirect, or web-based authentication are enabled.
```

Cause

The failure is due to a mutual exclusion restriction.

Action

Procedure

1. Check which one of the following are enabled: BYOD redirect, MAC authentication failure redirect, or web-based authentication.
2. Disabled the enabled authentication method found in step 1.
3. Run the `aaa authentication captive-portal enable` command.

Unable to enable feature

Symptom

One of the following messages is displayed:

- `BYOD redirect cannot be enabled when captive portal is enabled.`
- `MAC authentication failure redirect cannot be enabled when captive portal is enabled.`
- `Web-based authentication cannot be enabled when captive portal is enabled.`
- `V1 compatibility mode cannot be enabled when captive portal is enabled.`

Cause

You cannot enable these features when Captive Portal is already enabled. They are mutually exclusive.

Action

Procedure

1. You can either disable Captive Portal or avoid enabling these features.

Authenticated user redirected to login page

Symptom

User is redirected back to the login page to submit credentials even after getting fully authenticated.

Solution 1

Cause

The status is not changed to Known.

Action

Procedure

1. After the client submits the credentials, the CPPM service must change the Endpoint Status to `Known`.

Solution 2

Cause

The cache value is set.

Action

Procedure

1. Clear the CPPM Cache Timeout of the Endpoint Repository.

Unable to configure a URL hash key

Symptom

The following message is displayed:

```
Key exceeds the maximum length of 64 characters.
```

Cause

The URL hash key is not valid.

Action

Procedure

1. Select a key that is 64 or less ASCII text. For example:

```
switch(config)# aaa authentication captive-portal url-hash-key plaintext  
"8011A89FEAE0234BCCA"
```

authentication command

Use the following authentication commands to configure ClearPass Captive Portal.

Command	Description
<code>aaa authentication captive-portal enable</code>	Enables redirection to a Captive Portal server for additional client authentication.
<code>aaa authentication captive-portal disable</code> <code>or</code> <code>no aaa authentication captive-portal enable</code>	Disables redirection to a Captive Portal server for additional client authentication.
<code>aaa authentication captive-portal url-hash-key</code>	Configures a hash key used to verify the integrity of the portal URL.

show command

Use the following show commands to view the various configurations and certificates.

Command	Description
<code>show running-config</code>	Shows the running configuration.
<code>show config</code>	Shows the saved configuration.
<code>show ip</code>	Shows the switch IP addresses.
<code>show captive-portal</code>	Captive portal configuration.
<code>show port-access clients [port] [detailed]</code>	Consolidated client view; the <code>detailed</code> option shows the Access Policy that is applied. The IP address is only displayed if <code>dhcp-snooping</code> is enabled. For the summary view (without the <code>detailed</code> option), only the untagged VLAN is displayed.
<code>show radius authentication</code>	Displays NAS identifier and data on the configured RADIUS server and switch interactions with this server.
<code>show radius dyn-authorization</code>	Statistics for Radius CoA and Disconnect.
<code>show radius accounting</code>	Statistics for Radius accounting.
<code>show crypto pki local-certificate [summary]</code>	Installed certificates.

Debug command

Use the `debug` command to help you debug your issues.

Command	Description
<code>debug security captive-portal</code>	Enables debug logging for the Captive Portal sub-system.
<code>debug security port-access mac-based</code>	Enables debug logging for the MAC-auth sub-system.
<code>debug security port-access authenticator</code>	Enables debug logging for the 802.1X authenticator sub-system.
<code>debug security radius-server</code>	Enables debug logging for the Radius sub-system.
<code>debug destination session</code>	Prints debug messages to terminal.
<code>debug destination logging</code>	Sends debug messages to the syslog server.
<code>debug destination buffer</code>	Prints debug messages to a buffer in memory.

AirWave is a Network Management Solution (NMS) tool. Once connected to AirWave using the WebUI and CLI interfaces, you can:

- Configure your switches using Zero Touch Provisioning (ZTP)
- Configure your switches using the CLI
- Troubleshoot your switches
- Monitor your switches
- Upgrade your firmware for your switches

Once you have configured your switch, you can monitor, manage, and upgrade your hardware using the AirWave Management Platform.

Requirements

- DHCP server
- AirWave NMS
- HPE Aruba switches

Best Practices

- Implement ZTP in a secure and private environment. Any public access may compromise the security of the switch, as follows:
 - ZTP is enabled only on the factory default configuration of the switch. DHCP snooping is not enabled. You must manage the Rogue DHCP server.
 - The DHCP offer is in plain data without encryption. Therefore, the offer can be listened by any device on the network and they can in turn obtain the AirWave information.
 - The TLS certificate of the server is not validated by the switch during the HTTPs check-in to AirWave. The AirWave server is in the private environment of the switch.

Limitations

- ZTP is not supported through OOBM.
- The DNS/hostname in option 66 is not supported, only the IPv4 address.
- The switch does not validate peer certificate of the AirWave server as part of the TLS handshake.
- The HTTPS check-in to AirWave does not support HTTPS proxy.
- For non-ZTP cases, the AirWave check-in starts by validating the following condition: Primary or Management VLAN must be configured with the IP address and one of the interface must be UP. By default, `VLAN 1` is the primary VLAN.

Switch configuration

To configure your switch, follow these steps:

1. **Configure AirWave details in DHCP (preferred method)** on page 214.



If you are using existing HPE switches and using the DHCP server for the configuration or firmware management, you can configure the AirWave details in DHCP using this method:
Configure AirWave details in DHCP (alternate method) on page 219

2. If you are configuring the switch using a CLI, see **Configure a switch using the CLI** on page 227.

If you are using ZTP, the configuration is automatic and does not require any user interaction, see **Zero Touch Provisioning**

The switch contacts the AirWave server that is configured on the switch and initiates the check-in process.

Once you have configured the DHCP server, the AirWave details received from the DHCP options are stored in the switch configuration. This assures that the configuration is retained across reboots.

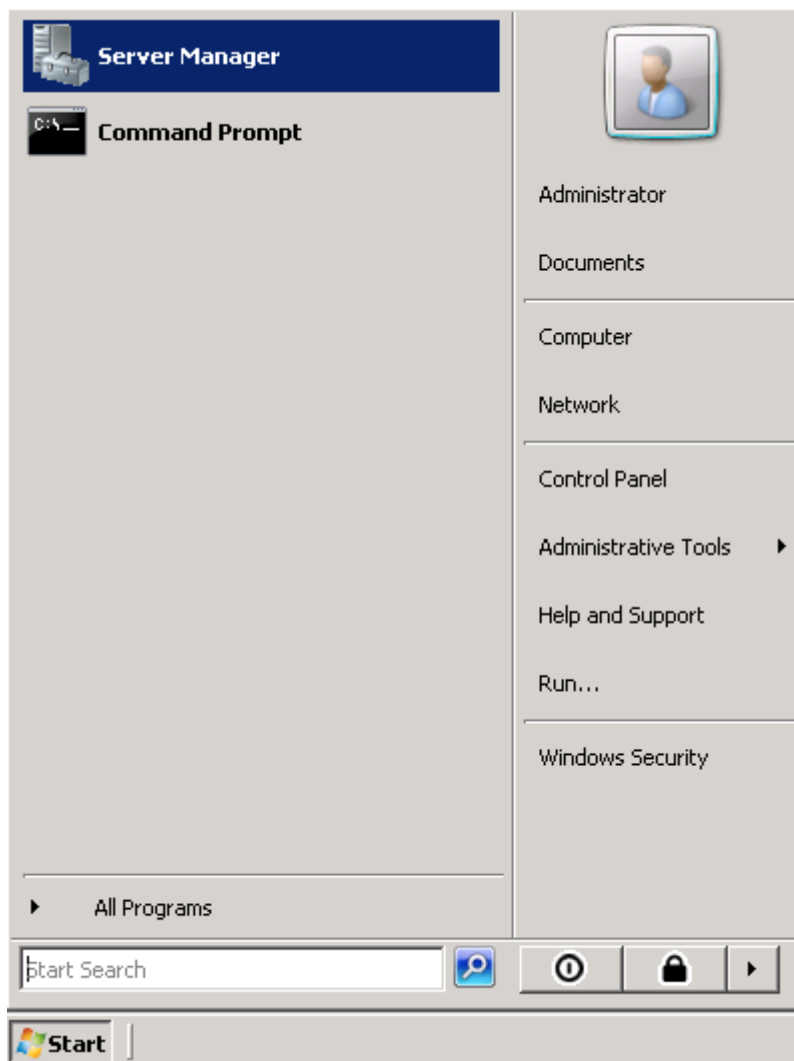
Once AirWave completes the switch check-in, it lists the first switch as `New Devices`. The first switch is used to create a new configuration template for the specific group and device type. With this new template, the required configuration is generated for the group. Subsequent switch of the specific type and joining the same group as the first device are added directly to the group and the configuration is pushed using the configuration template via a SSH connection.

Configure AirWave details in DHCP (preferred method)


To configure a DHCP server for AirWave, from a Windows Server 2008, do the following steps:

Procedure


1. From the **Start** menu, select **Server Manager**.



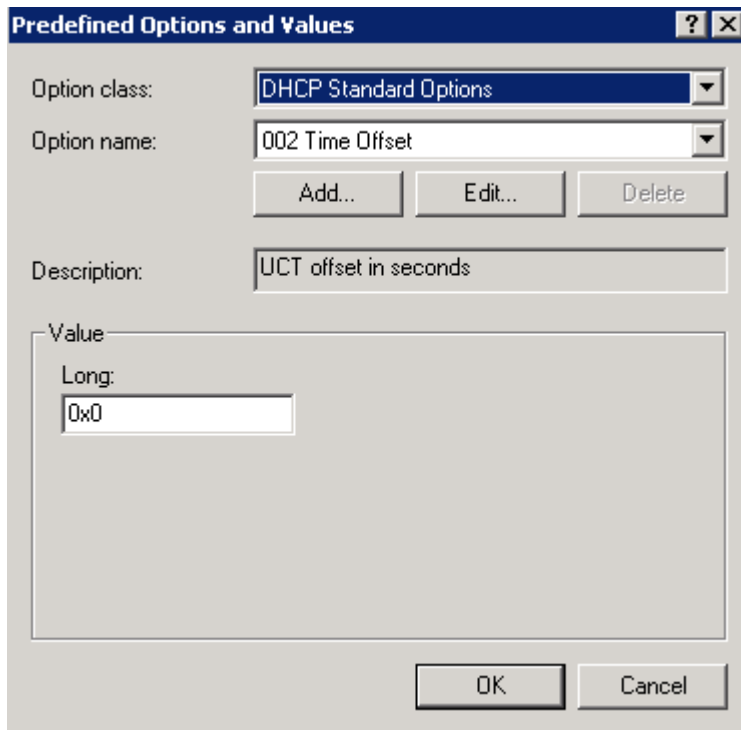
2. Select **Roles** -> **DHCP** -> **Server** -> **w2k8** -> **IPv4**.




3. Right click on **IPv4** and select **Set Predefined Options...**



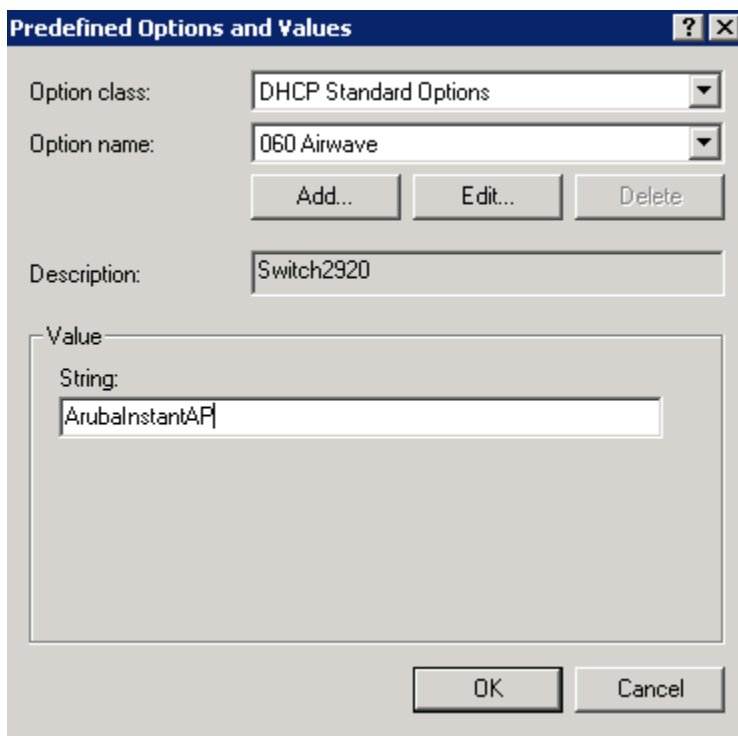
4. The Predefined Options and Values screen is displayed. Click **Add...**



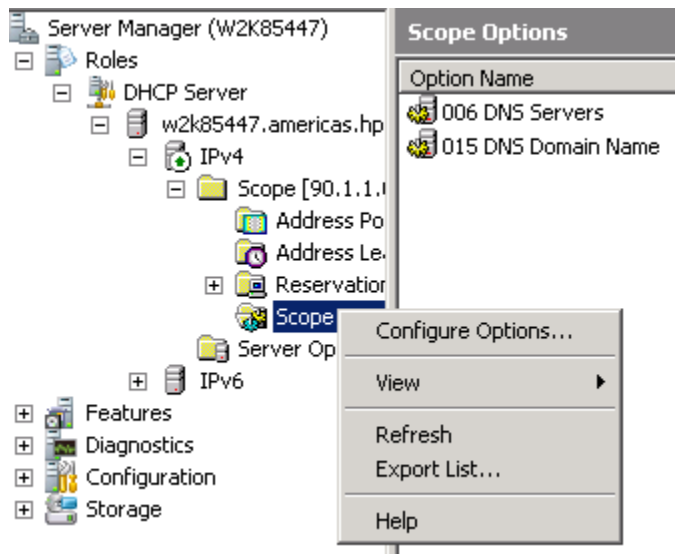
5. Enter the desired **Name** (any), **Data type** (select **String**), **Code** (enter **60**), and **Description** (any).



6. Click **OK**.
7. From the Predefined Options and Values screen, under Value, enter the String **ArubaInstantAP**. The string is case sensitive and must be `ArubaInstantAP`.



8. Click **OK**.
9. Under IPv4, expand **Scope**. Right click on **Scope Options** and select **Configure Options...**

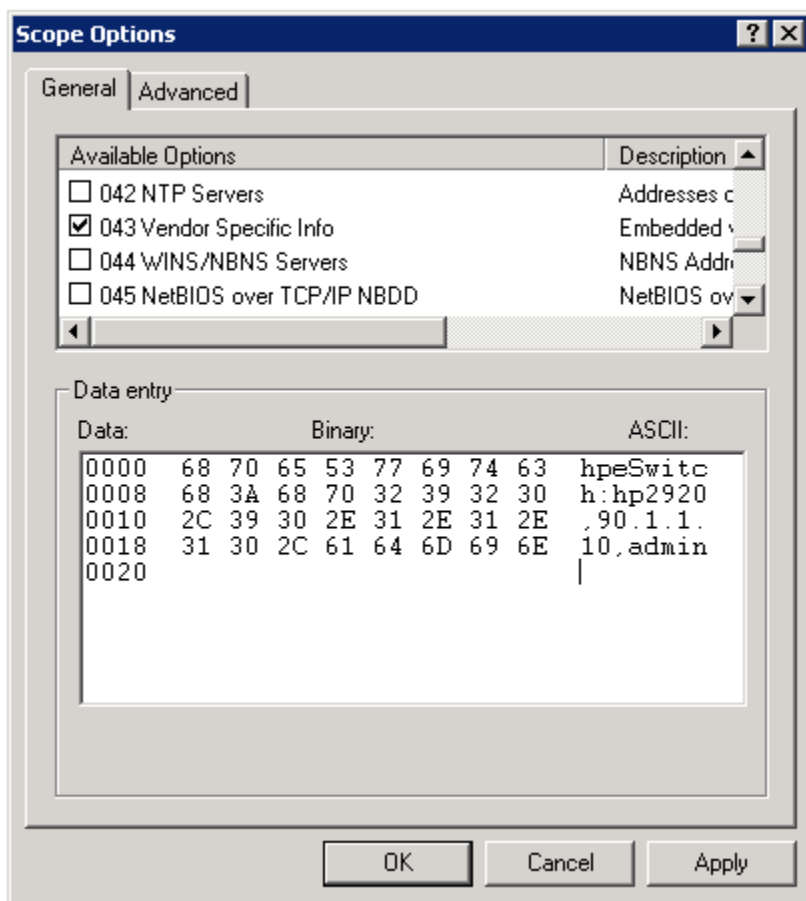


- Under the General tab, select **043 Vendor Specific Info**. The Data entry data appears. Under ASCII, enter **hpeSwitch:hp2920,90.1.1.10,admin**. The ASCII value has the following format:
`<Group>:<Topfolder>,<AMP IP>,<shared secret>`

If you need to add sub-folders, use the following format:

`<Group>:<Topfolder>:<folder1>,<AMP IP>,<shared secret>`


-



- Under the General tab, select **060 Airwave**. Click **OK**.



No changes are required to the 060 option.



13. You can verify the AirWave details as follows:

```
switch# show amp-server
switch# show run
```

Configure AirWave details in DHCP (alternate method)


To configure a DHCP server for ZTP and AirWave, from a Windows Server 2008, do the following steps:



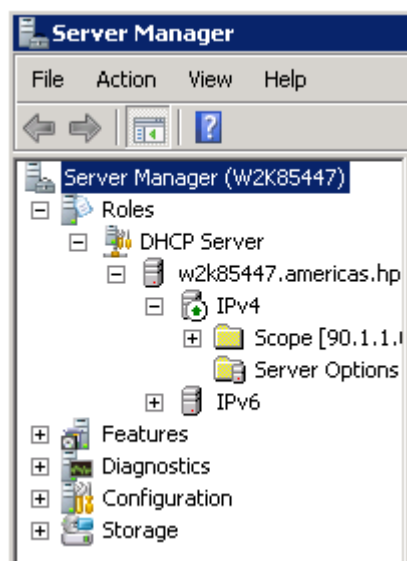
You must repeat these steps for every type of switch that needs to be configured for ZTP, selecting a different Vendor Class for each type of switch.

Procedure


1. From the **Start** menu, select **Server Manager**.




2. Select **Roles** -> **DHCP** -> **Server** -> **w2k8** -> **IPv4**.



3. Right click on **IPv4** and select **Define Vendor Classes...**



4. The DHCP Vendor Classes window is displayed. Click **Add...**



5. To get the vendor-specific value of a switch, go to the switch console and enter:

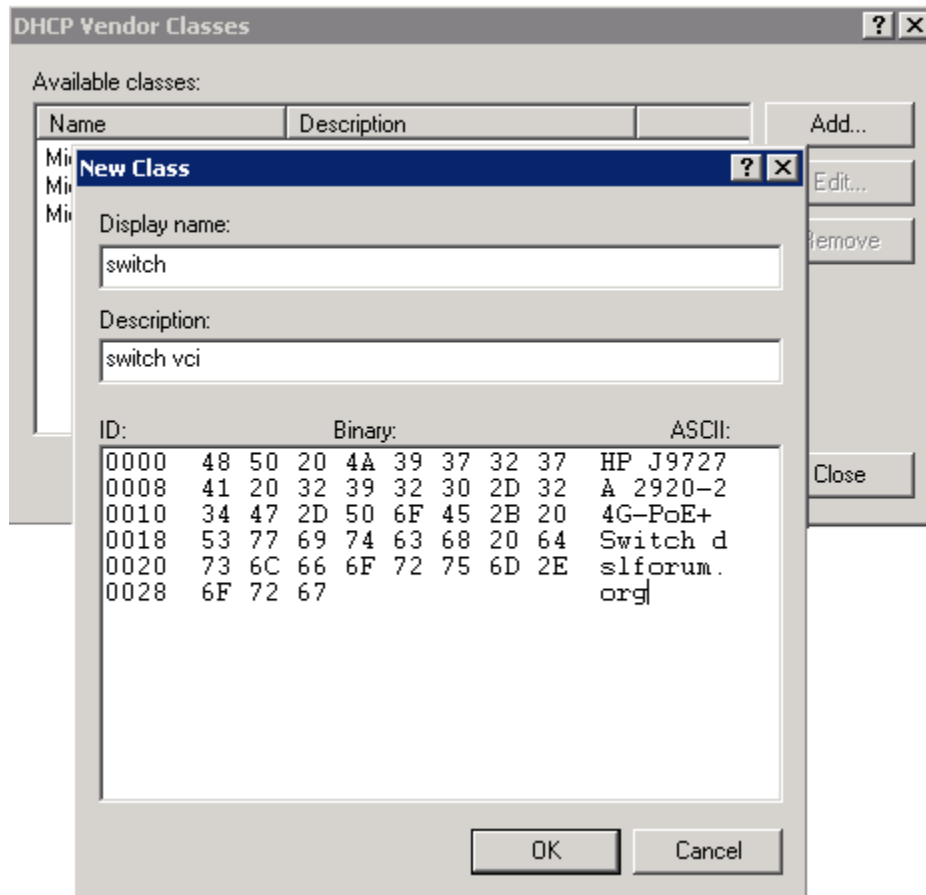
```
switch# show dhcp client vendor-specific
```

In our example, the command returns the following value:

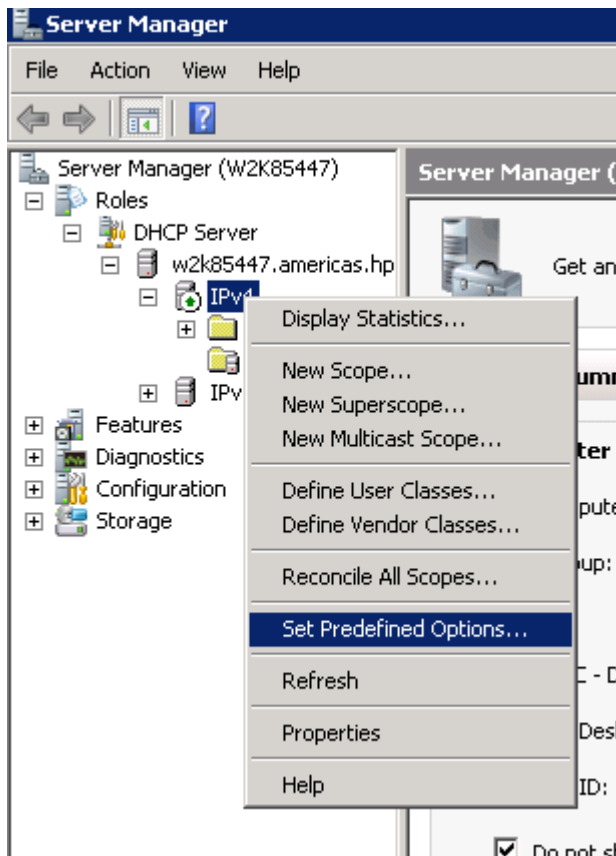
```
Vendor Class Id = HP J9729A 2920-24G-PoE+ Switch dslforum.org
```

Processing of Vendor Specific Configuration is enabled

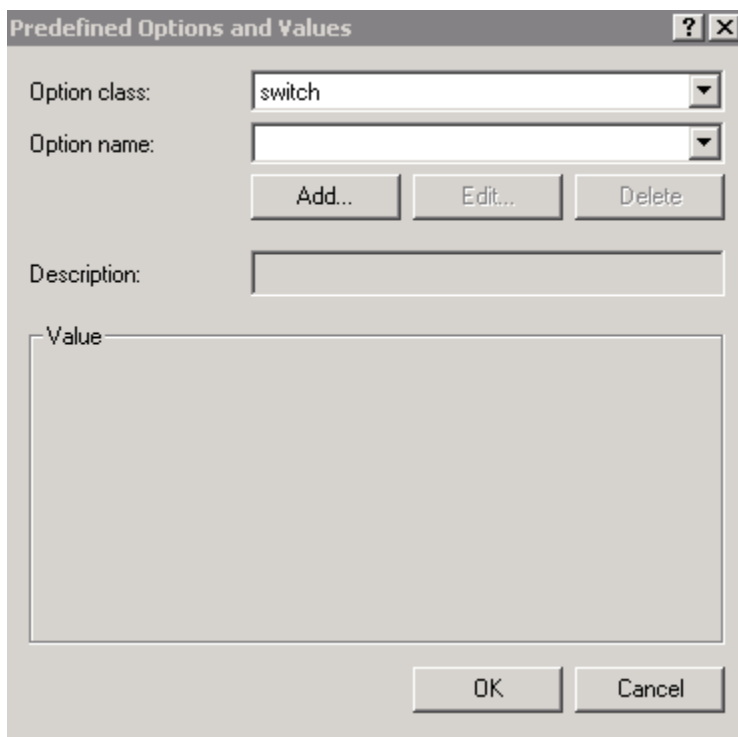
- From the New Class window, enter the desired **Display name** (any) and the **Description** (any). For the **ASCII** field, enter the exact value that you got by executing the `show` command performed in the previous step. In this example, **HP J9729A 2920-24G-PoE+ Switch dsiforum.org**.



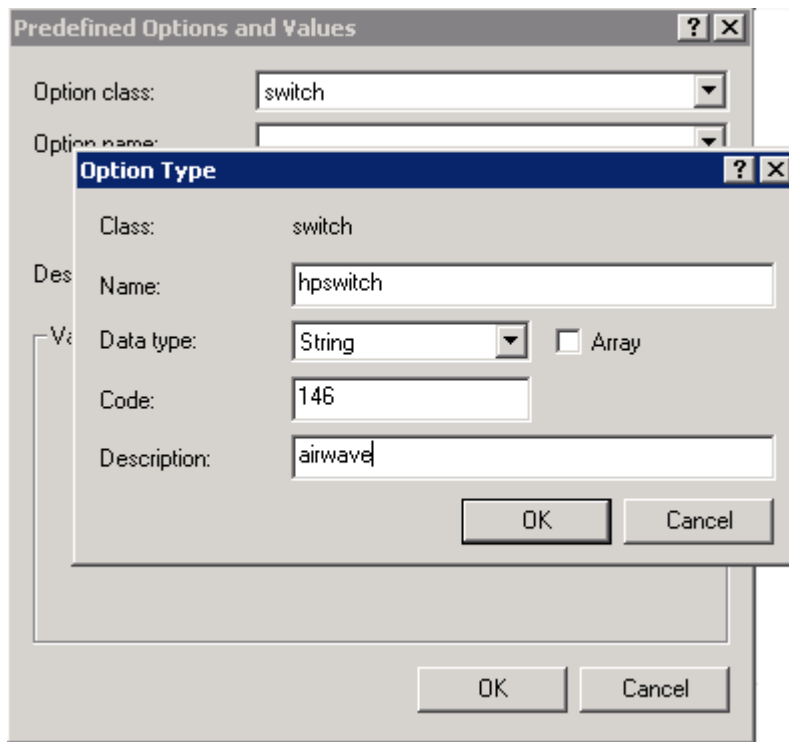
- Click **OK**.
- Right click on **IPv4** and select **Set Predefined Options....**



9. From the Predefined Options and Values window, select **Option class**. The Option Class displayed is the one that you configured under **DHCP Vendor Class**. In this example, the Option Class is **switch**.




10. Click **Add...**
11. From the Option Type window, enter the desired **Class** (any), the **Data type** (select **string**), the **Code** (enter **146**), and the **Description** (any).

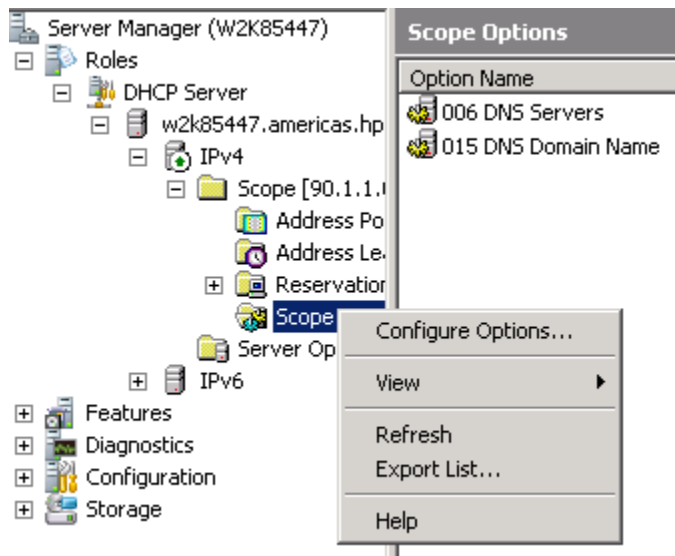


12. Click **OK**.
13. Under the Predefined Options and Values window, enter the Value String. In this example, we enter **hpeSwitch:hp2920,90.1.1.10,admin**. The String has the following format: <Group>:<Topfolder>,<AMP IP>,<shared secret>


If you need to add sub-folders, use the following format:<Group>:<Topfolder>:<folder1>,<AMP IP>,<shared secret>



14. Click **OK**.
15. Under **IPv4**, expand **Scope**. Right click on **Scope Options** and select **Configure Options...**



16. From the Scope Options window:
 - a. Select the **Advanced** tab.
 - b. Under Vendor class, select the desired switch. In this example, **switch**.
 - c. Select the **146 hpswitch** option.
 - d. Click **OK**.



17. You can verify the AirWave details as follows:

```
switch# show amp-server
switch# show run
```

Zero Touch Provisioning

The Zero Touch Provisioning (ZTP) solution enables the auto-configuration of your switches on the first boot without requiring any administrator's intervention at the switch. The switches use DHCP server option configurations to support ZTP.



If the switch does not contain the minimal configuration set, ZTP will get disabled. See [Image Upgrade](#).

Auto-configuration using ZTP

ZTP auto-configures your switches as follows:

Procedure

1. The switch boots up with the factory default configuration.

2. The switch sends out a DHCP discovery from the primary VLAN interface.
 - a. The preferred configuration method uses DHCP option 43 value as a string to parse Airwave configuration. Switch would expect a DHCP option 60 with value `ArubaInstantAP` along with DHCP option 43 to parse Airwave details
 - b. The alternate configuration method supports both encapsulated values from option 43 and direct value from option 43. Encapsulated vendor-specific sub options, with sub-option code 146 is for Airwave details.
3. After the AirWave details are verified and configured, the switch initiates the check-in into the AirWave server using the HTTPS communication.



The AirWave configuration must be in the following format:

```
<Group>:<Topfolder>:<folder1>,<AMP IP >,<shared secret>
```

4. After a successful registration, AirWave can monitor, configure, and troubleshoot the switches. Refer to *Aruba Networks and AirWave Switch Configuration Guide*.
5. Check-in failure retry is done every 60 seconds for 10 retries.
6. If the DHCP options are not configured for AirWave, the switch is left in its default state for manual configuration.

Disabling ZTP

Zero touch provisioning is disabled if you make any of the following changes to the switch's configuration:

- Enter the switch configuration mode using the `configure terminal` command.
- Enter into Menu and exit without doing any configuration
- Make any successful configuration that changes the running-configuration of the switch using a CLI, SNMP, REST APIs, menu interface, or the web GUI.
- If you upgrade with non-minimal configuration set from any 15.xx version to version 16.01, see [Image Upgrade](#)

Image Upgrade

If you upgrade from any 15.xx version to version 16.01, the following minimal set of configuration is validated to enable or disable the ZTP process:

- If the switch has any other VLAN apart from the default VLAN, ZTP gets disabled.
- In default VLAN, if the IPv4 address is not set as DHCP (default option is DHCP), ZTP gets disabled.
- In default VLAN, if IPv6 is enabled or configured, ZTP gets disabled.

If you have any other configuration during the upgrade, ZTP will be in the enabled state only.

Configure a switch using the CLI

Use the `amp-server` command to configure the AirWave IP address, group, folder, and shared secret. You must have the `manager` role to execute this command.

For example:

```
HP switch(config)# amp-server ip 172.16.185.23 group 2530 folder 2530 secret secret
```

The `show amp-server` command shows the configuration details:

```
switch# show amp-server
Airwave Configuration details
  AMP Server IP       : 172.16.185.23
  AMP Server Group    : 2530
```

```
AMP Server Folder      : 2530
AMP Server Secret     : secret
AMP Server Config status: Configured
```

Stacking and chassis switches

The ZTP and AirWave interaction for stacked switches is similar to the one for the standalone switch, with the exception that only the commander in the stack processes the ZTP and AirWave interaction.

Stacking supports the following features:

- Backplane Stacking (BPS) running on:
 - HPE 3800 Switch Series
 - HPE Aruba 2920 Switch Series
 - HPE Aruba 3810M Series
- Virtual Switching Framework (VSF) running on HPE Aruba 5400R Switch Series v3 modules
- Chassis running on HPE Aruba 5400R Switch Series v3 modules

Troubleshooting

Cause

You can troubleshoot switches by using the SSH connection and the device logs available in AirWave. For a list of all RMON message, refer to *HPE ArubaOS-Switch Event Log Message Reference Guide*.

You can enable the debug logging with the `debug ztp` command, see [debug ztp](#).

View AMP server messages

To display the AMP server debug messages, enter:

```
switch# debug ztp
```

To print the debug messages to the terminal, enter:

```
switch# debug destination session
```

Validation Rules

Validation	Error/Warning
Invalid AirWave IP address	Invalid input: 300.300.300.300
Group name exceeds max length	String %s too long. Allowed length is 32 characters.
Folder name exceeds max length	String %s too long. Allowed length is 128 characters.
Secret name exceeds max length	String %s too long. Allowed length is 32 characters.
AirWave IP address or Group or folder or secret is not configured.	Incomplete input: amp-server

View configuration details

to view the AirWave configuration details, use the `show amp-server` command. For example:

Airwave Configuration details

```
AMP Server IP      : 192.168.1.1
AMP Server Group   : HP_GROUP
AMP Server Folder  : folder
AMP Server Secret  : secret123
AMP Server Config Status: Configured
```

The `show amp-server` command displays the following values for the above configuration details. The `show running` command also displays the AirWave configuration details.

For example, to show details of the running configuration:

```
switch# show running-config
hostname "Aruba-2930F-24G"
module 1 type j9726a
snmp-server community "public" unrestricted
oobm
  ip address dhcp-bootp
  exit
vlan 1
  name "DEFAULT_VLAN"
  untagged 1-24
  ip address dhcp-bootp
  exit
amp-server ip 192.168.1.1 group "group" folder "folder" secret "secret123"
```

amp-server

Syntax

```
[no] amp-server ip <IP ADDRESS> group <GROUP> folder <FOLDER> secret <SECRET>
```

Description

The `amp-server` command configures the AirWave Management Platform (AMP) IP address, group, folder, and shared secret and triggers the device registration with AMP.

Only the `manager` role can execute this command.

Parameters

`ip`

AMP server IP address.

`group`

AMP server group name.

`folder`

AMP server folder name.

`secret`

AMP server shared secret string.

Options

no

The `no amp-server` command removes the configuration for the AMP server.

Permissions

Only the `manager` role can execute this command.

debug ztp

Syntax

```
[no] debug ztp
```

Description

Enables or disables ZTP debug logging.

Parameters

ztp

Zero Touch Provisioning.

Options

no

The `no debug ztp` command disables the ZTP debug logging.

Auto device detection and configuration

The auto device detection and configuration detects a directly connected Aruba AP dynamically and applies predefined configurations to ports on which the Aruba AP is detected.

You can create port configuration profiles, associate them to a device type, and enable or disable a device type. The only device type supported is `aruba-ap` and it is used to identify all the Aruba APs.

When a configured device type is connected on a port, the system automatically applies the corresponding port profile. Connected devices are identified using LLDP. When the LLDP information on the port ages out, the device profile is removed.

By default, the device profile feature is disabled. When you enable the device profile support for a device type, if no other device profile is mapped to the device type, the default device profile `default-ap-profile` is associated with the device type. You can modify the AP default device profile configuration but you cannot delete it. The `default-ap-profile` command supports only the AP device type.

Requirements

- Only APs directly connected to the switch will be detected.

Limitations

- Only one device type is supported, `aruba-ap`, and it is used to identify all the Aruba APs.
- You can modify the configuration parameters of the default profile, `default-ap-profile`, but you cannot delete it or change its name.
- The maximum value for `poe-max-power` is 33 W.
- If the port was part of any protocol VLANs prior to the device profile application, those VLANs will not be removed while applying the device profile.
- Egress rate limiting is not supported for devices running on:
 - Aruba 2530 Switch Series
 - Aruba 2620 Switch Series
- The `egress-bandwidth` is only supported for devices running on:
 - Aruba 2920 Switch Series
 - Aruba 2930F Switch Series
 - Aruba 5400R zl2 Switch Series v2 & v3 modules
 - Aruba 3800 Switch Series
- The `egress-bandwidth` option is not supported and not displayed in the CLI running on:
 - Aruba 2530 Switch Series
 - Aruba 2620 Switch Series

Feature Interactions

Profile Manager and 802.1X

Profile Manager interoperates with RADIUS when it is working in the client mode. When a port is blocked due to 802.1X authentication failure, the LLDP packets cannot come in on that port. Therefore, the Aruba AP cannot be

detected and the device profile cannot be applied. When the port gets authenticated, the LLDP packets comes in, the AP is detected, and the device profile is applied.

You must ensure that the RADIUS server will not supply additional configuration such as VLAN or CoS during the 802.1X authentication as they will conflict with the configuration applied by the Profile Manager. If the RADIUS server supplies any such configurations to a port, the device profile will not be applied on such ports.

Profile Manager and LMA/WMA/MAC-AUTH

If either LMA, WMA, or MAC-AUTH is enabled on an interface, all the MAC addresses reaching the port must be authenticated. If LMA, WMA, or MAC-AUTH is configured on an interface, the user can have more granular control and does not need the device profile configuration. Therefore, the device profile will not be applied on such interface.

Profile manager and Private VLANs

When the device profile is applied, a check is performed to verify if the VLAN addition violates any PVLAN requirements. The following PVLAN related checks are done before applying the VLANs configured in the device profile to an interface:

- A port can be a member of only one VLAN from a given PVLAN instance.
- A promiscuous port cannot be a member of a secondary VLAN.

Creating a profile and associate a device type

Procedure

1. Create a new profile:

```
switch# device-profile <profile-name>
```

2. Enable the `aruba-ap` device type:

```
switch# device-profile type aruba-ap enable
```

3. Associate the new profile to the `aruba-ap` device type:

```
switch# device-profile type aruba-ap associate <profile-name>
```

For example, to add the profile `abc` and associate it with the `aruba-ap` type, enter:.

```
switch# device-profile name abc
switch# device-profile type aruba-ap enable
switch# device-profile type aruba-ap associate abc
```

device-profile name

Syntax

```
[no] device-profile name <PROFILE-NAME> [untagged-vlan <VLAN-ID> |
tagged-vlan <VLAN-LIST> |
cos <COS-VALUE> |
ingress-bandwidth <Percentage> |
egress-bandwidth <Percentage> |
{poe-priority {critical | high | low} |
speed-duplex {auto | auto-10 | auto-100 | ...} |
poe-max-power <Watts>]
```

Description

This command is used to create an user-defined profile. A profile is a named collection of port settings applied as a group. You can modify the default profile, `default-ap-profile`, but you cannot delete it. You can create four additional profiles.

The `default-ap-profile` has the following values:

- `untagged-vlan: 1`
- `tagged-vlan: None`
- `ingress-bandwidth: 100`
- `egress-bandwidth: 100`
- `cos: 0`
- `speed-duplex: auto`
- `poe-max-power: 33`
- `poe-priority: critical`

You can modify these parameters. For example, you can execute `no untagged-vlan` to create a device profile with tagged only ports.

Parameters

`name`

Specifies the name of the profile to be configured. The profile names can be at most 32 characters long.

`cos`

The Class of Service (CoS) priority for traffic from the device.

`untagged-vlan`

The port is an untagged member of specified VLAN.

`tagged-vlan`

The port is a tagged member of the specified VLANs.

`ingress-bandwidth`

The ingress maximum bandwidth for the device port.

`egress-bandwidth`

The egress maximum bandwidth for the device port.

`poe-priority`

The PoE priority for the device port.

`speed-duplex`

The speed and duplex for the device port.

`poe-max-power`

The maximum PoE power for the device port.

Options

`no`

Removes the user-defined profiles.

Restrictions

- You can modify the configuration parameters of the default profile, `default-ap-profile`, but you cannot delete it or change its name.
- For Aruba 5400R Switch Series and Aruba 2930F Switch Series devices, the maximum value for `poe-max-power` is 30 W. For all other devices, the maximum value for `poe-max-power` is 33 W.

- Egress rate limiting is not supported for devices running on:
 - Aruba 2530 Switch Series
 - Aruba 2540 Switch Series
 - Aruba 2620 Switch Series
 - Aruba 2930F Switch Series
- The `egress-bandwidth` is only supported for Aruba 2920 Switch Series, Aruba 5400R Switch Series v2 & v3 modules, and Aruba 3800 Switch Series.
- The `egress-bandwidth` option is not supported and not displayed in the CLI for devices on: Aruba 2530 Switch Series and Aruba 2620 Switch Series.
- The profile configuration is only applicable to access points.

device-profile type

Syntax

```
device-profile type <DEVICE> [associate <PROFILE-NAME> | enable | disable ]
```

Description

This command specifies an approved device type in order to configure and attach a profile to it. The profile's configuration is applied to any port where a device of this type is connected.

Parameters

`type`

An approved device type in order to configure and attach a profile to it. The only device type supported is `aruba-ap` and it is used to identify all the Aruba APs.

APs.

`associate`

Associates a profile with a device type.

`enable`

Enables automatic profile association.

`disable`

Disables automatic profile association.

Options

`no`

Removes the device type association and disables the feature for the device type. By default, this feature is disabled.

Restrictions

Only one device type is supported, `aruba-ap`, and it is used to identify all the Aruba access points.

Rogue AP Isolation

The Rogue AP Isolation feature detects and blocks any unauthorized APs in the network. You can either log or block the rogue device. If the action requested is to log the rogue device, the MAC address of the rogue device is logged in the system logs (RMON). If the action is to block the rogue device, the traffic to and from the MAC address of the rogue device is blocked. The MAC is also logged in the system log.

When an Aruba AP detects a rogue AP on the network, it sends out the MAC address of the AP as well as the MAC of the clients connected to the AP to the switch using the ArubaOS-Switch proprietary LLDP TLV protocol.

The switch then adds a rule in its hardware table to block all the traffic originating from the rogue AP's MAC address.

The `rogue-ap-isolation` command configures the rogue AP isolation for the switch and gives the option to enable or disable the rogue AP isolation feature. The `rogue-ap-isolation action` command gives you the ability to block the traffic to or from the rogue device or log the MAC of the rogue device. When the action is set to block, the rogue MAC is logged as well. By default, the action is set to block.

The `rogue-ap-isolation whitelist` command lets you add devices detected as possible rogue APs to the whitelist. A maximum of 128 MAC addresses are supported for the whitelist.

The `clear rogue-aps` command clears the detected rogue AP device MAC address.

Limitations

- You can add a maximum of 128 MAC addresses to the whitelist.
- When a MAC is already authorized by any of the port security features such as LMA, WMA, or 802.1X, the MAC is logged but you cannot block it using the `rogue-ap-isolation` feature. A RMON event is logged to notify the user.
- When a MAC is already configured as an IP received MAC of a VLAN interface, the MAC is logged but you cannot block it by using the `rogue-ap-isolation` feature. A RMON event is logged to notify the user.
- When a MAC is already locked out via `lockout-mac` or locked down using the `static-mac` configuration, the MAC is logged but you cannot block it using the `rogue-ap-isolation` feature. A RMON event is logged to notify the user.
- The number of rogue MACs supported on a switch is a function of the value of `max-vlans` at boot time. Since the resources are shared with the `lockout-mac` feature, the scale is dependent on how many lockout addresses have been configured on the switch using the `lockout-mac` feature. The following table lists the scale when there are no lockout addresses configured on the switch:

Max VLAN	Supported MACs
0 < VLAN <= 8	200
8 < VLAN <= 16	100
16 < VLAN <= 256	64
256 < VLAN <= 1024	16
1024 < VLAN <= 2048	8
2048 < VLAN <= 4094	4

The switch will create an RMON log entry and the rogue MAC will be ignored when the limit is reached.



If the `max-vlans` value is changed to a different value, the scale of rogue MACs supported will not change until the next reboot.

Feature Interactions

MAC lockout and lockdown

The Rogue AP isolation feature uses the MAC lockout feature to block MACs in hardware. Therefore, any MAC blocked with the Rogue AP isolation feature cannot be added with the `lockout-mac` or `static-mac` command if the action type is set to `block`.

For example:

```
switch# lockout-mac 247703-7a8950
Cannot add the entry for the MAC address 247703-7a8950 because it is already
blocked by rogue-ap-isolation.
```

```
switch# static-mac 247703-7a8950 vlan 1 interface 1
Cannot add the entry for the MAC address 247703-7a8950 because it is already
blocked by rogue-ap-isolation.
```

Similarly, any MAC that was added with the `lockout-mac` or `static-mac` command and that is being detected as rogue will be logged, but not blocked in hardware as it already is set to block. If the MAC is removed from `lockout-mac` or `static-mac` but is still in the rogue device list, it will be blocked back in hardware if the action type is `block`.

LMA/WMA/802.1X/Port-Security

Any configuration using LMA, WMA, 802.1X, or Port-Security will not be blocked if the Rogue AP isolation feature is enabled. All these features act only when a packet with the said MAC is received on a port.

If `rogue-ap-isolation` blocks a MAC before it is configured to be authorized, packets from such MACs will be dropped until one of the following happens:

- Rogue action is changed to LOG.
- Rogue-AP isolation feature is disabled.
- The MAC is not detected as rogue anymore.
- LLDP is disabled on the port (or globally).

Once a MAC has been authorized by one of these features, it will not be blocked by Rogue AP isolation. A RMON will be logged to indicate the failure to block.

The Rogue AP module will retry to block any such MACs periodically. In the event of the MAC no longer being authorized, Rogue AP isolation will block the MAC again. No RMON is logged to indicate this event.

L3 MAC

The Rogue AP isolation feature will not block a MAC configured as an IP receive MAC address on a VLAN interface. This event will be logged in RMON if such MACs are detected as rogue.

Conversely, any MAC already blocked by Rogue AP isolation will not be allowed to be configured as an IP receive MAC address of a VLAN interface.

For example:

```
switch# vlan 1 ip-recv-mac-address 247703-3effbb
Cannot add an entry for the MAC address 247703-3effbb because it is already
blocked by rogue-ap-isolation.
```

Using the Rogue AP Isolation feature

Procedure

1. Check the feature state:

```
switch# show rogue-ap-isolation

Rogue AP Isolation
```

```
Rogue AP Status : Disabled
Rogue AP Action : Block

Rogue MAC Address Neighbour MAC Address
-----
```

2. Enable the feature:

```
switch# rogue-ap-isolation enable
switch# show rogue-ap-isolation

Rogue AP Isolation

Rogue AP Status : Enabled
Rogue AP Action : Block

Rogue MAC Address Neighbour MAC Address
-----
```

3. Change the action type from block to log:

```
switch# rogue-ap-isolation action log
switch# show rogue-ap-isolation

Rogue AP Isolation

Rogue AP Status : Enabled
Rogue AP Action : Log

Rogue MAC Address Neighbour MAC Address
-----
```

4. List the current whitelist entries:

```
switch# show rogue-ap-isolation whitelist

Rogue AP Whitelist Configuration

Rogue AP MAC
-----
```

5. Add a new whitelist entry:

```
switch# rogue-ap-isolation whitelist 005056-00326a
switch# show rogue-ap-isolation whitelist

Rogue AP Whitelist Configuration

Rogue AP MAC
-----
00:50:56:00:32:6a
```

rogue-ap-isolation

syntax

```
rogue-ap-isolation {enable | disable}
```

Description

Configures the rogue AP isolation for the switch.

Parameters

enable

Enables the rogue AP isolation.

disable

Disables the rogue AP isolation.

rogue-ap-isolation action

syntax

```
rogue-ap-isolation action {log | block}
```

Description

Configures the action to take for the rogue AP packets. This function is disabled by default.

Parameters

action

Configure the action to take for rogue AP packets. By default, the rogue AP packets are blocked.

Options

log

Logs traffic to or from any rogue access points.

block

Blocks and logs traffic to or from any rogue access points.

rogue-ap-isolation whitelist

syntax

```
[no] rogue-ap-isolation whitelist <MAC-ADDRESS>
```

Description

Configures the rogue AP Whitelist MAC addresses for the switch. Use this command to add to the whitelist the MAC addresses of approved access points or MAC addresses of clients connected to the rogue access points. These approved access points will not be added to the rogue AP list even if they are reported as rogue devices.

Parameters

MAC-ADDRESS

Specifies the MAC address of the device to be moved from the rogue AP list to the whitelist.

Options

no

Removes the MAC address individually by specifying the MAC.

Restrictions

You can add a maximum of 128 MAC addresses to the whitelist.

clear rogue-ap-isolation

syntax

```
clear rogue-ap-isolation { <MAC-ADDRESS> | all }
```

Description

Removes the MAC addresses from the rogue AP list.

Parameters

MAC-ADDRESS

Specifies the MAC address of the device to be moved from the rogue AP list.

all

Clears all MAC addresses from the rogue AP list.

Restrictions

The MAC addresses cleared using this option will be added back to the rogue list under the following cases:

1. The LLDP administrator status of the port on which the AP that reported the MAC is disabled and enabled back.
2. The data that is in the rogue AP TLV sent from the AP that informed the rogue MAC has changed.
3. To permanently ignore a MAC from being detected as rogue, add it to the whitelist.

Troubleshooting

Dynamic configuration not displayed when using “show running-config”

Symptom

The `show running-config` command does not display the dynamic configuration applied through the device profile.

Cause

The `show running-config` command shows only the permanent user configuration and parameters configured through device profile.

Action

Procedure

1. Use the specific `show device-profile` command to display the parameters dynamically configured through the device profile.

Switch does not detect the rogue AP TLVs

Symptom

The switch does not detect the rogue AP TLVs that could be sent from the neighboring device.

Cause

The LLDP administrator status of a port is moved from `txOnly` to `tx_rx` or `rx_only` within 120 seconds of the previous state change to `txOnly`.

Action

Procedure

1. Wait for 120 seconds before moving from the state `txOnly` to the state `tx_rx` or `rx_only`.

2. Move the administrator status to `disable` and then back to `tx_rx` or `rx_only`.

The `show run` command displays non-numerical value for `untagged-vlan`

Symptom

The `show run` command displays one of the following values for `untagged-vlan`:

- `no untagged-vlan`
- `untagged-vlan : None`

Cause

The `no device-profile` or the `no rogue-ap-isolation whitelist` command is executed to configure `untagged-vlan` to 0.

Action

Procedure

1. No action is required.

Show commands

Use the following show commands to view the various configurations and status.

Command	Description
<code>show device-profile</code>	Shows the device profile configuration and status.
<code>show device-profile config</code>	Shows the device profile configuration details for a single profile or all profiles.
<code>show device-profile status</code>	Shows currently applied device profiles.
<code>show rogue-ap-isolation</code>	Shows the following information: <ul style="list-style-type: none">• The status of the feature: enabled or disabled.• The current action type for the rogue MACs detected.• The list of MAC addresses detected as rogue and the MAC address of the AP that reported them.
<code>show rogue-ap-isolation whitelist</code>	Shows the rogue AP whitelist configuration.
<code>show run</code>	Shows the running configuration.

Validation Rules

Validation	Error/Warning/Prompt
device-profile profile-name default-ap-profile	Maximum tagged VLANs that can be associated with a device-profile is 256.
device-profile profile-name creation.	String too long. Allowed length is 32 characters.
device-profile profile-name creation.	Device profile <> already exists.
device-profile profile-name creation.	The maximum number of device profiles allowed is 5.
device-profile profile-name deletion.	Device profile <> does not exist.
device-profile profile-name deletion.	Cannot delete profile <> when associated with a device type.
device-profile profile-name deletion.	Default profile cannot be deleted.
device-profile profile-name modification via SNMP.	Default profile name cannot be changed.
device-profile profile-name creation/modification via SNMP.	Device profile index cannot be greater than 5.
untagged-vlan	Invalid VLAN.
untagged-vlan	Cannot configure the VLAN <> as an untagged VLAN because this is already used as a tagged VLAN.
tagged-vlan 1-1000	The maximum number of tagged VLANs in a profile is less than 512 or the maximum VLANs, MAX_VLANs, configurable in the system.
tagged-vlan	Cannot configure the VLAN <> as a tagged VLAN because this is already used as an untagged VLAN.
ingress-bandwidth	SNMP should return WRONG_VALUE_ERROR.
egress-bandwidth	SNMP should return WRONG_VALUE_ERROR.
cos	SNMP should return WRONG_VALUE_ERROR.
speed-duplex	SNMP should return WRONG_VALUE_ERROR.
poe-max-power	SNMP should return WRONG_VALUE_ERROR.
poe-priority	SNMP should return WRONG_VALUE_ERROR.
device-profile type aruba-ap profile-name	String <> too long. Allowed length is 32 characters.

Table Continued

Validation	Error/Warning/Prompt
device-profile type aruba-ap profile-name	Device profile <> does not exist.
device-profile type aruba-switch-router	Device type is not supported.
rogue-ap-whitelist	Whitelist MAC address already exists in the list.
rogue-ap-whitelist	Whitelist MAC address does not exist in the list.
rogue-ap-whitelist	The maximum number of whitelist MACs allowed is 128.
rogue-ap-whitelist <MAC>	Cannot add the whitelist entry because the specified MAC address is already configured as a lock-out MAC.
lock-out <MAC>	Cannot add the lock-out entry because the specified MAC address is already configured as a whitelist MAC.
lockout-mac <MAC-ADDRESS> OR static-mac <MAC-ADDRESS> vlan <vlan-id> interface <interface> OR vlan <vlan-id> ip-recv-mac-address <MAC-ADDRESS	Cannot add an entry for the MAC address <MAC-ADDRESS> because it is already blocked by rogue-ap-isolation.

LACP-MAD commands

Configuration command

The following command defines whether LACP is enabled on a port, and whether it is in active or passive mode when enabled. When LACP is enabled and active, the port sends LACP packets and listens to them. When LACP is enabled and passive, the port sends LACP packets only if it is spoken to. When LACP is disabled, the port ignores LACP packets. If the command is issued without a mode parameter, 'active' is assumed. During dynamic link aggregation using LACP, ports with the same key are aggregated as a single trunk. MAD passthrough applies only to trunks and not to physical ports.

```
switch# [no] interface <port-list> lacp [mad-passthrough <enable|disable>|active|passive|key <key>]
```

show commands

LACP-MAD supports the following show commands:

- show LACP-MAD passthrough configuration on LACP trunks

```
switch# show lacp [counters <port-list>] | local <port-list> | peer <port-list> | distributed | mad-passthrough [counters <port-list>]]
```

- show LACP-MAD passthrough counters on ports

```
switch# show lacp mad-passthrough counters <port-list>
```

clear command

Clear all LACP statistics including MAD passthrough counters. Resets LACP packets sent and received on all ports.

```
switch# clear lacp statistics
```

LACP-MAD overview

Link Aggregation Control Protocol-Multi-Active Detection (LACP-MAD) is a detection mechanism deployed by switches to recover from a breakup of the Vertical Switching Framework (VSF) stack due to link or other failure.

LACP-MAD is implemented by sending extended LACP data units (LACPDUs) with a type length value (TLV) that conveys the active ID of an VSF virtual device. The active ID is identical to the member ID of the master and is thus unique to the VSF virtual device. When LACP MAD detection is enabled, the members exchange their active IDs by sending extended LACPDUs.

- When the VSF virtual device operates normally, the active IDs in the extended LACPDUs sent by all members are the same, indicating that there is no multi-active collision.
- When there is a breakup in the VSF stack, the active IDs in the extended LACPDUs sent by the members in different VSF virtual devices are different, indicating that there are multi-active collisions.

LACP-MAD passthrough helps VSF-capable devices detect multi-access and take corrective action. These devices do not initiate transmission of LACP-MAD frames or participate in any MAD decision making process. These devices simply forward LACP-MAD TLVs received on one interface to the other interfaces on the trunk. LACP-MAD passthrough can be enabled for 24 LACP trunks. By default, LACP-MAD passthrough is disabled.

The following table lists the switch scalability values for the areas of VLANs, ACLs, hardware, ARP, and routing.

Subject	Maximum
IPv4 ACLs	
total named (extended or standard)	Up to 2048 (minus any IPv4 numeric standard or extended ACL assignments and any RADIUS-assigned ACLs) ¹
total numbered standard	Up to 99 ¹
total numbered extended	Up to 100 ¹
total ACEs in all IPv4 ACLs	Up to 3072 ¹
Layer-3	
VLANs with at least one IP Address	512
IP addresses per system	2048 IPv4 2048 IPv6 ²
IP addresses per VLAN	32 ³
Static routes (IPv4 and IPv6 combined)	256
IPv4 host hardware table	72 K (8K internal, 64K external)
IPv4 BMP hardware table	2 K
ARP	
ARP entries	25,000
Packets held for ARP resolution	25
Dynamic Routing	
Total routes supported	IPv4 only: 10,000 (including ARP) IPv4 and IPv6: 10 K (IPv4) and 3 K (IPv6) ⁴ IPv6 only: 5 K ⁵
IPv4 Routing Protocol	
RIP interfaces	128

Table Continued

IPv6 Routing Protocol

DHCPv6 Helper Addresses 32 unique addresses; multiple instances of same address counts as 1 towards maximum

¹

² These limits apply only to user-configured addresses and not to auto-configured link local and prefix IPv6 addresses. A maximum configuration could support up to 2048 user-configured and 2048 auto-configured IPv6 addresses for a total of 4096.

³ There can be up to 32 IPv4 and 32 user-configured IPv6 addresses on a single VLAN. In addition, each VLAN is limited to 3 auto-configured prefix-based IPv6 addresses.

⁴ Configured as an ABR for OSPF with four IPv4 areas and four IPv6 areas.

⁵ Configured as an ABR for OSPF with two IPv6 OSPF areas.

Overview

The switches support several methods for transferring files to and from a physically connected device, or via the network, including TFTP and Xmodem. This appendix explains how to download new switch software, upload or download switch configuration files and software images, and upload command files for configuring ACLs.

Downloading switch software

HPE Switch periodically provides switch software updates through the Switch Networking website. For more information, see the support and warranty booklet shipped with the switch, or visit <http://www.hpe.com/networking> and click on **software updates**.



This manual uses the terms **switch software** and **software image** to refer to the downloadable software files the switch uses to operate its networking features. Other terms sometimes include **Operating System**, or **OS**.

General software download rules

- Switch software that you download via the menu interface always goes to primary flash.
- After a software download, you must reboot the switch to implement the new software. Until a reboot occurs, the switch continues to run on the software it was using before the download.



Downloading new switch software does not change the current switch configuration. The switch configuration is contained in separate files that can also be transferred. See **Transferring switch configurations** on page 261.

In most cases, if a power failure or other cause interrupts a flash image download, the switch reboots with the image previously stored in primary flash. In the unlikely event that the primary image is corrupted (which may occur if a download is interrupted by a power failure), the switch goes into boot ROM mode. In this case, use the boot ROM console to download a new image to primary flash.

Using TFTP to download software from a server

This procedure assumes that:

- A software version for the switch has been stored on a TFTP server accessible to the switch. (The software file is typically available from the HPE Switch Networking website at <http://www.hpe.com/networking>.)
- The switch is properly connected to your network and has already been configured with a compatible IP address and subnet mask.
- The TFTP server is accessible to the switch via IP.

Before you use the procedure, do the following:

- Obtain the IP address of the TFTP server in which the software file has been stored.
- If VLANs are configured on the switch, determine the name of the VLAN in which the TFTP server is operating.
- Determine the name of the software file stored in the TFTP server for the switch (For example, E0820.swi).



If your TFTP server is a UNIX workstation, ensure that the case (upper or lower) that you specify for the filename is the same case as the characters in the software filenames on the server.

Downloading from a server to primary flash using TFTP (Menu)

Note that the menu interface accesses only the primary flash.

Procedure

1. In the console Main Menu, select **Download OS** to display the screen in **Figure 37: Example: of a download OS (software) screen (default values)** on page 248. (The term "OS" or "operating system" refers to the switch software):

Figure 37: Example: of a download OS (software) screen (default values)

```
----- CONSOLE - MANAGER MODE -----
                          Download OS

Current Firmware revision : K.11.00

Method [TFTP] : TFTP
TFTP Server :

Remote File Name :

Actions->  C   a   n   c   e   l       E   d   i   t       e   x   e   c   u   t   e       H   e   l   p

Select the file transfer method (TFTP and XMODEM are currently supported).
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

2. Press **[E]** (for **Edit**).
3. Ensure that the **Method** field is set to **TFTP** (the default).
4. In the **TFTP Server** field, enter the IP address of the TFTP server in which the software file has been stored.
5. In the **Remote File Name** field, enter the name of the software file (if you are using a UNIX system, remember that the filename is case-sensitive).
6. Press **[Enter]**, then **[X]** (for **eXecute**) to begin the software download.

The following screen appears:

Figure 38: Example: of the download OS (software) screen during a download

```
----- CONSOLE - MANAGER MODE -----
                          Download OS

Current Firmware revision : E.08.00
Method [TFTP] : TFTP
TFTP Server : 10.28.227.105

Remote File Name : K.11.00.swi

                          Received 370,000 bytes of OS download.
+-----+
|*****|
+-----+
```

A "progress" bar indicates the progress of the download. When the entire software file has been received, all activity on the switch halts and you will see **Validating and writing system software to FLASH...**

7. After the primary flash memory is updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch**).

You will see this prompt:

```
Continue reboot of system? : No
```

Press the space bar once to change **No** to **Yes**, then press **[Enter]** to begin the reboot.



When you use the menu interface to download a switch software, the new image is always stored in primary flash. Also, using the `Reboot Switch` command in the Main Menu always reboots the switch from primary flash. Rebooting the switch from the CLI provides more options. See "Rebooting the Switch" in the basic operation guide.

8. After you reboot the switch, confirm that the software downloaded correctly:

- a. From the Main Menu, select
 2. **Switch Configuration...**
 2. **Port/Trunk Settings**
- b. Check the **Firmware revision** line.


For troubleshooting information on download failures, see [Troubleshooting TFTP download failures](#) on page 249.

Troubleshooting TFTP download failures

Cause

When using the menu interface, if a TFTP download fails, the Download OS (Operating System, or software) screen indicates the failure as seen in the following figure.

Figure 39: Example: of message for download failure



Some of the causes of download failures include:

- Incorrect or unreachable address specified for the **TFTP Server** parameter. This may include network problems.
- Incorrect VLAN.
- Incorrect name specified for the **Remote File Name** parameter, or the specified file cannot be found on the TFTP server. This can also occur if the TFTP server is a UNIX machine and the case (upper or lower) for the filename on the server does not match the case for the filename entered for the **Remote File Name** parameter in the **Download OS** (Operating System, or software) screen.
- One or more of the switch's IP configuration parameters are incorrect.

- For a UNIX TFTP server, the file permissions for the software file do not allow the file to be copied.
- Another console session (through either a direct connection to a terminal device or through Telnet) was already running when you started the session in which the download was attempted.

To find more information on the cause of a download failure:

- Examine the messages in the switch's Event Log by executing the `show log tftp` command from the CLI.
- For descriptions of individual Event Log messages, see the latest version of the event log message reference guide for your switch, available on the HPE Switch website. (See "Getting Documentation From the Web".)



If an error occurs in which normal switch operation cannot be restored, the switch automatically reboots itself, and an appropriate message is displayed after the reboot.

Downloading from a server to flash using TFTP (CLI)

Syntax:

```
copy tftp flash <ip-address> <remote-file> [<primary | secondary>]
```

Automatically downloads a switch software file to primary or secondary flash. If you do not specify the flash destination, the TFTP download defaults to primary flash.

Example:

To download a switch software file named k0800.swi from a TFTP server with the IP address of 10.28.227.103 to primary flash:

Procedure

1. Execute `copy` as shown below:

The command to download an OS (switch software)

```
switch# copy tftp flash 10.28.227.103 k0800.swi
The primary OS Image will be deleted, continue [y/n]? y 1
01431K 2
```

- ¹This message means that the image you want to upload will replace the image currently in primary flash.
- ²Dynamic counter continually displays the number of bytes transferred.

When the switch finishes downloading the software file from the server, it displays this progress message:

```
Validating and Writing System Software to FLASH ...
```

2. When the download finishes, you must reboot the switch to implement the newly downloaded software image. To do so, use one of the following commands:

Syntax:

```
boot system flash {<primary | secondary>}
```

Boots from the selected flash.

Syntax:

```
reload
```

Boots from the flash image and startup-config file. A switch covered in this guide (with multiple configuration files), also uses the current startup-config file.

For more information on these commands, see "Rebooting the Switch" in the basic operation guide for your switch.

3. To confirm that the software downloaded correctly, execute `show system` and check the **Firmware revision** line.

For information on primary and secondary flash memory and the boot commands, see "Using Primary and Secondary Flash Image Options" in the basic operation guide for your switch.



If you use `auto-tftp` to download a new image in a redundant management system, the active management module downloads the new image to both the active and standby modules. Rebooting after the `auto-tftp` process completes reboots the entire system.

Using SCP and SFTP

For some situations you may want to use a secure method to issue commands or copy files to the switch. By opening a secure, encrypted SSH session and enabling `ip ssh file transfer`, you can then use a third-party software application to take advantage of SCP and SFTP. SCP and SFTP provide a secure alternative to TFTP for transferring information that may be sensitive (like switch configuration files) to and from the switch. Essentially, you are creating a secure SSH tunnel as a way to transfer files with SFTP and SCP channels.

Once you have configured your switch to enable secure file transfers with SCP and SFTP, files can be copied to or from the switch in a secure (encrypted) environment and TFTP is no longer necessary.

To use these commands, you must install on the administrator workstation a third-party application software client that supports the SFTP and/or SCP functions. Some examples of software that supports SFTP and SCP are PuTTY, Open SSH, WinSCP, and SSH Secure Shell. Most of these are freeware and may be downloaded without cost or licensing from the internet. There are differences in the way these clients work, so be sure you also download the documentation.

As described earlier in this chapter you can use a TFTP client on the administrator workstation to update software images. This is a plain-text mechanism that connects to a standalone TFTP server or another HPE switch acting as a TFTP server to obtain the software image files. Using SCP and SFTP allows you to maintain your switches with greater security. You can also roll out new software images with automated scripts that make it easier to upgrade multiple switches simultaneously and securely.

SFTP is unrelated to FTP, although there are some functional similarities. Once you set up an SFTP session through an SSH tunnel, some of the commands are the same as FTP commands. Certain commands are not allowed by the SFTP server on the switch, such as those that create files or folders. If you try to issue commands such as `create` or `remove` using SFTP, the switch server returns an error message.

You can use SFTP just as you would TFTP to transfer files to and from the switch, but with SFTP, your file transfers are encrypted and require authentication, so they are more secure than they would be using TFTP. SFTP works only with SSH version 2 (SSH v2).



SFTP over SSH version 1 (SSH v1) is not supported. A request from either the client or the switch (or both) using SSH v1 generates an error message. The actual text of the error message differs, depending on the client software in use. Some examples are:

```
Protocol major versions differ: 2 vs. 1
Connection closed

Protocol major versions differ: 1 vs. 2
Connection closed

Received disconnect from <ip-addr> : /usr/local/libexec/
sftp-server: command not supported
Connection closed
```

SCP is an implementation of the BSD `rcp` (Berkeley UNIX remote copy) command tunneled through an SSH connection.

SCP is used to copy files to and from the switch when security is required. SCP works with both SSH v1 and SSH v2. Be aware that the most third-party software application clients that support SCP use SSHv1.

The general process for using SCP and SFTP involves three steps:

Procedure

1. Open an SSH tunnel between your computer and the switch if you have not already done so.
(This step assumes that you have already set up SSH on the switch.)
2. Execute `ip ssh filetransfer` to enable secure file transfer.
3. Use a third-party client application for SCP and SFTP commands.

Enabling SCP and SFTP

For more information about secure copy and SFTP, see [Using SCP and SFTP](#) on page 251.

Procedure

1. Open an SSH session as you normally would to establish a secure encrypted tunnel between your computer and the switch.

For more detailed directions on how to open an SSH session, see "Configuring secure shell (SSH)" in the access security guide for your switch. Please note that this is a one-time procedure for new switches or connections. If you have already done it once you should not need to do it a second time.
2. To enable secure file transfer on the switch (once you have an SSH session established between the switch and your computer), open a terminal window and enter the following command:

```
switch(config)# ip ssh filetransfer
```

For information on disabling TFTP and auto-TFTP, see [Disabling TFTP and auto-TFTP for enhanced security](#) on page 252.

Disabling TFTP and auto-TFTP for enhanced security

Using the `ip ssh filetransfer` command to enable SFTP automatically disables TFTP and auto-TFTP (if either or both are enabled), as shown in [Switch configuration with SFTP enabled](#) on page 252.

Switch configuration with SFTP enabled

```
switch(config)# ip ssh filetransfer
Tftp and auto-tftp have been disabled. 1
switch(config)# sho run
```

Running configuration:

```
; J9091A Configuration Editor; Created on release #xx.15.xx
```

```
hostname "HP Switch"
module 1 type J8702A
module 2 type J702A
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24,B1-B24
  ip address 10.28.234.176 255.255.240.0
  exit
ip ssh filetransfer 2
no tftp-enable
```

```
password manager
password operator
```

¹ Enabling SFTP automatically disables TFTP and auto-tftp and displays this message.

² Viewing the configuration shows that SFTP is enabled and TFTP is disabled.

If you enable SFTP and then later disable it, TFTP and auto-TFTP remain disabled unless they are explicitly re-enabled.

Operating rules are:

- The TFTP feature is enabled by default, and can be enabled or disabled through the CLI, the Menu interface (see **Figure 40: Using the Menu interface to disable TFTP** on page 253), or an SNMP application. Auto-TFTP is disabled by default and must be configured through the CLI.

Figure 40: Using the Menu interface to disable TFTP

```
----- CONSOLE - MANAGER MODE -----
                Switch Configuration - System Information

System Name : ProCurve
System Contact :
System Location :

Inactivity Timeout (min) [0] : 0      MAC Age Time (sec) [300] : 300
Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled
Tftp-enable [Yes] : Yes ←

Time Zone [0] : 0
Daylight Time Rule [None] : None

Actions->  Cancel   Edit   Save   Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Enables/Disables TFTP.
Note: If SFTP is enabled, this field will be set to **No**. You cannot use this field to enable TFTP if SFTP is enabled. Attempting to do so produces an **Inconsistent value** message in the banner below the **Actions** line.

- While SFTP is enabled, TFTP and auto-TFTP cannot be enabled from the CLI. Attempting to enable either non-secure TFTP option while SFTP is enabled produces one of the following messages in the CLI:

```
SFTP must be disabled before enabling tftp.
```

```
SFTP must be disabled before enabling auto-tftp.
```

Similarly, while SFTP is enabled, TFTP cannot be enabled using an SNMP management application. Attempting to do so generates an "inconsistent value" message. (An SNMP management application cannot be used to enable or disable auto-TFTP.)

- To enable SFTP by using an SNMP management application, you must first disable TFTP and, if configured, auto-TFTP on the switch. You can use either an SNMP application or the CLI to disable TFTP, but you must use the CLI to disable auto-TFTP.

The following CLI commands disable TFTP and auto-TFTP on the switch.

Syntax:

```
no tftp-enable
```

This command disables all TFTP operation on the switch **except** for the auto-TFTP feature. To re-enable TFTP operation, use the `tftp-enable` command. When TFTP is disabled, the instances of `tftp` in the CLI copy command and the Menu interface "Download OS" screen become unavailable.



This command does **not** disable auto-TFTP operation. To disable an auto-TFTP command configured on the switch, use the `no auto-tftp` command described below to remove the command entry from the switch's configuration.

Syntax:

```
no auto-tftp
```

If auto-TFTP is configured on the switch, this command deletes the `auto-tftp` entry from the switch configuration, thus preventing auto-tftp operation if the switch reboots.



This command does not affect the current TFTP-enable configuration on the switch.

Enabling SSH V2 (required for SFTP)

```
switch(config)# ip ssh version 2
```



As a matter of policy, administrators should **not** enable the SSH V1-only or the SSH V1-or-V2 advertisement modes. SSHv1 is supported on only some legacy switches (such as the HPE Switch Series 2500 switches).

Confirming that SSH is enabled

```
switch(config)# show ip ssh
```

Once you have confirmed that you have enabled an SSH session (with the `show ip ssh` command), enter `ip ssh filetransfer` so that SCP and/or SFTP can run. You can then open your third-party software client application to begin using the SCP or SFTP commands to safely transfer files or issue commands to the switch.



Any attempts to use SCP or SFTP without using `ip ssh filetransfer` cause the SCP or SFTP session to fail. Depending on the client software in use, you will receive an error message on the originating console, for Example:

```
IP file transfer not enabled on the switch
```

Disabling secure file transfer

```
switch(config)# no ip ssh filetransfer
```

Authentication

Switch memory allows up to ten public keys. This means the authentication and encryption keys you use for your third-party client SCP/SFTP software can differ from the keys you use for the SSH session, even though both SCP and SFTP use a secure SSH tunnel.



SSH authentication is mutually exclusive with RADIUS servers.

Some clients, such as PSCP (PuTTY SCP), automatically compare switch host keys for you. Other clients require you to manually copy and paste keys to the `$HOME/.ssh/known_hosts` file. Whatever SCP/SFTP software tool you use, after installing the client software you must verify that the switch host keys are available to the client.

Because the third-party software utilities you may use for SCP/SFTP vary, you should refer to the documentation provided with the utility you select before performing this process.

SCP/SFTP operating notes

- Any attempts to use SCP or SFTP without using `ip ssh filetransfer` will cause the SCP or SFTP session to fail. Depending on the client software in use, you will receive an error message on the originating console, for Example:

IP file transfer not enabled on the switch

- There is a delay when SFTP is copying an image onto the switch, and although the command prompt returns in a couple of seconds, the switch may take approximately a minute and half writing the image to flash. You can keep entering the `show flash` command to see when the copy is complete and the flash is updated. You can also check the log for an entry similar to the following:

I 01/09/13 16:17:07 00150 update: Primary Image updated.

I 01/09/13 16:13:22 00636 ssh: sftp session from 15.22.22.03

- When an SFTP client connects, the switch provides a file system displaying all of its available files and folders. No file or directory creation is permitted by the user. Files may be only uploaded or downloaded, according to the permissions mask. All of the necessary files the switch needs are already in place on the switch. You do not need to (nor can you) create new files.
- The switch supports one SFTP session or one SCP session at a time.
- All files have read-write permission. Several SFTP commands, such as `create` or `remove`, are not allowed and return an error message. The switch displays the following files:

```
/
+---cfg
|  running-config
|  startup-config
+---log
|  crash-data
|  crash-data-a
|  crash-data-b
|  crash-data-c
|  crash-data-d
|  crash-data-e          "      "
|  crash-data-f ""
|  crash-data-g
|  crash-data-h          "      "
|  crash-data-I ""
|  crash-data-J ""
|  crash-data-K ""
|  crash-data-L "      "
|  crash-log
|  crash-log-a
|  crash-log-b
|  crash-log-c
|  crash-log-d
|  crash-log-e""
|  crash-log-f""
|  crash-log-g
|  crash-log-h"  "
|  crash-log-I"  "
|  crash-log-J"  "
|  crash-log-K"  "
|  crash-log-L"  "
|  event log
+---os
|  primary
```

```

|   secondary
|   \---ssh
|       +---mgr_keys
|           |   authorized_keys
|           |   \---oper_keys
|           |       authorized_keys
|   \---core
|       port_1-24.cor   core-dump for ports 1-24 (stackable switches only)
|       port_25-48.cor core-dump for ports 25-48 (stackable switches only)

```

Once you have configured your switch for secure file transfers with SCP and SFTP, files can be copied to or from the switch in a secure (encrypted) environment and TFTP is no longer necessary.

Troubleshooting SSH, SFTP, and SCP operations

Cause

You can verify secure file transfer operations by checking the switch's event log, or by viewing the error messages sent by the switch that most SCP and SFTP clients print out on their console.



Messages that are sent by the switch to the client depend on the client software in use to display them on the user console.

Broken SSH connection

If an ssh connection is broken at the wrong moment (for instance, the link goes away or spanning tree brings down the link), a fatal exception occurs on the switch. If this happens, the switch gracefully exits the session and produces an Event Log message indicating the cause of failure. The following three examples show the error messages that may appear in the log, depending on the type of session that is running (SSH, SCP, or SFTP):

```

ssh: read error Bad file number, session aborted I 01/01/90
00:06:11 00636 ssh: sftp session from ::ffff:10.0.12.35 W
01/01/90 00:06:26 00641 ssh:

```

```

sftp read error Bad file number, session aborted I 01/01/90
00:09:54 00637 ssh: scp session from ::ffff:10.0.12.35 W 01/
01/90

```

```

ssh: scp read error Bad file number, session aborted

```



The `Bad file number` is from the system error value and may differ depending on the cause of the failure. In the third Example:, the device file to read was closed as the device read was about to occur.

Attempt to start a session during a flash write

If you attempt to start an SCP (or SFTP) session while a flash write is in progress, the switch does not allow the SCP or SFTP session to start. Depending on the client software in use, the following error message may appear on the client console:

```

Received disconnect from 10.0.12.31: 2: Flash access in
progress

```

```

lost connection

```

Failure to exit from a previous session

This next Example: shows the error message that may appear on the client console if a new SCP (or SFTP) session is started from a client before the previous client session has been closed (the switch requires approximately ten seconds to timeout the previous session):


```
Received disconnect from 10.0.12.31: 2: Wait for previous
session to complete
```

```
lost connection
```

Attempt to start a second session

The switch supports only one SFTP session or one SCP session at a time. If a second session is initiated (For example, an SFTP session is running and then an SCP session is attempted), the following error message may appear on the client console:

```
Received disconnect from 10.0.12.31: 2: Other SCP/SFTP
session running
```

```
lost connection
```

Using Xmodem to download switch software from a PC or UNIX workstation

This procedure assumes that:

- The switch is connected via the Console RS-232 port to a PC operating as a terminal. (For information on connecting a PC as a terminal and running the switch console interface, see the installation and getting started guide you received with the switch.)
- The switch software is stored on a disk drive in the PC.
- The terminal emulator you are using includes the Xmodem binary transfer feature. (For example, in the HyperTerminal application included with Windows NT, you would use the **Send File** option in the **Transfer** drop-down menu.)

Downloading to primary flash using Xmodem (Menu)



The menu interface accesses only the primary flash.

Procedure

1. From the console Main Menu, select
7. Download OS
2. Press **[E]** (for **Edit**).
3. Use the Space bar to select **XMODEM** in the **Method** field.
4. Press **[Enter]**, then **[X]** (for **eXecute**) to begin the software download.

The following message appears:

Press enter and then initiate Xmodem transfer from the attached computer....

5. Press **[Enter]** and then execute the terminal emulator commands to begin Xmodem binary transfer.

For example, using HyperTerminal:

- a. Click on **Transfer**, then **Send File**.
- b. Enter the file path and name in the Filename field.
- c. In the Protocol field, select **Xmodem**.
- d. Click on the **[Send]** button.

The download then commences. It can take several minutes, depending on the baud rate set in the switch and in your terminal emulator.

6. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch**). You then see the following prompt:

Continue reboot of system? : No

Press the space bar once to change **No** to **Yes**, then press **[Enter]** to begin the reboot.

7. To confirm that the software downloaded correctly:
 - a. From the Main Menu, select
 1. **Status and Counters**
 1. **General System Information**
 - b. Check the **Firmware revision** line.

Downloading to primary or secondary flash using Xmodem and a terminal emulator (CLI)

Syntax:

```
copy xmodem flash [<primary | secondary>]
```

Downloads a software file to primary or secondary flash. If you do not specify the flash destination, the Xmodem download defaults to primary flash.

Example:

To download a switch software file named `E0822.swi` from a PC (running a terminal emulator program such as HyperTerminal) to primary flash:

Procedure

1. Execute the following command in the CLI:

```
switch# copy xmodem flash
Press 'Enter and start XMODEM on your host...
```

2. Execute the terminal emulator commands to begin the Xmodem transfer. For example, using HyperTerminal:
 - a. Click on **Transfer**, then **Send File**.
 - b. Type the file path and name in the Filename field.
 - c. In the Protocol field, select **Xmodem**.
 - d. Click on the **[Send]** button.

The download can take several minutes, depending on the baud rate used in the transfer.

3. When the download finishes, you must reboot the switch to implement the newly downloaded software. To do so, use one of the following commands:

Syntax:

```
boot system flash {<primary | secondary>}
```

Reboots from the selected flash

Syntax:

```
reload
```

Reboots from the flash image currently in use

For more information on these commands, see “Rebooting the Switches” in the basic operation guide for your switch.

4. To confirm that the software downloaded correctly:

```
switch# show system
```

Check the **Firmware revision** line. It should show the software version that you downloaded in the preceding steps.

If you need information on primary/secondary flash memory and the boot commands, see "Using Primary and Secondary Flash Image Options" in the basic operation guide for your switch.

Switch-to-switch download

You can use TFTP to transfer a software image between two switches of the same series. The CLI enables all combinations of flash location options. The menu interface enables you to transfer primary-to-primary or secondary-to-primary.

Switch-to-switch download to primary flash (Menu)

Using the menu interface, you can download a switch software file from either the primary or secondary flash of one switch to the primary flash of another switch of the same series.

Procedure

1. From the switch console Main Menu in the switch to receive the download, select **7. Download OS** screen.
2. Ensure that the **Method** parameter is set to **TFTP** (the default).
3. In the **TFTP Server** field, enter the IP address of the remote switch containing the software file you want to download.
4. For the **Remote File Name** , enter one of the following:
 - a. To download the software in the primary flash of the source switch, enter
flash

in lowercase characters.
 - b. To download the software in the secondary flash of the source switch, enter
/os/secondary

5. Press **[Enter]**, and then **[X]** (for **eXecute**) to begin the software download.

A "progress" bar indicates the progress of the download. When the entire switch software download has been received, all activity on the switch halts and the following messages appear:

Validating and writing system software to FLASH...

6. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch**). You then see this prompt:

Continue reboot of system? : No

Press the space bar once to change **No** to **Yes**, then press **[Enter]** to begin the reboot.

7. To confirm that the software downloaded correctly:
 - a. From the Main Menu, select
Status and Counters
General System Information
 - b. Check the **Firmware revision** line.

Downloading the OS from another switch (CLI)

Where two switches in your network belong to the same series, you can download a software image between them by initiating a `copy tftp` command from the destination switch. The options for this CLI feature include:

- Copy from primary flash in the source to either primary or secondary in the destination.
- Copy from either primary or secondary flash in the source to either primary or secondary flash in the destination.

Downloading from primary only (CLI)

Syntax:

```
copy tftp flash <ip-addr> flash [primary | secondary]
```

When executed in the destination switch, downloads the software flash in the source switch's primary flash to either the primary or secondary flash in the destination switch.

If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

To download a software file from primary flash in a switch with an IP address of 10.29.227.103 to the primary flash in the destination switch, you would execute the following command in the destination switch's CLI:

Switch-to-switch, from primary in source to either flash in destination

```
switch# copy tftp flash 10.29.227.13 flash
Device will be rebooted, do you want to continue [y/n]? y
00107K 1
```

- ¹Running Total of Bytes Downloaded

Downloading from either flash in the source switch to either flash in the destination switch (CLI)

Syntax:

```
copy tftp flash <ip-addr> {</os/primary> | </os/secondary>} [primary | secondary]
```

This command (executed in the destination switch) gives you the most options for downloading between switches. If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

To download a software file from secondary flash in a switch with an IP address of 10.28.227.103 to the secondary flash in a destination switch, you would execute the following command in the destination switch's CLI:

Switch-to-switch, from either flash in source to either flash in destination

```
switch# copy tftp flash 10.29.227.13 flash /os/secondary secondary
Device will be rebooted, do you want to continue [y/n]? y
00184K
```

Using AirWave to update switch software

AirWave can be used to update HPE switch products. For further information, refer to the **ZTP with Airwave network Management** chapter in this manual.

Using IMC to update switch software

IMC includes a software update utility for updating on HPE switch products. For further information, refer to the getting started guide and the administrator's guide, provided electronically with the application.

Copying software images



For details on how switch memory operates, including primary and secondary flash, see “Switch Memory and Configuration” in the basic operation guide for your switch.

TFTP: Copying a software image to a remote host (CLI)

Syntax:

```
copy flash tftp <ip-addr> <filename>
```

Copies the primary flash image to a TFTP server.

Example:

To copy the primary flash to a TFTP server having an IP address of 10.28.227.105:

```
switch# copy flash tftp 10.28.227.105 k0800.swi
```

where `k0800.swi` is the filename given to the flash image being copied.

Xmodem: Copying a software image from the switch to a serially connected PC or UNIX workstation (CLI)

To use this method, the switch must be connected via the serial port to a PC or UNIX workstation.

Syntax:

```
copy flash xmodem {[<pc> | unix>]}
```

Uses Xmodem to copy a designated configuration file from the switch to a PC or UNIX workstation.

Example:

To copy the primary flash image to a serially connected PC:

Procedure

1. Execute the following command:

```
switch# copy xmodem flash  
Press 'Enter' and start XMODEM on your host...
```

2. After you see the above prompt, press **[Enter]**.
3. Execute the terminal emulator commands to begin the file transfer.

Transferring switch configurations

Using the CLI commands described in the section beginning with **TFTP: Copying a configuration file to a remote host (CLI)** on page 262, you can copy switch configurations to and from a switch, or copy a software image to configure or replace an ACL in the switch configuration.



For greater security, you can perform all TFTP operations using SFTP, as described in the section **Using SCP and SFTP** on page 251.

You can also use the `include-credentials` command to save passwords, secret keys, and other security credentials in the running config file. For more information, see the section on "Saving Security Credentials in a Config File" in the access security guide for your switch.

TFTP: Copying a configuration file to a remote host (CLI)

Syntax:

```
copy {<startup-config | running-config>} tftp <ip-addr> <remote-file> [pc | unix]
```

```
copy config <filename> tftp <ip-addr> <remote-file> [pc | unix]
```

This command can copy a designated config file in the switch to a TFTP server. For more information, see "Multiple Configuration Files" in the basic operation guide for your switch.

Example:

To upload the current startup configuration to a file named `sw8200` in the `configs` directory on drive `"d"` in a TFTP server having an IP address of 10.28.227.105:

```
switch# copy startup-config tftp 10.28.227.105
d:\configs\sw8200
```

TFTP: Copying a configuration file from a remote host (CLI)

Syntax:

```
copy tftp {<startup-config | running-config> <ip-address> <remote-file>} [pc | unix]
```

```
copy tftp config <filename> <ip-address> <remote-file> [pc | unix]
```

This command can copy a configuration from a remote host to a designated config file in the switch. For more information, see "Multiple Configuration Files" in the basic operation guide for your switch.

For more information on flash image use, see "Using Primary and Secondary Flash Image Options" in the basic operation guide for your switch.

Example:

To download a configuration file named `sw8200` in the `configs` directory on drive `"d"` in a remote host having an IP address of 10.28.227.105:

```
switch# copy tftp startup-config 10.28.227.105
d:\configs\sw8200
```

TFTP: Copying a customized command file to a switch (CLI)

Using the `copy tftp` command with the `show-tech` option provides the ability to copy a customized command file to the switch. When the `show tech custom` command is executed, the commands in the custom file are executed instead of the hard-coded list of commands. If no custom file is found, the current hard-coded list is executed. This list contains commands to display data, such as the image stamp, running configuration, boot history, port settings, and so on.

Syntax:

```
copy tftp show-tech <ipv4 or ipv6 address> <filename>
```

Copies a customized command file to the switch (see [Using the copy tftp show-tech command to upload a customized command file](#) on page 263).

Using the copy tftp show-tech command to upload a customized command file

```
switch(config)# copy tftp show-tech 10.10.10.3 commandfile1
```

Syntax:

```
show tech custom
```

Executes the commands found in a custom file instead of the hard-coded list.



Exit the global config mode (if needed) before executing `show tech` commands.

You can include `show tech` commands in the custom file, with the exception of `show tech custom`. For example, you can include the command `show tech all`.

If no custom file is found, a message displays stating "No SHOW-TECH file found." (No custom file was uploaded with the `copy tftp show-tech` command.)

The show tech custom command

```
switch# show tech custom  
No SHOW-TECH file found.
```

Xmodem: Copying a configuration file to a serially connected PC or UNIX workstation (CLI)

To use this method, the switch must be connected via the serial port to a PC or UNIX workstation. You will need to:

- Determine a filename to use
- Know the directory path you will use to store the configuration file.

Syntax:

```
copy {<startup-config | running-config>} xmodem {<pc | unix>}
```

```
copy config <filename> xmodem {<pc | unix>}
```

Uses Xmodem to copy a designated configuration file from the switch to a PC or UNIX workstation. For more information, see "Multiple Configuration Files" in the basic operation guide for your switch.

Example:

To copy a configuration file to a PC serially connected to the switch:

1. Determine the file name and directory location on the PC.
2. Execute the following command:

```
switch# copy startup-config xmodem pc
Press 'Enter' and start XMODEM on your host...
```

3. After you see the above prompt, press **[Enter]**.
4. Execute the terminal emulator commands to begin the file transfer.

Xmodem: Copying a configuration file from a serially connected PC or UNIX workstation (CLI)

To use this method, the switch must be connected via the serial port to a PC or UNIX workstation on which is stored the configuration file you want to copy. To complete the copying, you need to know the name of the file to copy and the drive and directory location of the file.

Syntax:

```
copy xmodem startup-config {<pc | unix>}
```

```
copy xmodem config <filename> < {pc | unix}>
```

Copies a configuration file from a serially connected PC or UNIX workstation to a designated configuration file on the switch.

For more information, see "Multiple Configuration Files" in the basic operation guide for your switch.

Example:

To copy a configuration file from a PC serially connected to the switch:

Procedure

1. Execute the following command:

```
switch# copy xmodem startup-config pc
Device will be rebooted, do you want to continue [y/n]? y
Press 'Enter' and start XMODEM on your host...
```

2. After you see the above prompt, press **[Enter]**.
3. Execute the terminal emulator commands to begin the file transfer.
4. When the download finishes, you must reboot the switch to implement the newly downloaded software. To do so, use one of the following commands:

Syntax:

```
boot system flash [primary | secondary]
```

```
boot system flash [config < filename >]
```

Switches boot from the designated configuration file. For more information, see "Multiple Configuration Files" in the basic operation guide for your switch.

Syntax:

```
reload
```

Reboots from the flash image currently in use.

(For more on these commands, see "Rebooting the Switch" in the basic operation guide for your switch.)

Copying diagnostic data to a remote host, PC or UNIX workstation

You can use the CLI to copy the following types of switch data to a text file in a destination device:

Command output	Sends the output of a switch CLI command as a file on the destination device.
Event log	Copies the switch's Event Log into a file on the destination device.
Crash data	Software-specific data useful for determining the reason for a system crash.
Crash log	Processor-specific operating data useful for determining the reason for a system crash.
Flight data recorder (FDR) logs	Information that is "interesting" at the time of the crash, as well as when the switch is not performing correctly but has not crashed.

The destination device and copy method options are as follows (CLI keyword is in bold):

- Remote Host via **TFTP**.
- Serially connected PC or UNIX workstation via **Xmodem**.

Copying command output to a destination device (CLI)

Syntax:


```
copy command-output <"cli-command"> tftp <ip-address> <filepath-filename>
```

```
copy command-output <"cli-command"> xmodem
```

Example:

To use Xmodem to copy the output of `show config` to a serially connected PC:

Figure 41: Sending command output to a file on an attached PC



The command you specify must be enclosed in double quotation marks.

Copying Event Log output to a destination device (CLI)

Syntax:


```
copy event-log tftp <ip-address> <filepath_filename>
```

```
copy event-log xmodem <filename>
```

Example:

To copy the event log to a PC connected to the switch:

Figure 42: Sending event log content to a file on an attached PC



Copying crash data content to a destination device (CLI)

This command uses TFTP, USB or Xmodem to copy the Crash Data content to a destination device.

Syntax:

```
copy crash-data tftp <ip-address> <filename>
```

```
copy crash-data xmodem
```

These commands copy the crash data content to a remote host, attached USB device, or to a serially connected PC or UNIX workstation.

slot-id	a - h—Retrieves the crash log or crash data from the processor on the module in the specified slot
mm	Retrieves crash log or crash data from the switch's chassis processor. When "mm" is specified, crash files from both management modules are copied.
oobm	For switches that have a separate OOBM port, specifies that the transfer is through the OOBM interface. (Default is transfer through the data interface.)

You can copy individual slot information or the management module (mm) switch information. If you do not specify either, the command defaults to the mm data.

To copy the switch's crash data to a file in a PC:

Copying switch crash data content to a PC

```
switch(config)# copy crash-data xmodem pc
Press 'Enter' and start XMODEM on your host... 3
```

Transfer complete

- At this point press [Enter] and start the Xmodem command sequence in your terminal emulator.

³ At this point press [Enter] and start the Xmodem command sequence in your terminal emulator.

Overview

The switches have several built-in tools for monitoring, analyzing, and troubleshooting switch and network operation:

- **Status:** Includes options for displaying general switch information, management address data, port status, port and trunk group statistics, MAC addresses detected on each port or VLAN, and STP, IGMP, and VLAN data.
- **Counters:** Display details of traffic volume on individual ports.
- **Event Log:** Lists switch operating events ([Using the Event Log for troubleshooting switch problems](#) on page 314).
- **Alert Log:** Lists network occurrences detected by the switch—in the System > Logging screen of the WebAgent.
- **Configurable trap receivers:** Uses SNMP to enable management stations on your network to receive SNMP traps from the switch.
- **Port monitoring (mirroring):** Copy all traffic from the specified ports to a designated monitoring port ([Interface monitoring features](#) on page 289).
- **Chassis Locator LED:** The blue Locator LED lights up when you enter the `chassislocate` command.



Link test and ping test—analysis tools in troubleshooting situations—are described in [Troubleshooting](#) on page 292. See [Diagnostic tools](#) on page 351.

Accessing port and trunk group statistics

Use the CLI to view port counter summary reports, and to view detailed traffic summary for specific ports.

show interfaces

Syntax

```
show interfaces <PORT-LIST>
```

Description

Provides an overview of port activity for all ports on the switch or for the ports you specify. Displays the totals accumulated since the last boot or the last execution of the `clear statistics` command.

Parameters and options

<PORT-LIST> View port activity for specific ports.

Reset port counters

When troubleshooting network issues, you can clear all counters and statistics without rebooting the switch using the `clear statistics global` command or using the menu.

SNMP displays the counter and statistics totals accumulated since the last reboot, and it is not affected by the `clear statistics global` command or the `clear statistics <PORT-LIST>` command. Clearing statistics initiates an SNMP trap.

ⓘ Once cleared, statistics cannot be reintroduced.

clear statistics

Syntax

```
clear statistics [<PORT-LIST>|global]
```

Description

This command clears all counters and statistics for all interfaces except SNMP.

Parameters and options

- | | |
|--------------------------------|--|
| <code><PORT-LIST></code> | Clears the counters and statistics for specific ports. |
| <code>global</code> | Clears all counters and statistics for all interfaces except SNMP. |

Accessing port and trunk statistics (Menu)

Procedure

1. From the Main Menu, select **1. Status and Counters ...**, and then select **4. Port Counters**.

Figure 43: Example of port counters on the menu interface

```
===== CONSOLE - MANAGER MODE =====
                        Status and Counters - Port Counters
-----
```

Port	Total Bytes	Total Frames	Errors Rx	Drops Tx	Flow Ctrl
A1	195,072	323	0	0	off
A2	651,816	871	0	0	off
A3-Trk1	290,163	500	0	0	off
A4-Trk1	260,134	501	0	0	off
C1	859,363	5147	0	0	off
C2	674,574	1693	0	0	off
C3	26,554	246	0	0	off
C4	113,184	276	0	0	off
C5	0	0	0	0	off

```
-----
Actions->  Back  Show details  Reset  Help
Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

2. To view details about the traffic on a particular port, use the ↓ key to highlight that port number, and then select **Show Details**. For example, selecting port A2 displays a screen similar to the example below.

Figure 44: Example of the display for Show Details on a selected port

```
===== CONSOLE - MANAGER MODE =====
                        Status and Counters - Port Counters - Port A2
-----
```

Link Status	: Up		
Bytes Rx	: 630,746	Bytes Tx	: 21,070
Unicast Rx	: 568	Unicast Tx	: 285
Bcast/Mcast Rx	: 18	Bcast/Mcast Tx	: 0
FCS Rx	: 0	Drops Tx	: 0
Alignment Rx	: 0	Collisions Tx	: 0
Runts Rx	: 0	Late Colln Tx	: 0
Giants Rx	: 0	Excessive Colln	: 0
Total Rx Errors	: 0	Deferred Tx	: 0

```
-----
Actions->  Back  Reset  Help
Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

This screen also includes the **Reset** action for the current session.

MAC address tables

MAC address views and searches

You can view and search MAC addresses using the CLI or the menu.

show mac-address

Syntax

```
show mac-address [vlan <VLAN-ID>] [<PORT-LIST>] [<MAC-ADDR>]
```

Description

Lists all MAC addresses on the switch and their corresponding port numbers. You can also choose to list specific addresses and ports, or addresses and ports on a VLAN. The switches operate with a multiple forwarding database architecture.

List all learned MAC addresses on the switch and corresponding port numbers

```
switch# show mac-address
```

List all learned MAC addresses on one or more ports and corresponding port numbers

```
switch# show mac-address a1-a4,a6
```

List all learned MAC addresses on a VLAN and corresponding port numbers

```
switch# show mac-address vlan 100
```

List the port on which the switch learned a specific MAC address

To find the port on which the switch learns a MAC address of 080009-21ae84:

```
Select VLAN : DEFAULT VLAN
```

Using the menu to view and search MAC addresses

To determine which switch port on a selected VLAN the switch uses to communicate with a specific device on the network:

Procedure

1. From the Main Menu, select **1. Status and Counters ...**, and then select **5. VLAN Address Table**.
2. Use the arrow keys to scroll to the VLAN you want, and then press **Enter** on the keyboard to select it.

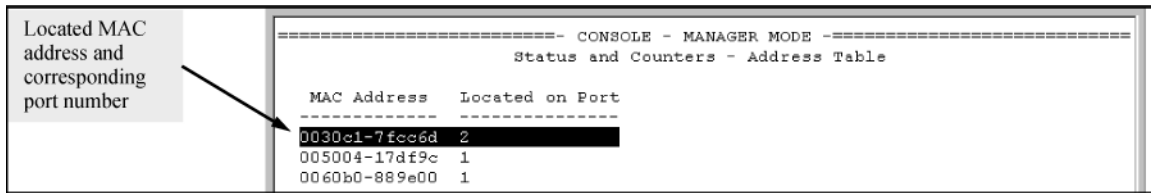
```
----- CONSOLE - MANAGER MODE -----
                Status and Counters - Address Table

  MAC Address   Located on Port
-----
0030c1-7f49c0  A3
0030c1-7fec40  A1
0030c1-b29ac0  A3
0060b0-17de5b  A3
0060b0-880a80  A2
0060b0-df1a00  A3
0060b0-df2a00  A3
0060b0-e9a200  A3
009027-e74f90  A3
080009-21ae84  A3
080009-62c411  A3
080009-6563e2  A3

Actions-> Back   Search   Next page   Prev page   Help
Return to previous screen.
Use up/down arrow keys to scroll to other entries, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

The switch then displays the MAC address table for that VLAN (**Figure 45: Example of the address table** on page 272.)

Figure 45: Example of the address table



```
----- CONSOLE - MANAGER MODE -----
                        Status and Counters - Address Table
-----
MAC Address      Located on Port
-----
0030c1-7fcc6d   2
005004-17df9c   1
0060b0-889e00   1
```

3. To page through the listing, use **Next page** and **Prev page** .

Finding the port connection for a specific device on a VLAN

This feature uses a device's MAC address that you enter to identify the port used by that device.

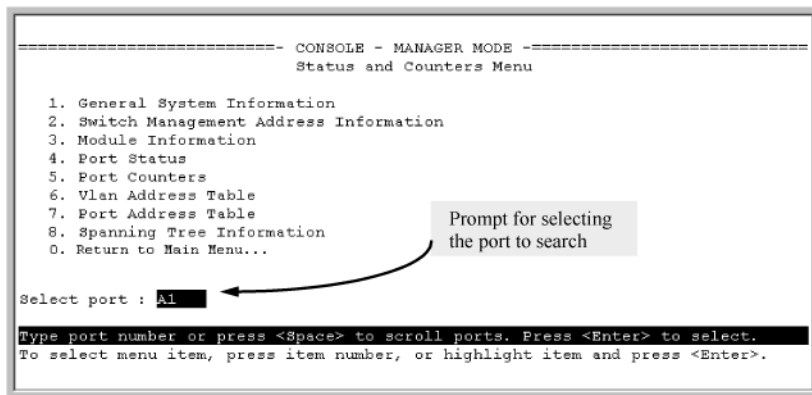
Procedure

1. Proceeding from **Figure 45: Example of the address table** on page 272, press **[S]** (for **Search**), to display the following prompt:

```
Enter MAC address: _
```

2. Enter the MAC address you want to locate and press **[Enter]**.
3. The address and port number are highlighted if found (**Figure 46: Example of menu indicating located MAC address** on page 272.) If the switch does not find the MAC address on the currently selected VLAN, it leaves the MAC address listing empty.

Figure 46: Example of menu indicating located MAC address



```
----- CONSOLE - MANAGER MODE -----
                        Status and Counters Menu

1. General System Information
2. Switch Management Address Information
3. Module Information
4. Port Status
5. Port Counters
6. Vlan Address Table
7. Port Address Table
8. Spanning Tree Information
0. Return to Main Menu...

Select port : A1
Type port number or press <Space> to scroll ports. Press <Enter> to select.
To select menu item, press item number, or highlight item and press <Enter>.
```

4. Press **[P]** (for **Prev page**) to return to the full address table listing.

Viewing and searching port-level MAC addresses

This feature displays and searches for MAC addresses on the specified port instead of for all ports on the switch.

Procedure

1. From the Main Menu, select:
 1. **Status and Counters ...**
 7. **Port Address Table**
2. Use the Space bar to select the port you want to list or search for MAC addresses, then press **[Enter]** to list the MAC addresses detected on that port.

Determining whether a specific device is connected to the selected port

Proceeding from Step 2, above:

Procedure

1. Press **[S]** (for **Search**), to display the following prompt:

```
Enter MAC address: _
```

2. Enter the MAC address you want to locate and press **[Enter]**.

The address is highlighted if found. If the switch does not find the address, it leaves the MAC address listing empty.

3. Press **[P]** (for **Prev page**) to return to the previous per-port listing.

MSTP data

show spanning-tree

Syntax

```
show spanning-tree
```

Description

Displays the global and regional spanning-tree status for the switch, and displays the per-port spanning-tree operation at the regional level.

Values for the following parameters appear only for ports connected to active devices: `Designated Bridge`, `Hello Time`, `PtP`, and `Edge`.

show spanning-tree command output

Figure 47: show spanning-tree command output

```
HP Switch(config)# show spanning-tree

Multiple Spanning Tree (MST) Information
-----
| STP Enabled      : Yes
| Force Version   : MSTP-operation
| IST Mapped VLANs : 1,66
|
| Switch MAC Address : 0004ea-5e2000
| Switch Priority   : 32768
| Max Age         : 20
| Max Hops        : 20
| Forward Delay   : 15
|
| Topology Change Count : 0
| Time Since Last Change : 2 hours
|-----
| CST Root MAC Address : 00022d-47367f
| CST Root Priority    : 0
| CST Root Path Cost  : 4000000
| CST Root Port       : A1
|-----
| IST Regional Root MAC Address : 00883-028300
| IST Regional Root Priority     : 32768
| IST Regional Root Path Cost    : 200000
| IST Remaining Hops             : 19
|-----
| Protected Ports : A4
| Filtered Ports  : A7-A10
|-----

Port Type | Cost | Priority | State | Designated Bridge | Hello Time | PTP | Edge
-----+-----+-----+-----+-----+-----+-----+-----
A1 100/1000T | Auto | 128 | Forwarding | 000883-028300 | 9 | Yes | No
A2 100/1000T | Auto | 128 | Blocked | 0001e7-948300 | 9 | Yes | No
A3 100/1000T | Auto | 128 | Forwarding | 000883-02a700 | 2 | Yes | No
A4 100/1000T | Auto | 128 | Disabled | . | . | . | .
A5 100/1000T | Auto | 128 | Disabled | . | . | . | .
. . . . .
. . . . .
```

Switch's Spanning Tree Configuration and Identity of VLANs Configured in the Switch for the IST Instance

Identifies the overall spanning-tree root for the network.

Lists the switch's MSTP root data for connectivity with other regions and STP or RSTP devices.

Identifies the spanning-tree root for the IST Instance for the region.

Internal Spanning Tree Data (IST Instance) for the region in which the Switch Operates

Identifies the ports with BPDU protection and BPDU filtering enabled.

Yes means the switch is operating the port as if it is connected to switch, bridge, or end node (but *not* a hub).

For Edge, No (admin-edge-port operation disabled) indicates the port is configured for connecting to a LAN segment that includes a bridge or switch. Yes indicates the port is configured for a host (end node) link. Refer to the **admin-edge-port** description under "Configuring MSTP Per-Port Parameters" on page 3-24.

IP IGMP status

show ip igmp

Syntax

```
show ip igmp <VLAN-ID> [config] [group <IP-ADDR>] [groups] [statistics]
```

Description

Global command that lists IGMP status for all VLANs configured in the switch, including:

- VLAN ID (VID) and name
- Querier address
- Active group addresses per VLAN
- Number of report and query packets per group
- Querier access port per VLAN

Parameters and options

config	Displays the IGMP configuration information, including VLAN ID, VLAN name, status, forwarding, and Querier information.
vlan-id	Per-VLAN command listing above, IGMP status for specified VLAN (VID).
group <IP-ADDR>	Lists the ports currently participating in the specified group, with port type, Access type, Age Timer data and Leave Timer data.
groups	Displays VLAN-ID, group address, uptime, expiration time, multicast filter type, and the last reporter for IGMP groups.
statistics	Displays IGMP operational information, such as VLAN IDs and names, and filtered and flooding statistics.

Output from show ip igmp config command

```
HP Switch(config)# show ip igmp config
```

```
IGMP Service
```

VLAN ID	VLAN Name	IGMP Enabled	Forward with High Priority	Querier Allowed	Querier Interval
1	DEFAULT_VLAN	No	No	Yes	125
2	VLAN2	Yes	No	Yes	125
12	New_Vlan	No	No	Yes	125

IGMP statistical information

```
switch(vlan-2)# show ip igmp statistics
```

IGMP Service Statistics

```
Total VLANs with IGMP enabled           : 1
Current count of multicast groups joined  : 1
```

IGMP Joined Groups Statistics

VLAN ID	VLAN Name	Filtered	Flood
2	VLAN2	2	1

VLAN information

show vlan

Syntax

```
show vlan <VLAN-ID>
```

Description

Lists the maximum number of VLANs to support, existing VLANs, VLAN status (static or dynamic), and primary VLAN.

Parameters and options

<VLAN-ID>

Lists the following for the specified VLAN:

- Name, VID, and status (static/dynamic)
- Per-port mode (tagged, untagged, forbid, no/auto)
- "Unknown VLAN" setting (Learn, Block, Disable)
- Port status (up/down)

List data on specific VLANs

The next three figures show how you can list data for the following VLANs:

Ports	VLAN	VID
A1-A12	DEFAULT_VLAN	1
A1, A2	VLAN-33	33
A3, A4	VLAN-44	44

Figure 48: Listing the VLAN ID (vid) and status for specific ports

```

HP Switch# show vlan ports A1-A2

Status and Counters = VLAN Information - for ports A1,A2

802.1Q VLAN ID Name                Status
-----
1          DEFAULT_VLAN             Static
33         VLAN-33                  Static

```

Because ports A1 and A2 are not members of VLAN-44, it does not appear in this listing.

Figure 49: Example of VLAN listing for the entire switch

```

HP Switch# show vlan
Status and Counters - VLAN Information

VLAN support : Yes
Maximum VLANs to support : 9
Primary VLAN: DEFAULT_VLAN

802.1Q VLAN ID Name                Status
-----
1          DEFAULT_VLAN             Static
33         VLAN-33                  Static
44         VLAN-44                  Static

```

Figure 50: Port listing for an individual VLAN

```

HP Switch(config)# show vlan 1

Status and Counters - VLAN Information - VLAN 1

VLAN ID : 1
Name : DEFAULT_VLAN
Status : Static
Voice : Yes
Jumbo : No

Port Information Mode      Unknown VLAN Status
-----
A1          Untagged Learn      Up
A2          Tagged   Learn      Up
A3          Untagged Learn      Up
A4          Untagged Learn      Down
A5          Untagged Learn      Up
A6          Untagged Learn      Up
A7          Untagged Learn      Up

```

Configuring a source switch in a local mirroring session

Enter the `mirror port` command on the source switch to configure an exit port on the same switch. To create the mirroring session, use the information gathered in [High-level overview of the mirror configuration process](#) on page 281.

Syntax

```
mirror 1 - 4 port exit-port-# [name name-str] no mirror 1- 4
```

Assigns the exit port to use for the specified mirroring session and must be executed from the global configuration level.

1 - 4	Identifies the mirroring session created by this command. (Multiple sessions on the switch can use the same exit port.)
name <i>name-str</i>	Optional alphanumeric name string used to identify the session (up to 15 characters)
port <i>exit-port-#</i>	Exit port for mirrored traffic in the remote session. This is the port to which a traffic analyzer or IDS is connected.

The `no` form of the command removes the mirroring session and any mirroring source previously assigned to that session.

Selecting all traffic on a port interface for mirroring according to traffic direction

Syntax

```
[no] interface port/trunk/mesh monitor [in | out | both] [mirror 1 - 4 | name-str] [{1 - 4 | name-str} | {1 - 4 | name-str} | {1 - 4 | name-str}] [no-tag-added]
```

Assigns a mirroring source to a previously configured mirroring session on a source switch by specifying the port, trunk, and/or mesh sources to use, the direction of traffic to mirror, and the session.

<code>interface <i>port/trunk/mesh</i></code>	Identifies the source ports, static trunks, and/or mesh on which to mirror traffic. Use a hyphen for a range of consecutive ports or trunks (a5-a8, Trk2-Trk4.) Use a comma to separate non-contiguous interfaces (b11, b14, Trk4, Trk7.)
<code>monitor all [<i>in</i> <i>out</i> <i>both</i>]</code>	For the interface specified by <i>port/trunk/mesh</i> , selects traffic to mirror based on whether the traffic is entering or leaving the switch on the interface: <ul style="list-style-type: none"> <code>in</code>: Mirrors entering traffic. <code>out</code>: Mirrors exiting traffic. <code>both</code>: Mirrors traffic entering and exiting. <p>If you enter the <code>monitor all</code> command without selection criteria or a session identifier, the command applies by default to session 1</p>

Table Continued

mirror [1 - 4 | *name-str*]

Assigns the traffic specified by the interface and direction to a session by number or—if configured—by name. The session must have been previously configured.

Depending on how many sessions are already configured on the switch, you can use the same command to assign the specified source to up to four sessions, for example, `interface a1 monitor all in mirror 1 2 4`.

- 1 - 4: Configures the selected port traffic to be mirrored in the specified session number.
- [*name name-str*]: Optional; configures the selected port traffic to be mirrored in the specified session name. The string can be used interchangeably with the session number when using this command to assign a mirroring source to a session.

[no-tag-added]

Prevents a VLAN tag from being added to the mirrored copy of an outbound packet sent to a local or remote mirroring destination.

The `no` form of the command removes a mirroring source assigned to the session, but does not remove the session itself. This enables you to repurpose a session by removing an unwanted mirroring source and adding another in its place.

Viewing all mirroring sessions configured on the switch

Syntax

show monitor

If a monitored source for a remote session is configured on the switch, the following information is displayed. Otherwise, the output displays: Mirroring is currently disabled.

Sessions	Lists the four configurable sessions on the switch.
Status	Displays the current status of each session: <ul style="list-style-type: none">• active: The session is configured.• inactive: Only the destination has been configured; the mirroring source is not configured.• not defined: Mirroring is not configured for this session.
Type	Indicates whether the mirroring session is local (<code>port</code>), remote (<code>IPv4</code>), or MAC-based (<code>mac</code>) for local or remote sessions.
Sources	Indicates how many monitored source interfaces are configured for each mirroring session.
Policy	Indicates whether the source is using a classifier-based mirroring policy to select inbound IPv4 or IPv6 traffic for mirroring.

If a remote mirroring endpoint is configured on the switch, the following information is displayed. Otherwise, the output displays: There are no Remote Mirroring endpoints currently assigned.

Type	Indicates whether the mirroring session is local (<code>port</code>), remote (<code>IPv4</code>), or MAC-based (<code>mac</code>) for local or remote sessions.
UDP Source Addr	The IP address configured for the source VLAN or subnet on which the monitored source interface exists. In the configuration of a remote session, the same UDP source address must be configured on the source and destination switches.
UDP port	The unique UDP port number that identifies a remote session. In the configuration of a remote session, the same UDP port number must be configured on the source and destination switches.
UDP Dest Addr	The IP address configured for the destination VLAN or subnet on which the remote exit port exists. In the configuration of a remote session, the same UDP destination address must be configured on the source and destination switches.
Dest Port	Identifies the exit port for a remote session on a remote destination switch.

Figure 51: *Displaying the currently configured mirroring sessions on the switch*

<pre>HP Switch# show monitor Network Monitoring</pre>						<p>Local and Remote Mirroring Sources:</p> <ul style="list-style-type: none"> • Session 1 is performing local mirroring using a classifier-based policy as traffic-selection criteria. • Session 2 is performing remote mirroring using MAC-based traffic-selection criteria. • Session 3 is not configured. • Session 4 is configured for remote mirroring from a non-policy source (for example, traffic direction), but is currently not mirroring any traffic. 																				
<table border="1"> <thead> <tr> <th>Sessions</th> <th>Status</th> <th>Type</th> <th>Sources</th> <th>Policy</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>active</td> <td>port</td> <td>1</td> <td>yes</td> </tr> <tr> <td>2</td> <td>active</td> <td>mac</td> <td>2</td> <td>no</td> </tr> <tr> <td>3</td> <td>not defined</td> <td></td> <td></td> <td></td> </tr> <tr> <td>4</td> <td>inactive</td> <td>IPv4</td> <td>0</td> <td>no</td> </tr> </tbody> </table>	Sessions	Status	Type	Sources	Policy		1	active	port	1	yes	2	active	mac	2	no	3	not defined				4	inactive	IPv4	0	no
Sessions	Status	Type	Sources	Policy																						
1	active	port	1	yes																						
2	active	mac	2	no																						
3	not defined																									
4	inactive	IPv4	0	no																						
<pre>Remote Mirroring - Remote Endpoints</pre>																										
<table border="1"> <thead> <tr> <th>Type</th> <th>UDP Source Addr</th> <th>UDP port</th> <th>UDP Dest Addr</th> </tr> </thead> <tbody> <tr> <td>IPv4</td> <td>10.10.30.1</td> <td>7950</td> <td>10.10.20.1</td> </tr> </tbody> </table>	Type	UDP Source Addr	UDP port	UDP Dest Addr	IPv4	10.10.30.1	7950	10.10.20.1	<p>Remote Mirroring Destination:</p> <p>The switch is configured as a remote mirroring destination (endpoint) for a source at 10.10.30.1, using port B10 as the exit port.</p>																	
Type	UDP Source Addr	UDP port	UDP Dest Addr																							
IPv4	10.10.30.1	7950	10.10.20.1																							

Viewing the mirroring configuration for a specific session

Syntax

```
show monitor [1 - 4 | name name-str]
```

Displays detailed configuration information for a specified local mirroring session on a source switch.

Session	Displays the number of the specified session.
Session Name	Displays the name of the session, if configured.
Policy	Indicates whether the source is using a classifier-based mirroring policy to select inbound IPv4 or IPv6 traffic for mirroring.
Mirroring Destination	For a local mirroring session, displays the port configured as the exit port on the source switch.

Table Continued

Monitoring Sources	For the specified local session, displays the source (port, trunk, or VLAN) interface and the MAC address (if configured) used to select mirrored traffic.
Direction	For the selected interface, indicates whether mirrored traffic is entering the switch (<i>in</i>), leaving the switch (<i>out</i>), or <i>both</i> .

Using the Menu to configure local mirroring

Menu and WebAgent limits

You can use the Menu and WebAgent to quickly configure or reconfigure local mirroring on session 1 and allow one of the following two mirroring source options:

- Any combination of source ports, trunks, and a mesh.
- One static, source VLAN interface.

The Menu and WebAgent also has these limits:

- Configure and display only session 1 and only as a local mirroring session for traffic in **both** directions on the specified interface. (Selecting inbound-only or outbound-only is not an option.)
- If session 1 has been configured in the CLI for local mirroring for inbound-only or outbound-only traffic on one or more interfaces, using the Menu to change the session 1 configuration **automatically reconfigures the session** to monitor both inbound and outbound traffic on the designated interface(s.)
- If session 1 has been configured in the CLI with an ACL/classifier-based mirroring policy or as a remote mirroring session, the Menu is not available for changing the session 1 configuration.
- The CLI (and SNMP) can be used to override any Menu configuration of session 1.

High-level overview of the mirror configuration process

Determine the mirroring session and destination

For a local mirroring session

Determine the port number for the exit port (such as A5, B10, and so forth), then go to **Configure the monitored traffic in a mirror session**.

Classifier-based mirroring configuration

1. Evaluate the types of traffic in your network and identify the traffic types that you want to mirror.
2. Create an IPv4 or IPv6 traffic class using the `class` command to select the packets that you want to mirror in a session on a preconfigured local or remote destination device. (See **Configuring classifier-based mirroring**).

A traffic class consists of match criteria, which consist of match and ignore commands.

- `match` commands define the values that header fields must contain for a packet to belong to the class and be managed by policy actions.
- `ignore` commands define the values which, if contained in header fields, exclude a packet from the policy actions configured for the class.



Be sure to enter match/ignore statements in the precise order in which you want their criteria to be used to check packets.

The following match criteria are supported in match/ignore statements for inbound IPv4/IPv6 traffic:

- IP source address (IPv4 and IPv6)
- IP destination address (IPv4 and IPv6)
- IP protocol (such as ICMP or SNMP)
- Layer 3 IP precedence bits
- Layer 3 DSCP codepoint
- Layer 4 TCP/UDP application port (including TCP flags)
- VLAN ID

Enter one or more match or ignore commands from the class configuration context to filter traffic and determine the packets on which policy actions will be performed. (See **Configuring classifier-based mirroring.**)

3. Create a mirroring policy to configure the session and destination device to which specified classes of inbound traffic are sent by entering the `policy mirror` command from the global configuration context.



Be sure to enter each class and its associated mirroring actions in the precise order in which you want packets to be checked and processed.

To configure the mirroring actions that you want to execute on packets that match the criteria in a specified class, enter one or more class action mirror commands from the policy configuration context. (See **Configuring classifier-based mirroring.**)

You can configure only one mirroring session (destination) for each class. However, you can configure the same mirroring session for different classes.

A packet that matches the match criteria in a class is mirrored to the exit (local or remote) port that has been previously configured for the session, where session is a value from 1 to 4 or a text string (if you configured the session with a name when you entered the `mirror` command.)

Prerequisite: The local or remote exit port for a session must be already configured before you enter the `mirror session` parameter in a class action statement:

- In a local mirroring session, the exit port is configured with the `mirror <session-number> port` command
- In a remote mirroring session, the remote exit port is configured with the `mirror endpoint ip` and `mirror <session-number> remote ip` commands.

Restriction: In a policy, you can configure only one mirroring session per class. However, you can configure the same session for different classes.

Mirroring is not executed on packets that match ignore criteria in a class.

The execution of mirroring actions is performed in the order in which the classes are numerically listed in the policy.

The complete no form of the `class action mirror` command or the `no <seq-number>` command removes a class and mirroring action from the policy configuration.

To manage packets that do not match the match or ignore criteria in any class in the policy, and therefore have no mirroring actions performed on them, you can enter an optional default class. The default class is placed at the end of a policy configuration and specifies the mirroring actions to perform on packets that are neither matched nor ignored.

4. (Optional) To configure a default-class in a policy, enter the `default-class` command at the end of a policy configuration and specify one or more actions to be executed on packets that are not matched and not ignored. (See **Configuring classifier-based mirroring.**)

Prerequisite: The local or remote exit port for a session must be already configured with a destination device before you enter the `mirror <session>` parameter in a default-class action statement.

5. Apply the mirroring policy to inbound traffic on a port (`interface service-policy in` command) or VLAN (`vlan service-policy in` command) interface.



After you apply a mirroring policy for one or more preconfigured sessions on a port or VLAN interface, the switch immediately starts to use the traffic-selection criteria and exit port to mirror traffic to the destination device connected to each exit port.

In a remote mirroring session that uses IPv4 encapsulation, if the remote switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic.

For this reason, Hewlett Packard Enterprise strongly recommends that you first configure the exit switch in a remote mirroring session, as described in [Configure a mirroring destination on a remote switch](#) and [Configure a mirroring session on the source switch](#), before you apply a mirroring service policy on a port or VLAN interface.

Restrictions: The following restrictions apply to a mirroring service policy:

- Only one mirroring policy is supported on a port or VLAN interface.
- If you apply a mirroring policy to a port or VLAN interface on which a mirroring policy is already configured, the new policy replaces the existing one.
- A mirroring policy is supported only on inbound traffic.

Because only one mirroring policy is supported on a port or VLAN interface, ensure that the policy you want to apply contains all the required classes and actions for your configuration.

Classifier-based mirroring restrictions

The following restrictions apply to mirroring policies configured with the classifier-based model:

- A mirroring policy is supported only on **inbound** IPv4 or IPv6 traffic.
- A mirroring policy is not supported on a meshed port interface. (Classifier-based policies are supported only on a port, VLAN, or trunk interface.)
- Only one classifier-based mirroring policy is supported on a port or VLAN interface. You can, however, apply a classifier-based policy of a different type, such as QoS.
- You can enter multiple `class action mirror` statements in a policy.
 - You can configure only one mirroring session (destination) for each class.
 - You can configure the same mirroring session for different classes.

- If a mirroring session is configured with a classifier-based mirroring policy on a port or VLAN interface, no other traffic-selection criteria (MAC-based or all inbound and/or outbound traffic) can be added to the session.

Figure 52: Mirroring configuration in which only a mirroring policy is supported

```
Switch-B(config)# mirror endpoint 10.10.40.4 9200 10.10.50.5 port a1
...
Switch-A(config)# mirror 1 remote ip 10.10.40.4 9200 10.10.50.5
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
Switch-A(config)# class ipv4 Data2
Switch-A(config-class)# match ip 10.28.31.1 any
Switch-A(config-class)# match ip any host 10.28.31.0/24
Switch-A(config-class)# exit
Switch-A(config)# policy mirror SalesData
Switch-A(config-policy)# class ipv4 Data2 action mirror 1
Switch-A(config-policy)# exit
Switch-A(config)# vlan 10 service-policy SalesData in
Switch-A(config)# vlan 10 monitor all out mirror 1
A prior mirror policy relationship exists with mirror session 1. Please remove.
```

Classifier-based policy used to select mirrored traffic in session 1

The configuration of additional traffic-direction criteria to select mirrored traffic is not supported in session 1.

- If a mirroring session is already configured with one or more traffic-selection criteria (MAC-based or all inbound and/or outbound traffic), the session does not support the addition of a classifier-based policy.

Figure 53: Mirroring configuration in which only traffic-selection criteria are supported

```
Switch-B(config)# mirror endpoint 10.10.40.4 9200 10.10.50.5 port a1
...
Switch-A(config)# mirror 1 remote ip 10.10.40.4 9200 10.10.50.5
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
Switch-A(config)# vlan 10 monitor all out mirror 1
Switch-A(config)# class ipv4 Data2
Switch-A(config-class)# match ip 10.28.31.1 any
Switch-A(config-class)# match ip any host 10.28.31.0/24
Switch-A(config-class)# exit
Switch-A(config)# policy mirror SalesData
Switch-A(config-policy)# class ipv4 Data2 action mirror 1
Switch-A(config-policy)# exit
Switch-A(config)# vlan 10 service-policy SalesData in
Mirror source VLAN exists on mirror session 1. Cannot add this mirror source.
```

Configuration of traffic-direction criteria to select all outbound traffic on VLAN 10 in mirror session 1

The configuration of an additional classifier-based policy to select mirrored traffic on VLAN 10 is not supported in session 1.

Mirroring configuration examples

Local mirroring using traffic-direction criteria

An administrator wants to mirror the inbound traffic from workstation "X" on port A5 and workstation "Y" on port B17 to a traffic analyzer connected to port C24 (see **Figure 54: Local mirroring topology** on page 285.) In this case, the administrator chooses "1" as the session number. (Any unused session number from 1 to 4 is valid.) Because the switch provides both the source and destination for the traffic to monitor, local mirroring can be used. In this case, the command sequence is:

- Configure the local mirroring session, including the exit port.
- Configure the monitored source interfaces for the session.

Figure 54: Local mirroring topology




Figure 55: Configuring a local mirroring session for all inbound and outbound port traffic

```

HP Switch(config)# mirror 1 port c24
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
HP Switch(config)# interface a5,b17 monitor all in mirror
1

```

Configures port C24 as the mirroring destination (exit port) for session 1.

Reminder to configure mirroring destination before configuring source.

Mirrors all inbound and outbound traffic on ports A5 and B17 to the mirroring destination configured for session 1.

Maximum supported frame size

The IPv4 encapsulation of mirrored traffic adds a 54-byte header to each mirrored frame. If a resulting frame exceeds the MTU allowed in the network, the frame is dropped or truncated.



Oversized mirroring frames are dropped or truncated, according to the setting of the `[truncation]` parameter in the `mirror` command. Also, remote mirroring does not allow downstream devices in a mirroring path to fragment mirrored frames.

If jumbo frames are enabled on the mirroring source switch, the mirroring destination switch and all downstream devices connecting the source switch to the mirroring destination must be configured to support jumbo frames.

Enabling jumbo frames to increase the mirroring path MTU

On 1-Gbps and 10-Gbps ports in the mirroring path, you can reduce the number of dropped frames by enabling jumbo frames on all intermediate switches and routers. (The MTU on the switches covered by this manual is 9220 bytes for frames having an 802.1Q VLAN tag, and 9216 bytes for untagged frames.)

Table 22: Maximum frame sizes for mirroring

	Frame type configuration	Maximum frame size	VLAN tag	Frame mirrored to local port	Frame mirrored to remote port	
				Data	Data	IPv4 header
Untagged	Non-jumbo (default config.)	1518	0	1518	1464	54
	Jumbo ¹ on all VLANs	9216	0	9216	9162	54
	Jumbo ¹ On all but source VLAN	1518	0	n/a ²	1464	54
Tagged	Non-jumbo	1522	4	1522	1468	54
	Jumbo ¹ on all VLANs	9220	4	9218	9164	54
	Jumbo ¹ On all but source VLAN	1522	4	n/a ²	1468	54

¹ Jumbo frames are allowed on ports operating at or above 1 Gbps


² For local mirroring, a non-jumbo configuration on the source VLAN dictates an MTU of 1518 bytes for untagged frames, and an MTU of 1522 for tagged frames, regardless of the jumbo configuration on any other VLANs on the switch.

Effect of downstream VLAN tagging on untagged, mirrored traffic

In a remote mirroring application, if mirrored traffic leaves the switch without 802.1Q VLAN tagging, but is forwarded through a downstream device that adds 802.1Q VLAN tags, the MTU for untagged mirrored frames leaving the source switch is reduced below the values shown in **Maximum frame sizes for mirroring**.

For example, if the MTU on the path to the destination is 1522 bytes, untagged mirrored frames leaving the source switch cannot exceed 1518 bytes. Likewise, if the MTU on the path to the destination is 9220 bytes, untagged mirrored frames leaving the source switch cannot exceed 9216 bytes.

Figure 56: Effect of downstream VLAN tagging on the MTU for mirrored traffic



Operating notes for traffic mirroring

- Mirroring dropped traffic

When an interface is configured to mirror traffic to a local or remote destination, packets are mirrored regardless of whether the traffic is dropped while on the interface. For example, if an ACL is configured on a VLAN with a `deny` ACE that eliminates packets from a Telnet application, the switch still mirrors the Telnet packets that are received on the interface and subsequently dropped.

- Mirroring and spanning tree

Mirroring is performed regardless of the STP state of a port or trunk. This means, for example, that inbound traffic on a port blocked by STP can still be monitored for STP packets during the STP setup phase.

- Tagged and untagged frames

For a frame entering or leaving the switch on a mirrored port, the mirrored copy retains the tagged or untagged state the original frame carried when it entered into or exited from the switch. (The tagged or untagged VLAN membership of ports in the path leading to the mirroring destination does not affect the tagged or untagged status of the mirrored copy itself.)

Thus, if a tagged frame arrives on a mirrored port, the mirrored copy is also tagged, regardless of the status of ports in the destination path. If a frame exits from the switch on a mirrored port that is a tagged member of a VLAN, the mirrored copy is also tagged for the same reason.

To prevent a VLAN tag from being added to the mirrored copy of an outbound packet sent to a mirroring destination, you must enter the `no-tag-added` parameter when you configure a port, trunk, or mesh interface to select mirrored traffic.

- Effect of IGMP on mirroring

If both inbound and outbound mirroring is operating when IGMP is enabled on a VLAN, two copies of mirrored IGMP frames may appear at the mirroring destination.

- Mirrored traffic not encrypted

Mirrored traffic undergoes IPv4 encapsulation, but mirrored encapsulated traffic is not encrypted.

- IPv4 header added

The IPv4 encapsulation of mirrored traffic adds a 54-byte header to each mirrored frame. If a resulting frame exceeds the maximum MTU allowed in the network, it is dropped or truncated (according to the setting of the `[truncation]` parameter in the `mirror` command.)

To reduce the number of dropped frames, enable jumbo frames in the mirroring path, including all intermediate switches and/or routers. (The MTU on the switch is 9220 bytes, which includes 4 bytes for the 802.1Q VLAN tag.)

- Intercepted or injected traffic

The mirroring feature does not protect against either mirrored traffic being intercepted or traffic being injected into a mirrored stream by an intermediate host.

- Inbound mirrored IPv4-encapsulated frames are not mirrored

The switch does not mirror IPv4-encapsulated mirrored frames that it receives on an interface. This prevents duplicate mirrored frames in configurations where the port connecting the switch to the network path for a mirroring destination is also a port whose inbound or outbound traffic is being mirrored.

For example, if traffic leaving the switch through ports B5, B6, and B7 is being mirrored through port B7 to a network analyzer, the mirrored frames from traffic on ports B5 and B6 will not be mirrored a second time as they pass through port B7.

- Switch operation as both destination and source

A switch configured as a remote destination switch can also be configured to mirror traffic to one of its own ports (local mirroring) or to a destination on another switch (remote mirroring.)

- Monitor command note

If session 1 is already configured with a destination, you can enter the `[no] vlan <VID>monitor` or `[no] interface <PORT> monitor` command without mirroring criteria and a mirror session number. In this case, the switch automatically configures or removes mirroring for inbound and outbound traffic from the specified VLAN or ports to the destination configured for session 1.

- Loss of connectivity suspends remote mirroring

When a remote mirroring session is configured on a source switch, the switch sends an ARP request to the configured destination approximately every 60 seconds. If the source switch fails to receive the expected ARP response from the destination for the session, transmission of mirrored traffic in the session halts. However, because the source switch continues to send ARP requests for each configured remote session, link restoration or discovery of another path to the destination enables the source switch to resume transmitting the session's mirrored traffic after a successful ARP response cycle occurs.

Note that if a link's connectivity is repeatedly interrupted ("link toggling"), little or no mirrored traffic may be allowed for sessions using that link. To verify the status of any mirroring session configured on the source switch, use the `show monitor` command.

Troubleshooting traffic mirroring

Cause

If mirrored traffic does not reach the configured remote destination (endpoint) switch or remote exit port, check the following configurations:

- If the destination for mirrored traffic is on a different VLAN than the source, routing must be correctly configured along the path from the source to the destination.
- All links on the path from the source switch to the destination switch must be active.



A mirroring exit port should be connected only to a network analyzer, IDS, or other network edge device that has no connection to other network resources. Configuring a mirroring exit port connection to a network can result in serious network performance problems, and is strongly discouraged by HPE Aruba.

Interface monitoring features

You can designate monitoring of inbound and outbound traffic on ports and static trunks. This allows the monitoring of individual ports, groups of contiguous ports, and static port trunks.

The switch monitors network activity by copying all traffic inbound and outbound on the specified interfaces to the designated monitoring port, to which a network analyzer can be attached.

If a tagged packet arrives on a monitored port, the packet will remain tagged when it goes out a monitored port even if that port is configured as untagged. If the packet is untagged, it will remain untagged going out the monitor port. The monitor port state (tagged or untagged) does not affect the tagging of the packet. However, egress mirroring does not reflect the tagged or untagged characteristic to the mirror port, instead it reflects the tagged or untagged characteristic of the mirror port.



When both inbound and outbound monitoring is done, and IGMP is enabled on any VLAN, you may get two copies of IGMP packets on the monitored port.



VLANs and port trunks cannot be used as a monitoring port.

The switch can monitor static LACP trunks, but not dynamic LACP trunks.

It is possible, when monitoring multiple interfaces in networks with high traffic levels, to copy more traffic to a monitor port than the link can support. In this case, some packets may not be copied to the monitor port.

Configuring port and static trunk monitoring (Menu)

This procedure describes configuring the switch for monitoring when monitoring is disabled. (If monitoring has already been enabled, the screens will appear differently than shown in this procedure.)

1. From the console Main Menu, select:**2. Switch Configuration... 3. Network Monitoring Port**
2. In the Actions menu, press **[E]** (for Edit).
3. If monitoring is currently disabled (the default) then enable it by pressing the Space bar (or **[Y]**) to select Yes.
4. Press the down arrow key to display a screen similar to the following and move the cursor to the **Monitoring Port** parameter.
5. Use the Space bar to select the port to use for monitoring.
6. Use the down arrow key to move the cursor to the **Action** column for the individual ports and position the cursor at a port you want to monitor.
7. Press the Space bar to select **Monitor** for each port and trunk that you want monitored. (Use the down arrow key to move from one interface to the next in the **Action** column.)
8. When you finish selecting ports to monitor, press **[Enter]**, then press **[S]** (for **Save**) to save your changes and exit from the screen.
9. Return to the Main Menu.

Configuring port and static trunk monitoring (CLI)

You must use the following configuration sequence to configure port and static trunk monitoring in the CLI:

1. Assign a monitoring (mirror) port.
2. Designate the port(s) and/or static trunk(s) to monitor.

Displaying the monitoring configuration

Syntax:

```
show monitor
```

This command lists the port assigned to receive monitored traffic and the ports and/or trunks being monitored. For example, if you assign port 5 as the monitoring port and configure the switch to monitor ports 2-4, `show monitor` displays the following:

Monitored port listing

```
switch(config)# show monitor

Network Monitoring Port

Mirror Port: 44

Monitoring sources 5
-----
2
3
4
```

- Port receiving monitored traffic.
- Monitored ports

Configuring the monitor port

Syntax:

```
[no] mirror-port [ <port-num> ]
```

This command assigns or removes a monitoring port, and must be executed from the global configuration level. Removing the monitor port disables port monitoring and resets the monitoring parameters to their factory-default settings.

For example, to assign port 6 as the monitoring port:

```
switch(config)# mirror-port 6
```

To turn off monitoring:

```
switch(config)# no mirror-port
```

Selecting or removing monitoring source interfaces

After you configure a monitor port you can use either the global configuration level or the interface context level to select ports or static trunks as monitoring sources. You can also use either level to remove monitoring sources.

Syntax:

```
[no] interface <monitor-list> monitor
```

<monitor-list> Includes port numbers and static trunk names such as **4, 7, 5-8, trk1**

⁴ Port receiving monitored traffic.

⁵ Monitored ports

Identifies the switch elements to monitor through the currently configured monitor port. You can monitor the ports and static trunks available on the switch.

Elements in the monitor list can include port numbers and static trunk names at the same time.

For example, with a port such as port 6 configured as the monitoring (mirror) port, you would use either of the following commands to select these interfaces for monitoring:

- 1–3, and 5
- Trunk 2

Selecting ports and static trunks as monitoring sources

```
switch(config)# int 1-3, 5, trk2, monitor
```

Disabling monitoring at the interface context and the global config level

```
switch(eth-1-3, 5)# no int 5 monitor 6  
switch(eth-1-3, 5)# no monitor
```

```
switch(config)# no int 5 monitor 7  
switch(config)# no int 1-3, 5 monitor
```

- These two commands show how to disable monitoring at the interface context level for a single port or all ports in an interface context level.
- These two commands show how to disable monitoring at the global config level for a single port or a group of ports.

⁶ These two commands show how to disable monitoring at the interface context level for a single port or all ports in an interface context level.

⁷ These two commands show how to disable monitoring at the global config level for a single port or a group of ports.

Overview

This chapter addresses performance-related network problems that can be caused by topology, switch configuration, and the effects of other devices or their configurations on switch operation. (For switch-specific information on hardware problems indicated by LED behavior, cabling requirements, and other potential hardware-related problems, see the installation guide you received with the switch.)



NOTE

HPE periodically places switch software updates on the HPE Switch Networking website. HPE recommends that you check this website for software updates that may have fixed a problem you are experiencing.

For information on support and warranty provisions, see the Support and Warranty booklet shipped with the switch.

Troubleshooting approaches

Cause

Use these approaches to diagnose switch problems:

- Check the HPE website for software updates that may have solved your problem: <http://www.hpe.com/networking>
- Check the switch LEDs for indications of proper switch operation:
 - Each switch port has a Link LED that should light whenever an active network device is connected to the port.
 - Problems with the switch hardware and software are indicated by flashing the Fault and other switch LEDs. For a description of the LED behavior and information on using the LEDs for troubleshooting, see the installation guide shipped with the switch.
- Check the network topology/installation. For topology information, see the installation guide shipped with the switch.
- Check cables for damage, correct type, and proper connections. You should also use a cable tester to check your cables for compliance to the relevant IEEE 802.3 specification. For correct cable types and connector pin-outs, see the installation guide shipped with the switch.
- Use HPE PCM+ to help isolate problems and recommend solutions.
- Use the Port Utilization Graph and Alert Log in the WebAgent included in the switch to help isolate problems. These tools are available through the WebAgent:
 - Port Utilization Graph
 - Alert log
 - Port Status and Port Counters screens
 - Diagnostic tools (Link test, Ping test, configuration file browser)
- For help in isolating problems, use the easy-to-access switch console built into the switch or Telnet to the switch console. For operating information on the Menu and CLI interfaces included in the console, see chapters 3 and 4. These tools are available through the switch console:
 - Status and Counters screens
 - Event Log
 - Diagnostics tools (Link test, Ping test, configuration file browser, and advanced user commands)

Browser or Telnet access problems

Cannot access the WebAgent

- Access may be disabled by the Web Agent Enabled parameter in the switch console. Check the setting on this parameter by selecting:

2. Switch Configuration

1. System Information

- The switch may not have the correct IP address, subnet mask, or gateway. Verify by connecting a console to the switch's Console port and selecting:

2. Switch Configuration

5. IP Configuration

Note: If DHCP/Bootp is used to configure the switch, the IP addressing can be verified by selecting:

1. Status and Counters...

2. Switch Management Address Information

Also check the DHCP/Bootp server configuration to verify correct IP addressing.

- If you are using DHCP to acquire the IP address for the switch, the IP address "lease time" may have expired so that the IP address has changed. For more information on how to "reserve" an IP address, see the documentation for the DHCP application that you are using.
- If one or more IP-authorized managers are configured, the switch allows inbound telnet access only to a device having an authorized IP address. For more information on IP Authorized managers, see the access security guide for your switch.
- Java™ applets may not be running on the web browser. They are required for the switch WebAgent to operate correctly. Refer to the online Help on your web browser for how to run the Java applets.

Cannot Telnet into the switch console from a station on the network

- Off-subnet management stations can lose Telnet access if you enable routing without first configuring a static (default) route. That is, the switch uses the IP default gateway only while operating as a Layer 2 device. While routing is enabled on the switch, the IP default gateway is not used. You can avoid this problem by using the ip route command to configure a static (default) route before enabling routing. For more information, see "IP Routing Features" in the multicast and routing guide for your switch.
- Telnet access may be disabled by the `Inbound Telnet Enabled` parameter in the System Information screen of the menu interface:

2. Switch Configuration

1. System Information

- The switch may not have the correct IP address, subnet mask, or gateway. Verify by connecting a console to the switch's Console port and selecting:

2. Switch Configuration

5. IP Configuration

- If you are using DHCP to acquire the IP address for the switch, the IP address "lease time" may have expired so that the IP address has changed. For more information on how to "reserve" an IP address, see the documentation for the DHCP application that you are using.
- If one or more IP-authorized managers are configured, the switch allows inbound telnet access only to a device having an authorized IP address. For more information on IP Authorized managers, see the access security guide for your switch.

Unusual network activity

Network activity that fails to meet accepted norms may indicate a hardware problem with one or more of the network components, possibly including the switch. Such problems can also be caused by a network loop or simply too much traffic for the network as it is currently designed and implemented. Unusual network activity is usually indicated by the LEDs on the front of the switch or measured with the switchconsole interface or with a network management tool such as HPE PCM+. For information on using LEDs to identify unusual network activity, see the installation guide you received with the switch.

A topology loop can also cause excessive network activity. The Event Log "FFI" messages can be indicative of this type of problem.

General problems

The network runs slow; processes fail; users cannot access servers or other devices

Broadcast storms may be occurring in the network. These may be caused by redundant links between nodes.

- If you are configuring a port trunk, finish configuring the ports in the trunk before connecting the related cables. Otherwise you may inadvertently create a number of redundant links (that is, topology loops) that will cause broadcast storms.
- Turn on STP to block redundant links
- Check for FFI messages in the Event Log

Duplicate IP addresses

This is indicated by this Event Log message:

```
ip: Invalid ARP source: IP address on IP address
```

where both instances of *IP address* are the same address, indicating that the switch's IP address has been duplicated somewhere on the network.

Duplicate IP addresses in a DHCP network

If you use a DHCP server to assign IP addresses in your network, and you find a device with a valid IP address that does not appear to communicate properly with the server or other devices, a duplicate IP address may have been issued by the server. This can occur if a client has not released a DHCP-assigned IP address after the intended expiration time and the server "leases" the address to another device. This can also happen, for example, if the server is first configured to issue IP addresses with an unlimited duration, and then is subsequently configured to issue IP addresses that will expire after a limited duration. One solution is to configure "reservations" in the DHCP server for specific IP addresses to be assigned to devices having specific MAC addresses. For more information, see the documentation for the DHCP server.

One indication of a duplicate IP address in a DHCP network is this Event Log message:

```
ip: Invalid ARP source: <IP-address>  
    on <IP-address>
```

where both instances of *IP-address* are the same address, indicating that the IP address has been duplicated somewhere on the network.

The switch has been configured for DHCP/Bootp operation, but has not received a DHCP or Bootp reply

When the switch is first configured for DHCP/Bootp operation, or if it is rebooted with this configuration, it immediately begins sending request packets on the network. If the switch does not receive a reply to its DHCP/Bootp requests, it continues to periodically send request packets, but with decreasing frequency. Thus, if a DHCP

or Bootp server is not available or accessible to the switch when DHCP/Bootp is first configured, the switch may not immediately receive the desired configuration.

After verifying that the server has become accessible to the switch, reboot the switch to re-start the process.

802.1Q Prioritization problems

Ports configured for non-default prioritization (level 1 to 7) are not performing the specified action

If the ports were placed in a trunk group after being configured for non-default prioritization, the priority setting was automatically reset to zero (the default). Ports in a trunk group operate only at the default priority setting.

Addressing ACL problems

ACLs are properly configured and assigned to VLANs, but the switch is not using the ACLs to filter IP layer 3 packets

Procedure

1. The switch may be running with IP routing disabled. To ensure that IP routing is enabled, execute `show running` and look for the IP routing statement in the resulting listing. For Example:

Indication that routing is enabled

```
switch(config)# show running
Running configuration:
; J9091A Configuration Editor; Created on release #XX.15.06
hostname " HPswitch "
ip default-gateway 10.33.248.1
ip routing 1
logging 10.28.227.2
snmp-server community "public" Unrestricted
ip access-list extended "Controls for VLAN 20"
permit tcp 0.0.0.0 255.255.255.255 10.10.20.98 0.0.0.0 eq 80
permit tcp 0.0.0.0 255.255.255.255 10.10.20.21 0.0.0.0 eq 80
deny tcp 0.0.0.0 255.255.255.255 10.10.20.1 0.0.0.255 eq 80
deny tcp 10.10.20.1? 0.0.0.0 10.10.10.100 0.0.0.0 eq 20 log
deny tcp 10.10.20.20 0.0.0.0 10.10.10.100 0.0.0.0 eq 20 log
deny tcp 10.10.20.43 0.0.0.0 10.10.10.100 0.0.0.0 eq 20 log
permit ip 10.10.20.1 0.0.0.255 10.10.10.100 0.0.0.0
deny ip 10.10.30.1 0.0.0.255 10.10.10.100 0.0.0.0
permit ip 10.10.30.1 0.0.0.255 10.10.10.1 0.0.0.255
exit
```

- 1

Indicates that routing is enabled, a requirement for ACL operation. (There is an exception. Refer to the **Note**, below.)



If an ACL assigned to a VLAN includes an ACE referencing an IP address on the switch itself as a packet source or destination, the ACE screens traffic to or from this switch address regardless of whether IP routing is enabled. This is a security measure designed to help protect the switch from unauthorized management access.

If you need to configure IP routing, execute the `ip routing` command.

2. ACL filtering on the switches applies only to routed packets and packets having a destination IP address (DA) on the switch itself.

Also, the switch applies assigned ACLs only at the point where traffic enters or leaves the switch on a VLAN. Ensure that you have correctly applied your ACLs ("in" and/or "out") to the appropriate VLANs.

The switch does not allow management access from a device on the same VLAN

The implicit `deny any` function that the switch automatically applies as the last entry in any ACL always blocks packets having the same DA as the switch's IP address on the same VLAN. That is, bridged packets with the switch itself as the destination are blocked as a security measure.

To preempt this action, edit the ACL to include an ACE that permits access to the switch's DA on that VLAN from the management device.

Error (Invalid input) when entering an IP address

When using the "host" option in the Command syntax, ensure that you are not including a mask in either dotted decimal or CIDR format. Using the "host" option implies a specific host device and therefore does not permit any mask entry.

Correctly and incorrectly specifying a single host

```
Switch(config)# access-list 6 permit host 10.28.100.100 1
Switch(config)# access-list 6 permit host 10.28.100.100 255.255.255.2552
Invalid input: 255.255.255.255
Switch(config)# access-list 6 permit host 10.28.100.100/32 3
Invalid input: 10.28.100.100/32
```

- ¹Correct.
- ²Incorrect. No mask needed to specify a single host.
- ³Incorrect. No mask needed to specify a single host.

Apparent failure to log all "deny" matches

Where the `log` statement is included in multiple ACEs configured with a "deny" option, a large volume of "deny" matches generating logging messages in a short period of time can impact switch performance. If it appears that the switch is not consistently logging all "deny" matches, try reducing the number of logging actions by removing the `log` statement from some ACEs configured with the "deny" action.

The switch does not allow any routed access from a specific host, group of hosts, or subnet

The implicit `deny any` function that the switch automatically applies as the last entry in any ACL may be blocking all access by devices not specifically permitted by an entry in an ACL affecting those sources. If you are using the ACL to block specific hosts, a group of hosts, or a subnet, but want to allow any access not specifically permitted, insert `permit any` as the last explicit entry in the ACL.

The switch is not performing routing functions on a VLAN

Two possible causes of this problem are:

- Routing is not enabled. If `show running` indicates that routing is not enabled, use the `ip routing` command to enable routing.
- An ACL may be blocking access to the VLAN (on a switch covered in this guide). Ensure that the switch's IP address on the VLAN is not blocked by one of the ACE entries in an ACL applied to that VLAN. A common mistake is to either not explicitly permit the switch's IP address as a DA or to use a wildcard ACL mask in a `deny` statement that happens to include the switch's IP address. For an Example: of this problem, see section

"General ACL Operating Notes" in the "Access Control Lists (ACLs)" of the latest access security guide for your switch.

Routing through a gateway on the switch fails

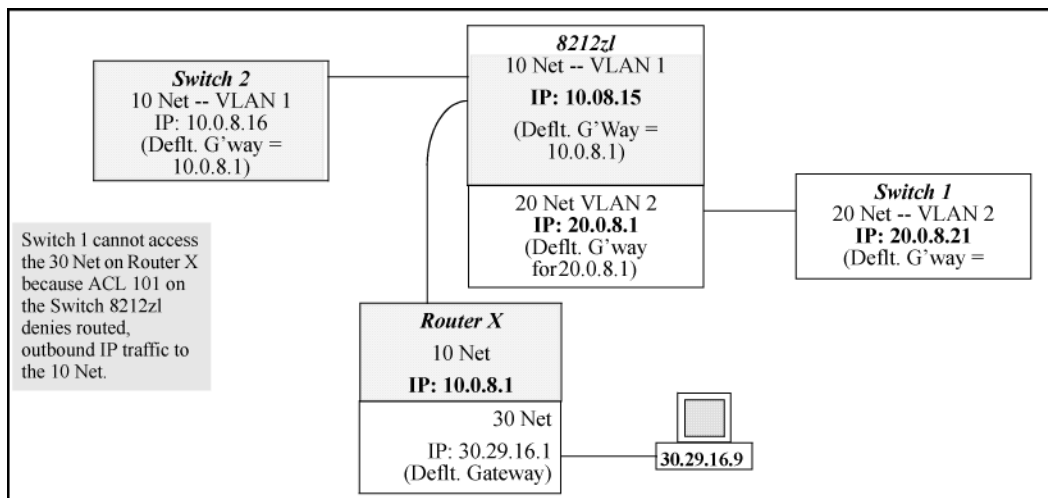
Configuring a "deny" ACE that includes a gateway address can block traffic attempting to use the gateway as a next-hop.

Remote gateway case

Configuring ACL "101" (example below) and applying it outbound on VLAN 1 in the figure below includes the router gateway (10.0.8.1) needed by devices on other networks. This can prevent the switch from sending ARP and other routing messages to the gateway router to support traffic from authorized remote networks.

In **Figure 57: Inadvertently blocking a gateway** on page 297, this ACE (see data in bold below) denies access to the 10 Net's 10.0.8.1 router gateway needed by the 20 Net (Subnet mask is 255.255.255.0). **See: example**

Figure 57: *Inadvertently blocking a gateway*



To avoid inadvertently blocking the remote gateway for authorized traffic from another network (such as the 20 Net in this Example):

Procedure

1. Configure an ACE that specifically permits authorized traffic from the remote network.
2. Configure narrowly defined ACEs to block unwanted IP traffic that would otherwise use the gateway; such ACEs might deny traffic for a particular application, particular hosts, or an entire subnet.
3. Configure a "permit any" ACE to specifically allow any IP traffic to move through the gateway.

ACE blocking an entire subnet

```
switch(config)# access-list config
ip access-list extended "101"
  deny ip 0.0.0.0 255.255.255.255 10.0.8.30 0.0.0.255
  permit ip 0.0.0.0 255.255.255.255 0.0.0.00 255.255.255.255
exit
```

Local gateway case

If you use the switch as a gateway for traffic you want routed between subnets, use these general steps to avoid blocking the gateway for authorized applications:

Procedure

1. Configure gateway security first for routing with specific permit and deny statements.
2. Permit authorized traffic.
3. Deny any unauthorized traffic that you have not already denied in step [1](#) on page 298.

IGMP-related problems

IP multicast (IGMP) traffic that is directed by IGMP does not reach IGMP hosts or a multicast router connected to a port

IGMP must be enabled on the switch and the affected port must be configured for "Auto" or "Forward" operation.

IP multicast traffic floods out all ports; IGMP does not appear to filter traffic

The IGMP feature does not operate if the switch or VLAN does not have an IP address configured manually or obtained through DHCP/Bootp. To verify whether an IP address is configured for the switch or VLAN, do one of the following:

- **Try using the WebAgent:** If you can access the WebAgent, then an IP address is configured.
- **Try to telnet to the switch console:** If you can Telnet to the switch, an IP address is configured.
- **Use the switch console interface:** From the Main Menu, check the Management Address Information screen by clicking on:
 1. Status and Counters
 2. Switch Management Address Information

LACP-related problems

Unable to enable LACP on a port with the `interface <port-number> lacp` command

In this case, the switch displays the following message:

Operation is not allowed for a trunked port.

You cannot enable LACP on a port while it is configured as a static Trunk port. To enable LACP on a static-trunked port:

Procedure

1. Use the `no trunk <port-number>` command to disable the static trunk assignment.
2. Execute `interface <port-number> lacp` .



Removing a port from a trunk without first disabling the port can create a traffic loop that can slow down or halt your network. Before removing a port from a trunk, Hewlett Packard Enterprise recommends that you either disable the port or disconnect it from the LAN.

Port-based access control (802.1X)-related problems



To list the 802.1X port-access Event Log messages stored on the switch, use `show log 802`.

See also [Radius-related problems](#) on page 301.

The switch does not receive a response to RADIUS authentication requests

In this case, the switch attempts authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use `ping` to ensure that the switch has access to the configured RADIUS servers.
- Verify that the switch is using the correct encryption key (RADIUS secret key) for each server.
- Verify that the switch has the correct IP address for each RADIUS server.
- Ensure that the `radius-server timeout` period is long enough for network conditions.

The switch does not authenticate a client even though the RADIUS server is properly configured and providing a response to the authentication request

If the RADIUS server configuration for authenticating the client includes a VLAN assignment, ensure that the VLAN exists as a static VLAN on the switch. See "How 802.1X Authentication Affects VLAN Operation" in the access security guide for your switch.

During RADIUS-authenticated client sessions, access to a VLAN on the port used for the client sessions is lost

If the affected VLAN is configured as untagged on the port, it may be temporarily blocked on that port during an 802.1X session. This is because the switch has temporarily assigned another VLAN as untagged on the port to support the client access, as specified in the response from the RADIUS server. See "How 802.1X Authentication Affects VLAN Operation" in the access security guide for your switch.

The switch appears to be properly configured as a supplicant, but cannot gain access to the intended authenticator port on the switch to which it is connected

If `aaa authentication port-access` is configured for Local, ensure that you have entered the local **login** (operator-level) username and password of the authenticator switch into the `identity` and `secret` parameters of the supplicant configuration. If instead, you enter the enable (manager-level) username and password, access will be denied.

The supplicant statistics listing shows multiple ports with the same authenticator MAC address

The link to the authenticator may have been moved from one port to another without the supplicant statistics having been cleared from the first port. See "Note on Supplicant Statistics" in the chapter on Port-Based and User-Based Access Control in the access security guide for your switch.

The `show port-access authenticator <port-list>` command shows one or more ports remain open after they have been configured with `control unauthorized`

802.1X is not active on the switch. After you execute `aaa port-access authenticator active`, all ports configured with `control unauthorized` should be listed as Closed.

Authenticator ports remain "open" until activated

```
switch(config)# show port-access authenticator e 9
Port Access Authenticator Status
  Port-access authenticator activated [No] : No
      Access Authenticator Authenticator
Port Status Control  State Backend  State
-----
  9      Open  1    FU           Force Auth  Idle
```

```
Switch(config)# show port-access authenticator active
Switch(config)# show port-access authenticator e 9
Port Access Authenticator Status
Port-access authenticator activated [No] : Yes
Access Authenticator Authenticator
Port Status Control State Backend State
-----
9 Closed FU Force Unauth Idle
```

- ¹Port A9 shows an “Open” status even though Access Control is set to Unauthorized (Force Auth). This is because the port-access authenticator has not yet been activated.

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch

Use `show radius` to verify that the encryption key (RADIUS secret key) the switch is using is correct for the server being contacted. If the switch has only a global key configured, it either must match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, it overrides the global key and must match the server key.

Displaying encryption keys

```
switch(config)# show radius
Status and Counters - General RADIUS Information
Deadtime(min) : 0
Timeout(secs) : 5
Retransmit Attempts : 3
Global Encryption Key : My-Global-Key
Dynamic Authorization UDP Port : 3799

Server IP Addr Auth Acct DM/ Time
Port Port CoA Window Encryption Key
-----
10.33.18.119 1812 1813 119-only-key
```

Also, ensure that the switch port used to access the RADIUS server is not blocked by an 802.1X configuration on that port. For example, `show port-access authenticator <port-list>` gives you the status for the specified ports. Also, ensure that other factors, such as port security or any 802.1X configuration on the RADIUS server are not blocking the link.

The authorized MAC address on a port that is configured for both 802.1X and port security either changes or is re-acquired after execution of `aaa port-access authenticator <port-list> initialize`

If the port is force-authorized with `aaa port-access authenticator <port-list> control authorized` command and port security is enabled on the port, then executing `initialize` causes the port to clear the learned address and learn a new address from the first packet it receives after you execute `initialize`.

A trunked port configured for 802.1X is blocked

If you are using RADIUS authentication and the RADIUS server specifies a VLAN for the port, the switch allows authentication, but blocks the port. To eliminate this problem, either remove the port from the trunk or reconfigure the RADIUS server to avoid specifying a VLAN.

QoS-related problems

Loss of communication when using VLAN-tagged traffic

If you cannot communicate with a device in a tagged VLAN environment, ensure that the device either supports VLAN tagged traffic or is connected to a VLAN port that is configured as `Untagged`.

Radius-related problems

The switch does not receive a response to RADIUS authentication requests

In this case, the switch attempts authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use `ping` to ensure that the switch has access to the configured RADIUS server.
- Verify that the switch is using the correct encryption key for the designated server.
- Verify that the switch has the correct IP address for the RADIUS server.
- Ensure that the `radius-server timeout` period is long enough for network conditions.
- Verify that the switch is using the same UDP port number as the server.



Because of an inconsistency between the Windows XP 802.1x supplicant timeout value and the switch default timeout value, which is 5, when adding a backup RADIUS server, set the switch `radius-server timeout` value to 4. Otherwise, the switch may not failover properly to the backup RADIUS server.

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch

Use `show radius` to verify that the encryption key the switch is using is correct for the server being contacted. If the switch has only a global key configured, it either must match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, it overrides the global key and must match the server key.

Global and unique encryption keys

```
Switch(config)# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key : My-Global-Key 1
  Dynamic Authorization UDP Port : 3799

  Server IP Addr      Auth Port  Acct Port  DM/CoA  Time Window  Encryption Key
  -----
  10.33.18.119       1812     1813      -        -        -        -        119-only-key 2
```

- 1

Global RADIUS Encryption Key

- 2

Unique RADIUS Encryption Key for the RADIUS server at 10.33.18.119

MSTP and fast-uplink problems



If you enable MSTP, Hewlett Packard Enterprise recommends that you leave the remainder of the MSTP parameter settings at their default values until you have had an opportunity to evaluate MSTP performance in your network. Because incorrect MSTP settings can adversely affect network performance, you should avoid making changes without having a strong understanding of how MSTP operates. To learn the details of MSTP operation, see the IEEE802.1s standard.

Broadcast storms appearing in the network

This can occur when there are physical loops (redundant links) in the topology. Where this exists, you should enable MSTP on all bridging devices in the topology to detect the loop.

STP blocks a link in a VLAN even though there are no redundant links in that VLAN

In 802.1Q-compliant switches, MSTP blocks redundant physical links even if they are in separate VLANs. A solution is to use only one, multiple-VLAN (tagged) link between the devices. Also, if ports are available, you can improve the bandwidth in this situation by using a port trunk. See "Spanning Tree Operation with VLANs" in "Static Virtual LANs (VLANs)" in the advanced traffic management guide for your switch.

Fast-uplink troubleshooting

Some of the problems that can result from incorrect use of fast-uplink MSTP include temporary loops and generation of duplicate packets.

Problem sources can include:

- Fast-uplink is configured on a switch that is the MSTP root device.
- Either the `Hello Time` or the `Max Age` setting (or both) is too long on one or more switches. Return the `Hello Time` and `Max Age` settings to their default values (2 seconds and 20 seconds, respectively, on a switch).
- A "downlink" port is connected to a switch that is further away (in hop count) from the root device than the switch port on which fast-uplink MSTP is configured.
- Two edge switches are directly linked to each other with a fast-uplink (Mode = `Uplink`) connection.
- Fast uplink is configured on both ends of a link.
- A switch serving as a backup MSTP root switch has ports configured for fast-uplink MSTP and has become the root device because of a failure in the original root device.

SSH-related problems

Switch access refused to a client

Even though you have placed the client's public key in a text file and copied the file (using the `copy tftp pub-key-file` command) into the switch, the switch refuses to allow the client to have access. If the source SSH client is an SSHv2 application, the public key may be in the PEM format, which the switch (SSHv1) does not interpret. Check the SSH client application for a utility that can convert the PEM-formatted key into an ASCII-formatted key.

Executing IP SSH does not enable SSH on the switch

The switch does not have a host key. Verify by executing `show ip host-public-key`. If you see the message

```
ssh cannot be enabled until a host key is configured (use 'crypto' command).
```

you need to generate an SSH key pair for the switch. To do so, execute `crypto key generate` (see "Generating the switch's public and private key pair" in the SSH chapter of the access security guide for your switch.)

Switch does not detect a client's public key that does appear in the switch's public key file (`show ip client-public-key`)

The client's public key entry in the public key file may be preceded by another entry that does not terminate with a new line (CR). In this case, the switch interprets the next sequential key entry as simply a comment attached to the preceding key entry. Where a public key file has more than one entry, ensure that all entries terminate with a new line (CR). While this is optional for the last entry in the file, not adding a new line to the last entry creates an error potential if you either add another key to the file at a later time or change the order of the keys in the file.

An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages

Download failed: overlength key in key file.

Download failed: too many keys in key file.

Download failed: one or more keys is not a valid RSA public key.

The public key file you are trying to download has one of the following problems:

- A key in the file is too long. The maximum key length is 1024 characters, including spaces. This could also mean that two or more keys are merged together instead of being separated by a <CR> <LF>.
- There are more than ten public keys in the key file.
- One or more keys in the file is corrupted or is not a valid rsa public key.

Client ceases to respond ("hangs") during connection phase

The switch does not support data compression in an SSH session. Clients often have compression turned on by default, but then disable it during the negotiation phase. A client that does not recognize the compression-request FAILURE response may fail when attempting to connect. Ensure that compression is turned **off** before attempting a connection to prevent this problem.

TACACS-related problems

Event Log

When troubleshooting TACACS+ operation, check the switch's Event Log for indications of problem areas.

All users are locked out of access to the switch

If the switch is functioning properly, but no username/password pairs result in console or Telnet access to the switch, the problem may be caused by how the TACACS+ server and/or the switch are configured. Use one of the following methods to recover:

- Access the TACACS+ server application and adjust or remove the configuration parameters controlling access to the switch.
- If the above method does not work, try eliminating configuration changes in the switch that have not been saved to flash (boot-up configuration) by causing the switch to reboot from the boot-up configuration (which includes only the configuration changes made prior to the last `write memory` command.) If you did not use `write memory` to save the authentication configuration to flash, pressing the `Reset` button reboots the switch with the boot-up configuration.

- Disconnect the switch from network access to any TACACS+ servers and then log in to the switch using either Telnet or direct console port access. Because the switch cannot access a TACACS+ server, it defaults to local authentication. You can then use the switch's local Operator or Manager username/password pair to log on.
- As a last resort, use the `Clear/Reset` button combination to reset the switch to its factory default boot-up configuration. Taking this step means you will have to reconfigure the switch to return it to operation in your network.

No communication between the switch and the TACACS+ server application

If the switch can access the server device (that is, it can `ping` the server), a configuration error may be the problem. Some possibilities include:

- The server IP address configured with the switch's `tacacs-serverhost` command may not be correct. (Use the switch's `show tacacs-server` command to list the TACACS+ server IP address.)
- The encryption key configured in the server does not match the encryption key configured in the switch (by using the `tacacs-server key` command). Verify the key in the server and compare it to the key configured in the switch. (Use `show tacacs-server` to list the global key. Use `show config` or `show config running` to list any server-specific keys.)
- The accessible TACACS+ servers are not configured to provide service to the switch.

Access is denied even though the username/password pair is correct

Some reasons for denial include the following parameters controlled by your TACACS+ server application:

- The account has expired.
- The access attempt is through a port that is not allowed for the account.
- The time quota for the account has been exhausted.
- The time credit for the account has expired.
- The access attempt is outside of the time frame allowed for the account.
- The allowed number of concurrent logins for the account has been exceeded.

For more help, see the documentation provided with your TACACS+ server application.

Unknown users allowed to login to the switch

Your TACACS+ application may be configured to allow access to unknown users by assigning them the privileges included in a *default user* profile. See the documentation provided with your TACACS+ server application.

System allows fewer login attempts than specified in the switch configuration

Your TACACS+ server application may be configured to allow fewer login attempts than you have configured in the switch with the `aaa authentication num-attempts` command.

TimeP, SNTP, or Gateway problems

The switch cannot find the time server or the configured gateway

TimeP, SNTP, and Gateway access are through the primary VLAN, which in the default configuration is the `DEFAULT_VLAN`. If the primary VLAN has been moved to another VLAN, it may be disabled or does not have ports assigned to it.

VLAN-related problems

Monitor port

When using the monitor port in a multiple-VLAN environment, the switch handles broadcast, multicast, and unicast traffic output from the monitor port as follows:

- If the monitor port is configured for tagged VLAN operation on the same VLAN as the traffic from monitored ports, the traffic output from the monitor port carries the same VLAN tag.
- If the monitor port is configured for untagged VLAN operation on the same VLAN as the traffic from the monitored ports, the traffic output from the monitor port is untagged.
- If the monitor port is not a member of the same VLAN as the traffic from the monitored ports, traffic from the monitored ports does not go out the monitor port.

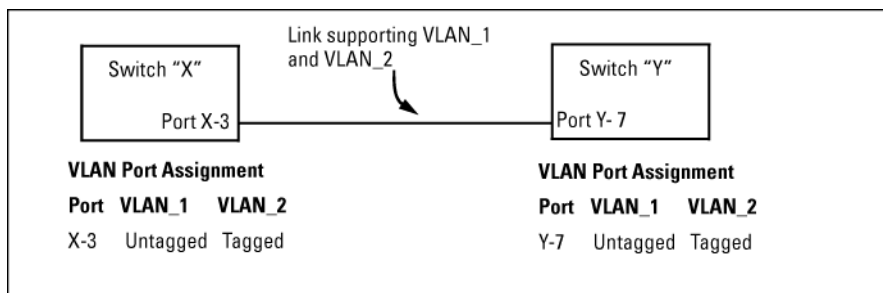
None of the devices assigned to one or more VLANs on an 802.1Q-compliant switch are being recognized

If multiple VLANs are being used on ports connecting 802.1Q-compliant devices, inconsistent VLAN IDs may have been assigned to one or more VLANs. For a given VLAN, the same VLAN ID must be used on all connected 802.1Q-compliant devices.

Link configured for multiple VLANs does not support traffic for one or more VLANs

One or more VLANs may not be properly configured as "Tagged" or "Untagged." A VLAN assigned to a port connecting two 802.1Q-compliant devices must be configured the same on both ports. For example, VLAN_1 and VLAN_2 use the same link between switch "X" and switch "Y," as shown in **Figure 58: Example: of correct VLAN port assignments on a link** on page 305.

Figure 58: Example: of correct VLAN port assignments on a link



- If VLAN_1 (VID=1) is configured as "Untagged" on port 3 on switch "X," it must also be configured as "Untagged" on port 7 on switch "Y." Make sure that the VLAN ID (VID) is the same on both switches.
- Similarly, if VLAN_2 (VID=2) is configured as "Tagged" on the link port on switch "A," it must also be configured as "Tagged" on the link port on switch "B." Make sure that the VLAN ID (VID) is the same on both switches.


Duplicate MAC addresses across VLANs

The switches operate with multiple forwarding databases. Thus, duplicate MAC addresses occurring on different VLANs can appear where a device having one MAC address is a member of more than one 802.1Q VLAN, and the switch port to which the device is linked is using VLANs (instead of MSTP or trunking) to establish redundant links to another switch. If the other device sends traffic over multiple VLANs, its MAC address consistently appears in multiple VLANs on the switch port to which it is linked.

Be aware that attempting to create redundant paths through the use of VLANs causes problems with some switches. One symptom is that a duplicate MAC address appears in the Port Address Table of one port and then later appears on another port. While the switches have multiple forwarding databases and thus do not have this problem, some switches with a single forwarding database for all VLANs may produce the impression that a connected device is moving among ports because packets with the same MAC address but different VLANs are

received on different ports. You can avoid this problem by creating redundant paths using port trunks or spanning tree.

Figure 59: Example: of duplicate MAC address



Fan failure

When two or more fans fail, a two-minute timer starts. After two minutes, the switch is powered down and must be rebooted to restart it. This protects the switch from possible overheating.

Hewlett Packard Enterprise recommends that you replace a failed fan tray assembly within one minute of removing it.

Viewing transceiver information

This feature provides the ability to view diagnostic monitoring information for transceivers with Diagnostic Optical Monitoring (DOM) support. The following table indicates the support level for specific transceivers:

Product #	Description	Support ¹
J8436A	10GbE X2-SC SR Optic	V
J8437A	10GbE X2-SC LR Optic	V
J8440B	10GbE X2-CX4 Xcver	NA
J8440C	10GbE X2-CX4 Xcver	NA
J4858A	Gigabit-SX-LC Mini-GBIC	V
J4858B	Gigabit-SX-LC Mini-GBIC	V
J4858C	Gigabit-SX-LC Mini-GBIC	V (some)
J9054B	100-FX SFP-LC Transceiver	N

Table Continued

Product #	Description	Support ¹
J8177C	Gigabit 1000Base-T Mini-GBIC	NA
J9150A	10GbE SFP+ SR Transceiver	D
J9151A	10GbE SFP+ LR Transceiver	D
J9152A	10GbE SFP+ LRM Transceiver	D
J9153A	10GbE SFP+ ER Transceiver	D
J9144A	10GbE X2-SC LRM Transceiver	D
J8438A	10Gbe X2-SC ER Transceiver	D

¹ Support indicators:

- V - Validated to respond to DOM requests
- N - No support of DOM
- D - Documented by the component suppliers as supporting DOM
- NA - Not applicable to the transceiver (copper transceiver)



Not all transceivers support Digital Optical Monitoring. If DOM appears in the Diagnostic Support field of the `show interfaces transceiver detail` command, or the `hpicfTransceiverMIB` `hpicfXcvrDiagnostics` MIB object, DOM is supported for that transceiver.

Viewing information about transceivers (CLI)

Syntax:

```
show interfaces transceiver [port-list] [detail]
```

Displays information about the transceivers. If a port is specified, displays information for the transceiver in that port.

[detail]	Displays detailed transceiver information.
----------	--

MIB support

The `hpicfTransceiver` MIB is available for displaying transceiver information.

Viewing transceiver information

The transceiver information displayed depends on the `show` command executed.

The output for `show interfaces transceiver [port-list]` is shown below. You can specify multiple ports, separated by commas, and the information for each transceiver will display.

Output for a specified transceiver

```
switch(config)# show interfaces transceiver 21
```

Transceiver Technical information:

Port	Type	Product Number	Serial Number	Part Number
21	1000SX	J4858C	MY050VM9WB	1990-3657

If there is no transceiver in the port specified in the command, the output displays as shown below.

Output when no transceiver is present in specified interface

```
switch(config)# show interfaces transceiver 22
```

No Transceiver found on interface 22

When no ports are specified, information for all transceivers found is displayed.

Output when no ports are specified

```
switch(config)# show interfaces transceiver
```

Transceiver Technical information:

Port	Type	Product Number	Serial Number	Part Number
21	1000SX	J4858C	MY050VM9WB	1990-3657
22	1000SX	J4858B	P834DIP2	

You can specify `all` for `port-list` as shown below.

Output when "all" is specified

```
switch(config)# show interfaces transceiver all
```

No Transceiver found on interface 1

No Transceiver found on interface 2

.

.

.

No Transceiver found on interface 24

Transceiver Technical information:

Port	Type	Product Number	Serial Number	Part Number
21	1000SX	J4858C	MY050VM9WB	1990-3657
22	1000SX	J4858B	P834DIP2	

Information displayed with the detail parameter

When the `show interfaces transceiver [port-list] detail` command is executed, the following information displays.

Table 23: General transceiver information

Parameter	Description
Interface Index	The switch interface number
Transceiver-type	Pluggable transceiver type
Transceiver model	Pluggable transceiver model
Connector-type	Type of connector of the transceiver
Wavelength	For an optical transceiver: the central wavelength of the laser sent, in nm. If the transceiver supports multiple wavelengths, the values will be separated by a comma.
Transfer Distance	Link-length supported by the transceiver in meters. The corresponding transfer medium is shown in brackets following the transfer distance value, For example, 50um multimode fiber. If the transceiver supports multiple transfer media, the values are separated by a comma.
Diagnostic Support	Shows whether the transceiver supports diagnostics: None Supported DOM Supported VCT Supported
Serial Number	Serial number of the transceiver

The information in the next three tables is only displayed when the transceiver supports DOM.

Table 24: DOM information

Parameter	Description
Temperature	Transceiver temperature (in degrees Centigrade)
Voltage	Supply voltage in transceiver (Volts)
Bias	Laser bias current (mA)
RX power	Rx power (mW and dBm))
TX power	Tx power (mW and dBm)

The alarm information for GBIC/SFP transceivers is shown in this table.

Table 25: Alarm and error information (GBIC/SFP transceivers only)

Alarm	Description
RX loss of signal	Incoming (RX) signal is lost
RX power high	Incoming (RX) power level is high
RX power low	Incoming (RX) power level is low
TX fault	Transmit (TX) fault
TX bias high	TX bias current is high
TX bias low	TX bias current is low
TX power high	TX power is high
TX power low	TX power is low
Temp high	Temperature is high
Temp low	Temperature is low
Voltage High	Voltage is high
Voltage Low	Voltage is low

The alarm information for XENPAK transceivers is shown in this table.

Table 26: Alarm and error information (XENPAK transceivers)

Alarm	Description
WIS local fault	WAN Interface Sublayer local fault
Receive optical power fault	Receive optical power fault
PMA/PMD receiver local fault	Physical Medium Attachment/Physical Medium Dependent receiver local fault
PCS receiver local fault	Physical Coding Sublayer receiver local fault
PHY XS receive local fault	PHY Extended Sublayer receive local fault
RX power high	RX power is high
RX power low	RX power is low
Laser bias current fault	Laser bias current fault

Table Continued

Alarm	Description
Laser temperature fault	Laser temperature fault
Laser output power fault	Laser output power fault
TX fault	TX fault
PMA/PMD transmitter local fault	PMA/PMD transmitter local fault
PCS Transmit local fault	PCS transmit local fault
PHY XS transmit local fault	PHY SX transmit local fault
TX bias high	TX bias current is high
TX bias low	TX bias current is low
TX power high	TX power is high
TX power low	TX power is low
Temp high	Temperature is high
Temp low	Temperature is low

An Example: of the output for the show interfaces transceiver [port-list] detail for a 1000SX transceiver is shown below.

Detailed information for a 1000SX Mini-GBIC transceiver

```
switch(config)# show interfaces transceiver 21 detail
```

```
Transceiver in 21
Interface index      : 21
Type                 : 1000SX
Model                : J4858C
Connector type      : LC
Wavelength           : 850nm
Transfer distance    : 300m (50um), 150m (62.5um),
Diagnostic support   : DOM
Serial number        : MY050VM9WB

Status
Temperature          : 50.111C
Voltage              : 3.1234V
TX Bias              : 6mA
TX Power             : 0.2650mW, -5.768dBm
RX Power             : 0.3892mW, -4.098dBm
```

```
Time stamp          : Mon Mar 7 14:22:13 2011
```

An Example: of the output for a 10GbE-LR transceiver is shown below.

Detailed information for a 10GbE-LR transceiver

```
switch(config)# show interfaces transceiver 23 detail
```

```
Transceiver in 23
Interface Index   : 24
Type              : 10GbE-LR
Model             : J8437A
Connector type    : SC
Wavelength        : Channel #0: 1310nm, #1:0nm, #2:0nm, #3:0nm
Transfer distance : 10000m (SM)
Diagnostic support: DOM
Serial number      : ED456SS987
```

```
Status
Temperature       : 32.754C
TX Bias           : 42.700mA
TX Power          : 0.5192mW, -2.847dBm
RX Power          : 0.0040mW, -23.979dBm
```

Recent Alarms:

```
Rx power low alarm
Rx power low warning
```

Recent errors:

```
Receive optical power fault
PMA/PMD receiver local fault
PMA/PMD transmitter local fault
PCS receive local fault
PHY XS transmit local fault
```

```
Time stamp : Mon Mar 7 16:26:06 2013
```

Viewing transceiver information for copper transceivers with VCT support

This feature provides the ability to view diagnostic monitoring information for copper transceivers with Virtual Cable Test (VCT) support. The cable quality of the copper cables connected between transceivers can be ascertained using the transceiver cable diagnostics. Results of the diagnostics are displayed with the appropriate CLI show commands and with SNMP using the hpicfTransceiver MIB.

The J8177C 1000Base-T Mini-GBIC is supported.

Testing the Cable

Enter the `test cable-diagnostics` command in any context to begin cable diagnostics for the transceiver. The diagnostic attempts to identify cable faults. The tests may take a few seconds to complete for each interface. There is the potential of link loss during the diagnostic.

Syntax:

```
test cable-diagnostics [port-list]
```

Invokes cable diagnostics and displays the results.

Output from test cable-diagnostics command

```
HP Switch # test cable-diagnostics a23-a24
```


The 'test cable-diagnostics' command will cause a loss of link and will take a few seconds per interface to complete.

Continue (Y/N)? y

MDI Port	Cable Pair	Distance Status	Pair to Fault	Pair Skew	MDI Polarity	MDI Mode
A23	1-2	OK	0 m	6 ns	Normal	MDIX
	3-6	OK	0 m	0 ns	Normal	
	4-5	OK	0 m	6 ns	Normal	MDIX
	7-8	OK	0 m	6 ns	Normal	
A24	1-2	Short	2 m			
	3-6	Impedance	3 m			
	4-5	Impedance	3 m			
	7-8	Open	1 m			

Copper cable diagnostic test results

```
switch# show interfaces transceiver a23 detail
```

```
Transceiver in A23
Interface Index   : 23
Type              : 1000T-sfp
Model             : J8177C
Connector Type    : RJ45
Wavelength        : n/a
Transfer Distance : 100m (copper),
Diagnostic Support : VCT
Serial Number     : US051HF099

Link Status       : Up
Speed             : 1000
Duplex            : Full
```

Port	MDI Pair	Cable Status	Distance to Fault	Pair Skew	MDI Polarity	MDI Mode
A23	1-2	OK	0 m	6 ns	Normal	MDIX
	3-6	OK	0 m	0 ns	Normal	
	4-5	OK	0 m	6 ns	Normal	MDIX
	7-8	OK	0 m	6 ns	Normal	

Test Last Run : Fri Apr 22 20:33:23 2011

General transceiver information

Parameter	Description
Interface Index	The switch interface number
Transceiver-type	Pluggable transceiver type
Transceiver model	Pluggable transceiver model
Connector-type	Type of connector of the transceiver

Table Continued

Parameter	Description
Wavelength	For an optical transceiver: the central wavelength of the laser sent, in nm. If the transceiver supports multiple wavelengths, the values will be separated by a comma. An electrical transceiver value is displayed as N/A.
Transfer Distance	Link-length supported by the transceiver in meters. The corresponding transfer medium is shown in brackets following the transfer distance value, For example, 50um multimode fiber. If the transceiver supports multiple transfer media, the values are separated by a comma.
Diagnostic Support	Shows whether the transceiver supports diagnostics: None Supported DOM Supported VCT Supported
Serial Number	Serial number of the transceiver
Link Status	Link up or down
Speed	Speed of transceiver in Mbps
Duplex	Type of duplexing
Cable Status	Values are OK, Open, Short, or Impedance
Distance to Fault	The distance in meters to a cable fault (accuracy is +/- 2 meters); displays 0 (zero) if there is no fault
Pair Skew	Difference in propagation between the fastest and slowest wire pairs
Pair Polarity	Signals on a wire pair are polarized, with one wire carrying the positive signal and one carrying the negative signal.
MDI Mode	The MDI crossover status of the two wire pairs (1&2, 3&6, 4&5, 7&8), will be either MDI or MDIX

Using the Event Log for troubleshooting switch problems

The Event Log records operating events in single- or double-line entries and serves as a tool to isolate and troubleshoot problems.

Once the log has received 2000 entries, it discards the oldest message each time a new message is received. The Event Log window contains 14 log entry lines. You can scroll through it to view any part of the log.

Once the log has received 2000 entries, it discards the oldest message each time a new message is received. The Event Log window contains 14 log-entry lines. You can scroll through it to view any part of the log.



The Event Log is **erased** if power to the switch is interrupted or if you enter the `boot system` command. The contents of the Event Log are **not** erased if you:

- Reboot the switch by choosing the **Reboot Switch** option from the menu interface.
- Enter the `reload` command from the CLI.

Event Log entries

As shown in **Figure 60: Format of an event log entry** on page 315, each Event Log entry is composed of six or seven fields, depending on whether numbering is turned on or not:

Figure 60: *Format of an event log entry*

Severity	Date	Time	Event number	System Module	Management Module	Event Message
M	10/28/09	21:45:42	03002	system:	AM1:	System reboot due to Reset Switch

Item	Description
Severity	One of the following codes (from highest to lowest severity): M —(major) indicates that a fatal switch error has occurred. E —(error) indicates that an error condition occurred on the switch. W —(warning) indicates that a switch service has behaved unexpectedly. I —(information) provides information on normal switch operation. D —(debug) is reserved for HPE internal diagnostic information.
Date	The date in the format mm/dd/yy when an entry is recorded in the log.
Time	The time in the format hh:mm:ss when an entry is recorded in the log.
Event number	The number assigned to an event. You can turn event numbering on and off with the <code>[no] log-number</code> command.
System module	The internal module (such as "ports:" for port manager) that generated a log entry. If VLANs are configured, a VLAN name also appears for an event that is specific to an individual VLAN.
Event message	A brief description of the operating event.

Table 27: Event Log system modules

System module	Description	Documented in HPE Switch hardware/software guide
802.1x	<p>802.1X authentication: Provides access control on a per-client or per-port basis:</p> <ul style="list-style-type: none"> • Client-level security that allows LAN access to 802.1X clients (up to 32 per port) with valid user credentials • Port-level security that allows LAN access only on ports on which a single 802.1X-capable client (supplicant) has entered valid RADIUS user credentials 	<i>Access Security Guide</i>
addrmgr	Address Table Manager: Manages MAC addresses that the switch has learned and are stored in the switch's address table.	<i>Management and Configuration Guide</i>
auth	Authorization: A connected client must receive authorization through web, AMC, RADIUS-based, TACACS+-based, or 802.1X authentication before it can send traffic to the switch.	<i>Access Security Guide</i>
cdp	Cisco Discovery Protocol: Supports reading CDP packets received from neighbor devices, enabling a switch to learn about adjacent CDP devices. HPE does not support the transmission of CDP packets to neighbor devices.	<i>Management and Configuration Guide</i>
console	Console interface used to monitor switch and port status, reconfigure the switch, and read the event log through an in-band Telnet or out-of-band connection.	<i>Installation and Getting Started Guide</i>

Table Continued

System module	Description	Documented in HPE Switch hardware/software guide
cos	Class of Service (CoS): Provides priority handling of packets traversing the switch, based on the IEEE 802.1p priority carried by each packet. CoS messages also include QoS events. The QoS feature classifies and prioritizes traffic throughout a network, establishing an end-to-end traffic priority policy to manage available bandwidth and improve throughput of important data.	<i>Advanced Traffic Management Guide</i>
dhcp	Dynamic Host Configuration Protocol (DHCP) server configuration: Switch is automatically configured from a DHCP (Bootp) server, including IP address, subnet mask, default gateway, Timep Server address, and TFTP server address.	<i>Management and Configuration Guide</i>
dhcp v6c	DHCP for IPv6 prefix assignment	<i>IPv6 Configuration Guide</i>
dhcpr	DHCP relay: Forwards client-originated DHCP packets to a DHCP network server.	<i>Advanced Traffic Management Guide</i>
download	Download operation for copying a software version or files to the switch.	<i>Management and Configuration Guide</i>
dma	Direct Access Memory (DMA): Transmits and receives packets between the CPU and the switch.	—
fault	Fault Detection facility, including response policy and the sensitivity level at which a network problem should generate an alert.	<i>Management and Configuration Guide</i>
ffi	Find, Fix, and Inform: Event or alert log messages indicating a possible topology loop that causes excessive network activity and results in the network running slow. FFI messages include events on transceiver connections with other network devices.	<i>Installation and Getting Started Guide</i> <i>Management and Configuration Guide</i>

Table Continued

System module	Description	Documented in HPE Switch hardware/software guide
garp	Generic Attribute Registration Protocol (GARP), defined in the IEEE 802.1D-1998 standard.	<i>Advanced Traffic Management Guide</i>
gvrp	GARP VLAN Registration Protocol (GVRP): Manages dynamic 802.1Q VLAN operations, in which the switch creates temporary VLAN membership on a port to provide a link to another port in the same VLAN on another device.	<i>Advanced Traffic Management Guide</i>
hpesp	Management module that maintains communication between switch ports.	<i>Installation and Getting Started Guide</i>
igmp	Internet Group Management Protocol: Reduces unnecessary bandwidth usage for multicast traffic transmitted from multimedia applications on a per-port basis.	<i>Multicast and Routing Guide</i>
ip	IP addressing: Configures the switch with an IP address and subnet mask to communicate on the network and support remote management access; configures multiple IP addresses on a VLAN; enables IP routing on the switch.	<i>Management and Configuration Guide</i>
iplock	IP Lockdown: Prevents IP source address spoofing on a per-port and per-VLAN basis by forwarding only the IP packets in VLAN traffic that contain a known source IP address and MAC address binding for the port.	<i>Access Security Guide</i>
ipx	Novell Netware protocol filtering: On the basis of protocol type, the switch can forward or drop traffic to a specific set of destination ports on the switch.	<i>Access Security Guide</i>
lACP	LACP trunks: The switch can either automatically establish an 802.3ad-compliant trunk group or provide a manually configured, static LACP trunk.	<i>Management and Configuration Guide</i>

Table Continued

System module	Description	Documented in HPE Switch hardware/software guide
ldbal	Load balancing in LACP port trunks or 802.1s Multiple Spanning Tree protocol (MSTP) that uses VLANs in a network to improve network resource utilization and maintain a loop-free environment. Load-balancing messages also include switch meshing events. The switch meshing feature provides redundant links, improved bandwidth use, and support for different port types and speeds.	<i>Management and Configuration Guide</i> <i>Advanced Traffic Management Guide</i>
lldp	Link-Layer Discovery Protocol: Supports transmitting LLDP packets to neighbor devices and reading LLDP packets received from neighbor devices, enabling a switch to advertise itself to adjacent devices and to learn about adjacent LLDP devices.	<i>Management and Configuration Guide</i>
macauth	Web and MAC authentication: Port-based security employed on the network edge to protect private networks and the switch itself from unauthorized access using one of the following interfaces: <ul style="list-style-type: none"> • Web page login to authenticate users for access to the network • RADIUS server that uses a device's MAC address for authentication 	<i>Access Security Guide</i>
maclock	MAC lockdown and MAC lockout <ul style="list-style-type: none"> • MAC lockdown prevents station movement and MAC address "hijacking" by requiring a MAC address to be used only on an assigned port on the switch. MAC Lockdown also restricts the client device to a specific VLAN. • MAC lockout blocks a specific MAC address so that the switch drops all traffic to or from the specified address. 	<i>Access Security Guide</i>

Table Continued

System module	Description	Documented in HPE Switch hardware/software guide
mgr	HPE PCM and PCM+: Windows-based network management solutions for managing and monitoring performance of HPE switches. PCM messages also include events for configuration operations.	<i>Management and Configuration Guide</i>
netinet	Network Internet: Monitors the creation of a route or an Address Resolution Protocol (ARP) entry and sends a log message in case of failure.	<i>Advanced Traffic Management Guide</i>
pagp	Ports Aggregation Protocol (PAgP): Obsolete. Replaced by LACP (802.3ad).	—
ports	Port status and port configuration features, including mode (speed and duplex), flow control, broadcast limit, jumbo packets, and security settings. Port messages include events on POE operation and transceiver connections with other network devices.	<i>Installation and Getting Started Guide</i> <i>Management and Configuration Guide</i> <i>Access Security Guide</i>
radius	RADIUS (Remote Authentication Dial-In User Service) authentication and accounting: A network server is used to authenticate user-connection requests on the switch and collect accounting information to track network resource usage.	<i>Access Security Guide</i>
snmp	Simple Network Management Protocol: Allows you to manage the switch from a network management station, including support for security features, event reporting, flow sampling, and standard MIBs.	<i>Management and Configuration Guide</i>
sntp	Simple Network Time Protocol: Synchronizes and ensures a uniform time among interoperating devices.	<i>Management and Configuration Guide</i>

Table Continued

System module	Description	Documented in HPE Switch hardware/software guide
ssh	<p>Secure Shell version 2 (SSHv2): Provides remote access to management functions on a switch via encrypted paths between the switch and management station clients capable of SSH operation.</p> <p>SSH messages also include events from the Secure File Transfer Protocol (SFTP) feature. SFTP provides a secure alternative to TFTP for transferring sensitive information, such as switch configuration files, to and from the switch in an SSH session.</p>	<i>Access Security Guide</i>
ssl	<p>Secure Socket Layer Version 3 (SSLv3), including Transport Layer Security (TLSv1) support: Provides remote web access to a switch via encrypted paths between the switch and management station clients capable of SSL/TLS operation.</p>	<i>Access Security Guide</i>
stack	<p>Stack management: Uses a single IP address and standard network cabling to manage a group (up to 16) of switches in the same IP subnet (broadcast domain), resulting in a reduced number of IP addresses and simplified management of small workgroups for scaling your network to handle increased bandwidth demand.</p>	<i>Advanced Traffic Management Guide</i>
stp	<p>Multiple-instance spanning tree protocol/MSTP (802.1s): Ensures that only one active path exists between any two nodes in a group of VLANs in the network. MSTP operation is designed to avoid loops and broadcast storms of duplicate messages that can bring down the network.</p>	<i>Advanced Traffic Management Guide</i>

Table Continued

System module	Description	Documented in HPE Switch hardware/software guide
system	<p>Switch management, including system configuration, switch bootup, activation of boot ROM image, memory buffers, traffic and security filters.</p> <p>System messages also include events from management interfaces (menu, CLI, and HPE PCM+) used to reconfigure the switch and monitor switch status and performance.</p>	<p><i>Management and Configuration Guide</i></p> <p><i>Access Security Guide</i></p>
tacacs	<p>TACACS+ authentication: A central server is used to control access to the switches (and other TACACS-aware devices) in the network through a switch's console port (local access) or Telnet (remote access).</p>	<p><i>Access Security Guide</i></p>
tcp	<p>Transmission Control Protocol: A transport protocol that runs on IP and is used to set up connections.</p>	<p><i>Advanced Traffic Management Guide</i></p>
telnet	<p>Session established on the switch from a remote device through the Telnet virtual terminal protocol.</p>	<p><i>Management and Configuration Guide</i></p>
tftp	<p>Trivial File Transfer Protocol: Supports the download of files to the switch from a TFTP network server.</p>	<p><i>Management and Configuration Guide</i></p>
timep	<p>Time Protocol: Synchronizes and ensures a uniform time among interoperating devices.</p>	<p><i>Management and Configuration Guide</i></p>
update	<p>Updates (TFTP or serial) to HPE switch software and updates to running-config and start-up config files</p>	<p><i>Management and Configuration Guide</i></p>

Table Continued

System module	Description	Documented in HPE Switch hardware/software guide
vlan	<p>Static 802.1Q VLAN operations, including port-and protocol-based configurations that group users by logical function instead of physical location</p> <ul style="list-style-type: none"> A port-based VLAN creates a layer-2 broadcast domain comprising member ports that bridge IPv4 traffic among themselves. <p>VLAN messages include events from management interfaces (menu, CLI, and HPE PCM+) used to reconfigure the switch and monitor switch status and performance.</p>	<i>Advanced Traffic Management Guide</i>
xmodem	Xmodem: Binary transfer feature that supports the download of software files from a PC or UNIX workstation.	<i>Management and Configuration Guide</i>

Using the Menu

To display the Event Log from the Main Menu, select `Event Log`. The following example shows a sample event log display.

An event log display

```
Switch 5406z1                               25-Oct-2013  18:02:52
=====CONSOLE - MANAGER MODE -
=====
M 10/25/13 16:30:02 sys: 'Operator cold reboot from CONSOLE session.'
I 10/25/13 17:42:51 00061 system: -----
-
I 10/25/13 17:42:51 00063 system: System went down : 10/25/13 16:30:02
I 10/25/13 17:42:51 00064 system: Operator cold reboot from CONSOLE session.
W 10/25/13 17:42:51 00374 chassis: WARNING: SSC is out of Date: Load 8.2 or
newer
I 10/25/13 17:42:51 00068 chassis: Slot D Inserted
I 10/25/13 17:42:51 00068 chassis: Slot E Inserted
I 10/25/13 17:42:51 00068 chassis: Slot F Inserted
I 10/25/13 17:42:51 00690 udpf: DHCP relay agent feature enabled
I 10/25/13 17:42:51 00433 ssh: Ssh server enabled
I 10/25/13 17:42:51 00400 stack: Stack Protocol disabled
I 10/25/13 17:42:51 00128 tftp: Enable succeeded
I 10/25/13 17:42:51 00417 cdp: CDP enabled

----  Log events stored in memory 1-751. Log events on screen 690-704.

  Actions->   Back   Next page   Prev page   End   Help

Return to previous screen.
```

Use up/down arrow to scroll one line, left/right arrow keys to change action selection, and <Enter> to execute action.

The **log status line** below the recorded entries states the total number of events stored in the event log and which logged events are currently displayed.

To scroll to other entries in the Event Log, either preceding or following the currently visible portion, press the keys indicated at the bottom of the display (Back,Nextpage, Prev page, or End) or the keys described in the following table.

Event Log control keys

Key	Action
[N]	Advances the display by one page (next page).
[P]	Rolls back the display by one page (previous page).
[V]	Advances display by one event (down one line).
[^]	Rolls back display by one event (up one line).
[E]	Advances to the end of the log.
[H]	Displays Help for the Event Log.

Using the CLI

Syntax:

By default, the `show logging` command displays the log messages recorded since the last reboot in chronological order:

-a	Displays all recorded log messages, including those before the last reboot.
-b	Displays log events as the time since the last reboot instead of in a date/time format.
-r	Displays all recorded log messages, with the most recent entries listed first (reverse order).
-s	Displays the active management module (AM) and standby management module (SM) log events.
-t	Displays the log events with a granularity of 10 milliseconds.
-m	Displays only major log events.
-e	Displays only error event class.
-p	Displays only performance log events.
-w	Displays only warning log events.
-i	Displays only informational log events.
-d	Displays only debug log events.

Table Continued

<code>filter</code>	Displays only log filter configuration and status information.
<code><option-str></code>	Displays all Event Log entries that contain the specified text. Use an <code><option-str></code> value with <code>-a</code> or <code>-r</code> to further filter <code>show logging</code> command output.

Example:

To display all Event Log messages that have "system" in the message text or module name, enter the following command:

```
switch# show logging -a system
```

To display all Event Log messages recorded since the last reboot that have the word "system" in the message text or module name, enter:

```
switch# show logging system
```

Clearing Event Log entries

Syntax:

Removes all entries from the event log display output.

Use the `clear logging` command to hide, but not erase, Event Log entries displayed in `show logging` command output. Only new entries generated after you enter the command will be displayed.

To redisplay all hidden entries, including Event Log entries recorded prior to the last reboot, enter the `show logging -a` command.

Turning event numbering on

Syntax:

```
[no] log-numbers
```

Turns event numbering on and off

Using log throttling to reduce duplicate Event Log and SNMP messages

A recurring event can generate a series of duplicate Event Log messages and SNMP traps in a relatively short time. As a result, the Event Log and any configured SNMP trap receivers may be flooded with excessive, exactly identical messages. To help reduce this problem, the switch uses **log throttle periods** to regulate (throttle) duplicate messages for recurring events, and maintains a counter to record how many times it detects duplicates of a particular event since the last system reboot.

When the first instance of a particular event or condition generates a message, the switch initiates a log throttle period that applies to all recurrences of that event. If the logged event recurs during the log throttle period, the switch increments the counter initiated by the first instance of the event, but does not generate a new message.

If the logged event repeats again after the log throttle period expires, the switch generates a duplicate of the first message, increments the counter, and starts a new log throttle period during which any additional instances of the event are counted, but not logged. Thus, for a particular recurring event, the switch displays only one message in the Event Log for each log throttle period in which the event reoccurs. Also, each logged instance of the event message includes counter data showing how many times the event has occurred since the last reboot. The switch manages messages to SNMP trap receivers in the same way.

Log throttle periods

The length of the log throttle period differs according to an event's severity level:

Severity level	Log throttle period
I (Information)	6000 Seconds
W (Warning)	600 Seconds
D (Debug)	60 Seconds
M (Major)	6 Seconds

Example:

Suppose that you configure VLAN 100 on the switch to support PIM operation, but do not configure an IP address. If PIM attempts to use VLAN 100, the switch generates the first instance of the following Event Log message and counter.



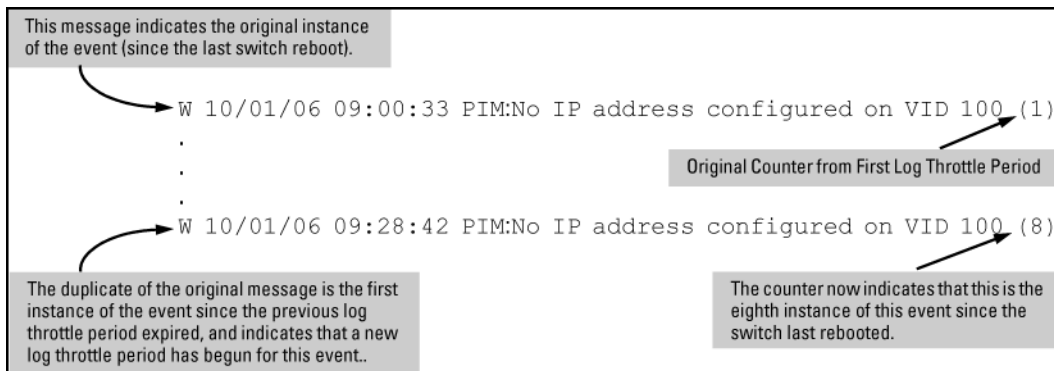
In **The first instance of an event message and counter** on page 326 the counter (1) indicates that this is the first instance of this event since the switch last rebooted.

The first instance of an event message and counter

```
W 10/01/12 09:00:33 PIM:No IP address configured on VID 100 (1)
```


If PIM operation causes the same event to occur six more times during the initial log throttle period, there are no further entries in the Event Log. However, if the event occurs again after the log throttle period has expired, the switch repeats the message (with an updated counter) and starts a new log throttle period.

Figure 61: Duplicate messages over multiple log throttling periods



Note that if the same type of event occurs under different circumstances, the switch handles these as unrelated events for the purpose of Event Log messages. For example, if PIM operation simultaneously detects that VLANs 100 and 205 are configured without IP addresses, you see log messages similar to the following:

Figure 62: Example: of log messages generated by unrelated events of the same type



Example: of event counter operation

Suppose the switch detects the following after a reboot:

- Three duplicate instances of the PIM "Send error" during the first log throttle period for this event
- Five more instances of the same Send error during the second log throttle period for this event
- Four instances of the same Send error during the third log throttle period for this event

In this case, the duplicate message appears three times in the Event Log (once for each log throttle period for the event being described), and the duplicate message counter increments as shown in the following table. (The same operation applies for messages sent to any configured SNMP trap receivers.)

Table 28: How the duplicate message counter increments

Instances during 1st log throttle period	Instances during 2nd log throttle period	Instances during 3rd log throttle period	Duplicate message counter ¹
3			1
	5		4
		4	9

¹ This value always comprises the first instance of the duplicate message in the current log throttle period plus all previous occurrences of the duplicate message occurring since the switch last rebooted.

Reporting information about changes to the running configuration

Syslog can be used for sending notifications to a remote syslog server about changes made to the running configuration. The notifications in the syslog messages are sent in ASCII format and contain this information:

- Notice-Type: Describes the syslog notification as a "running config change".
- Event-ID: Identifier for the running config change event that occurred on the switch.
- Config-Method: The source for the running config change.
- Device-Name: The managed device.
- User-Name: User who made the running config change.
- Remote-IP-Address: IP address of a remote host from which the user is connected.

Syntax:

```
[no] logging notify <running-config-change> [transmission-interval <0-4294967295>
```

Enables sending the running configuration change notifications to the syslog server.

The `no` form of the command disables sending the running configuration changes to the syslog server.

Default: Disabled

<code><running-config-change ></code>	Mandatory option for the notify parameter. Specifies the type of notification to send.
<code>transmission-interval <0-4294967295></code>	Specifies the time interval (in seconds) between the transmission of two consecutive notifications. Running config changes occurring within the specified interval will not generate syslog notifications.

A value of zero means there is no limit; a notification is sent for every running config change.

Default: Zero

Sending running config changes to the syslog server

```
switch(config)# logging notify running-config-change  
transmission-interval 10
```

Debug/syslog operation

While the Event Log records switch-level progress, status, and warning messages on the switch, the debug/system logging (**syslog**) feature provides a way to record Event Log and debug messages on a remote device. For example, you can send messages about routing misconfigurations and other network protocol details to an external device, and later use them to debug network-level problems.

Debug/syslog messaging

The debug/syslog feature allows you to specify the types of Event Log and debug messages that you want to send to an external device. You can perform the following operations:

- Use the `debug` command to configure messaging reports for the following event types:
 - Events recorded in the switch's Event Log
 - LLDP events
 - SSH events
- Use the `logging` command to select a subset of Event Log messages to send to an external device for debugging purposes according to:
 - Severity level
 - System module

Hostname in syslog messages

The syslog now messages the sender identified by hostname.

The hostname field identifies the switch that originally sends the syslog message. Configurable through the CLI and SNMP, the format of the hostname field supports the following formats:

- `ip-address`: The IP address of the sending interface will be used as the message origin identifier. This is the default format for the origin identifier. The IP address of the sending interface (in dotted decimal notation) is the default format.
- `hostname`: The hostname of the sending switch will be used as the message origin identifier.
- `none`: No origin identifier will be embedded in the syslog message. Nilvalue is used as defined by “-”.

This configuration is system-wide, not per syslog server.



There is no support in this feature for menu interface, WebUI or a fully qualified domain name. There are no changes in this feature to PCM or IDM. There are no new log events added in this feature.

Logging origin-id

Use the `logging origin-id` command to specify the content for the hostname field.

Syntax:

```
logging origin-id [ip-address|hostname|none]
```

```
[no] logging origin-id [ip-address|hostname|none]
```

To reset the hostname field content back to default (IP-address), use the `no` form of the command.

filter	Creates a filter to restrict which events are logged.
IP-ADDR	Adds an IPv4 address to the list of receiving syslog servers.
IPV6-ADDR	Adds an IPv6 address to the list of receiving syslog servers.
origin-id	Sends the Syslog messages with the specified origin-id.
notify	Notifies the specified type sent to the syslog server(s).
priority-descr	A text string associated with the values of facility, severity, and system-module.
severity	Event messages of the specified severity or higher sent to the syslog server.
system-module	Event messages of the specified system module (subsystem) sent to the syslog server.
hostname	Sets the hostname of the device as the origin-id.
none	Disables origin-id in the syslog message.

Add an IP address to the list of receiving syslog servers.

Use of `no` without an IP address specified will remove all IP addresses from the list of syslog receivers. If an IP address is specified, that receiver will be removed. Both link-local with zone ID and global IPv6 addresses are supported.

- Specify syslog server facility with the option `<facility>`. The command `no logging <facility>` sets the facility back to defaults.
- Specify filtering rules.
- Specify severity for event messages to be filtered to the syslog server with the option `<severity>`. The command `no logging <severity>` sets the severity back to default.
- Event messages of specified system module will be sent to the syslog server. Using `no` sends messages from all system modules. Messages are first filtered by selected severity.
- Specify syslog server transport layer with options `[udp] | [tcp] | [tls]`.
- Specify syslog server port number with options `[udp PORT-NUM] | [tcp PORT-NUM] | [tls PORT-NUM]`.
- Specify notification types to be sent to the syslog server.
- Use the option `transmission-interval` to control the egress rate limit for transmitting notifications, 0 value means there is no rate limit. The values are in seconds. Only one syslog message is allowed for transmission within specified time interval.
- Specify the origin information for the syslog messages with the option `origin-id`.



When the syslog server receives messages from the switch, the IPv6 address of the switch is partly displayed.

Example:

Configured Host Ipv6 Address: 2001::1

Expected Syslog message:

Syslog message: USER.INFO: Oct 11 02:40:02 2001::1 00025 ip: ST1CMDR: VLAN60: ip address 30.1.1.1/24 configured on vlan 60

Actual Truncated syslog message:

Syslog message: USER.INFO: Oct 11 02:40:02 2001:: 00025 ip: ST1CMDR: VLAN60: ip address 30.1.1.1/24 configured on vlan 60

Use the command in the following example to set the origin-id to the hostname.

Setting the origin-id to the hostname

```
switch(config)# logging origin-id hostname
```

The following syslog message will occur:

<14> Jan 1 00:15:35 HP-2910al-24G 00076 ports: port 2 is now on-line

Use the command in the following example to set the origin-id to none (nilvalue).

Setting the origin-id to none (nilvalue)

```
switch(config)# logging origin-id none
```

The following syslog message will occur:

<14> Jan 1 00:15:35 - 00076 ports: port 2 is now on-line

Use any of the commands in the following example to set the origin-id to ip-address (default).

Setting the origin-id to ip-address (default)

```
switch(config)# logging origin-id ip-address
```

```
switch(config)# no logging origin-id hostname
```

```
switch(config)# no logging origin-id none
```

The following syslog message will occur:

<14> Jan 1 00:15:35 169.254.230.236 00076 ports: port 2 is now on-line

Viewing the identification of the syslog message sender

Use the commands `show debug` or `show running-config` to display the identification of the syslog message sender. The default option for `origin-id` is `ip-address`. The command `show running-config` will not display the configured option when `origin-id` is set to the default value of `ip address`.

When `hostname` or `none` is configured using `logging origin-id`, the same displays as part of the `show running-config` command.

Syntax:

```
show debug
```

Default option is `ip-address`.

The following shows the output of the `show debug` command when configured without `login origin-id`.

Output of the show debug command when configured without login origin-id

```
Debug Logging
  Origin identifier: Outgoing Interface IP
  Destination:      None
```

```
Enabled debug types:
  None are enabled.
```

The command `logging origin-id hostname` will produce the syslog message shown in the following example.

Syslog message for logging origin-id hostname

```
Debug Logging
  Origin identifier: Hostname
  Destination:      None
```

```
Enabled debug types:
  None are enabled.
```

The command `logging origin-id none` will produce the syslog message shown in the following example.

Syslog message for logging origin-id none

```
Debug Logging
  Origin identifier: none
  Destination:      None
```

```
Enabled debug types:
  None are enabled.
```

Syntax:

```
show running-config
```

The following example shows the output of the `show running-config` command.

Output of the show running-config command

```
The command logging origin-id hostname will display the
following:
logging origin-id hostname
```

The command `logging origin-id none` will display as the following:

logging origin-id none

SNMP MIB

SNMP support will be provided through the following MIB objects.

HpicfSyslogOriginId = textual-convention

Description This textual convention enumerates the origin identifier of syslog message.

Syntax: integer

	ip-address
	hostname
	none

Status current

hpicfSyslogOriginId OBJECT-TYPE

Description Specifies the content of a Hostname field in the header of a syslog message.

Syntax:

	HpicfSyslogOriginId
	Max-access
	• read-write
	Status
	• current
	Default
	• ip-address

Debug/syslog destination devices

To use debug/syslog messaging, you must configure an external device as the logging destination by using the `logging` and `debug destination` commands. For more information, see [Debug destinations](#) on page 340 and [Configuring a syslog server](#) on page 341.

A debug/syslog destination device can be a syslog server and/or a console session. You can configure debug and logging messages to be sent to:

- Up to six syslog servers
- A CLI session through a direct RS-232 console connection, or a Telnet or SSH session

Debug/syslog configuration commands

Event notification logging	—	Automatically sends switch-level event messages to the switch's Event Log. Debug and syslog do not affect this operation, but add the capability of directing Event Log messaging to an external device.
logging command	<code><syslog-ip-addr></code>	Enables syslog messaging to be sent to the specified IP address. IPv4 and IPv6 are supported.
	<code>facility</code>	(Optional) The <code>logging facility</code> command specifies the destination (facility) subsystem used on a syslog server for debug reports.
	<code>priority-desc</code>	A text string associated with the values of facility, severity, and system-module.
	<code>severity</code>	Sends Event Log messages of equal or greater severity than the specified value to configured debug destinations. (The default setting is to send Event Log messages from all severity levels.)
	<code>system-module</code>	Sends Event Log messages from the specified system module to configured debug destinations. The severity filter is also applied to the system-module messages you select. The default setting is to send Event Log messages from all system modules. To restore the default setting, enter the <code>no logging system-module <system-module></code> or <code>logging system-module all-pass</code> commands.
	<code>all</code>	Sends debug logging to configured debug destinations for all ACL, Event Log, IP-OSPF, and IP-RIP options.

Table Continued

	destination	logging: Disables or re-enables syslog logging on one or more syslog servers configured with the <code>logging syslog-ip-addr</code> command.session: Assigns or re-assigns destination status to the terminal device that was most recently used to request debug output.buffer: Enables syslog logging to send the debug message types specified by the <code>debug <debug-type></code> command to a buffer in switch memory.
	event	Sends standard Event Log messages to configured debug destinations. (The same messages are also sent to the switch's Event Log, regardless of whether you enable this option.)
	ip	forwarding: Sends IPv4 forwarding messages to the debug destinations.packet: Sends IPv4 packet messages to the debug destinations.rip: Sends RIP event logging to the debug destinations.
	ipv6	dhcpv6-client: Sends DHCPv6 client debug messages to the configured debug destination.forwarding: Sends IPv6 forwarding messages to the debug destination(s)nd: Sends IPv6 debug messages for IPv6 neighbor discovery to the configured debug destinations.packet: Sends IPv6 packet messages to the debug destinations.
	lldp	Sends LLDP debug messages to the debug destinations.
	ssh	Sends SSH debug messages at the specified level to the debug destination. The levels are fatal, error, info, verbose, debug, debug2, and debug3.

Using the Debug/Syslog feature, you can perform the following operations:

- Configure the switch to send Event Log messages to one or more Syslog servers. In addition, you can configure the messages to be sent to the User log facility (default) or to another log facility on configured Syslog servers.
- Configure the switch to send Event Log messages to the current management- access session (serial-connect CLI, Telnet CLI, or SSH).
- Disable all Syslog debug logging while retaining the Syslog addresses from the switch configuration. This allows you to configure Syslog messaging and then disable and re-enable it as needed.
- Display the current debug configuration. If Syslog logging is currently active, the list of configured Syslog servers is displayed.
- Display the current Syslog server list when Syslog logging is disabled.

Configuring debug/syslog operation

Procedure

1. To use a syslog server as the destination device for debug messaging, follow these steps:
 - a. Enter the `logging <syslog-ip-addr>` command at the global configuration level to configure the syslog server IP address and enable syslog logging. Optionally, you may also specify the destination subsystem to be used on the syslog server by entering the `logging facility` command. If no other syslog server IP addresses are configured, entering the `logging` command enables both debug messaging to a syslog server and the event debug message type. As a result, the switch automatically sends Event Log messages to the syslog server, regardless of other debug types that may be configured.
 - b. Re-enter the `logging` command in Step 1a to configure additional syslog servers. You can configure up to a total of six servers. (When multiple server IP addresses are configured, the switch sends the debug message types that you configure in **Step 3** to all IP addresses.)
2. To use a CLI session on a destination device for debug messaging:
 - a. Set up a serial, Telnet, or SSH connection to access the switch's CLI.
 - b. Enter the `debug destination session` command at the manager level.
3. Enable the types of debug messages to be sent to configured syslog servers, the current session device, or both by entering the `debug <debug-type>` command and selecting the desired options.

Repeat this step if necessary to enable multiple debug message types.

By default, Event Log messages are sent to configured debug destination devices. To block Event Log messages from being sent, enter the `no debug event` command.

4. If necessary, enable a subset of Event Log messages to be sent to configured syslog servers by specifying a severity level, a system module, or both using the following commands:

```
switch(config)# logging severity <debug | major | error | warning | info>
switch(config)# logging system-module <system-module>
```

To display a list of valid values for each command, enter `logging severity` or `logging system-module` followed by `?` or pressing the Tab key.

The severity levels in order from the highest to lowest severity are major, error, warning, info, and debug. For a list of valid values for the `logging system-module <system-module>` command, see **Event Log system modules**.

5. If you configure system-module, severity-level values, or both to filter Event Log messages, when you finish troubleshooting, you may want to reset these values to their default settings so that the switch sends all Event Log messages to configured debug destinations (syslog servers, CLI session, or both).

To remove a configured setting and restore the default values that send all Event Log messages, enter one or both of the following commands:

```
switch(config)# no logging severity <debug | major | error | warning | info>
switch(config)# no logging system-module <system-module>
```



If you configure a severity-level, system-module, logging destination, or logging facility value and save the settings to the startup configuration (For example, by entering the `write memory` command), the debug settings are saved after a system reboot (power cycle or reboot) and re-activated on the switch. As a result, after switch startup, one of the following situations may occur:

- Only a partial set of Event Log messages may be sent to configured debug destinations.
- Messages may be sent to a previously configured syslog server used in an earlier debugging session.

Viewing a debug/syslog configuration

Use the `show debug` command to display the currently configured settings for:

- Debug message types and Event Log message filters (severity level and system module) sent to debug destinations
- Debug destinations (syslog servers or CLI session) and syslog server facility to be used

Syntax:

```
show debug
```

Displays the currently configured debug logging destinations and message types selected for debugging purposes. (If no syslog server address is configured with the `logging <syslog-ip-addr>` command, no `show debug` command output is displayed.)

Output of the show debug command

```
switch(config)# show debug

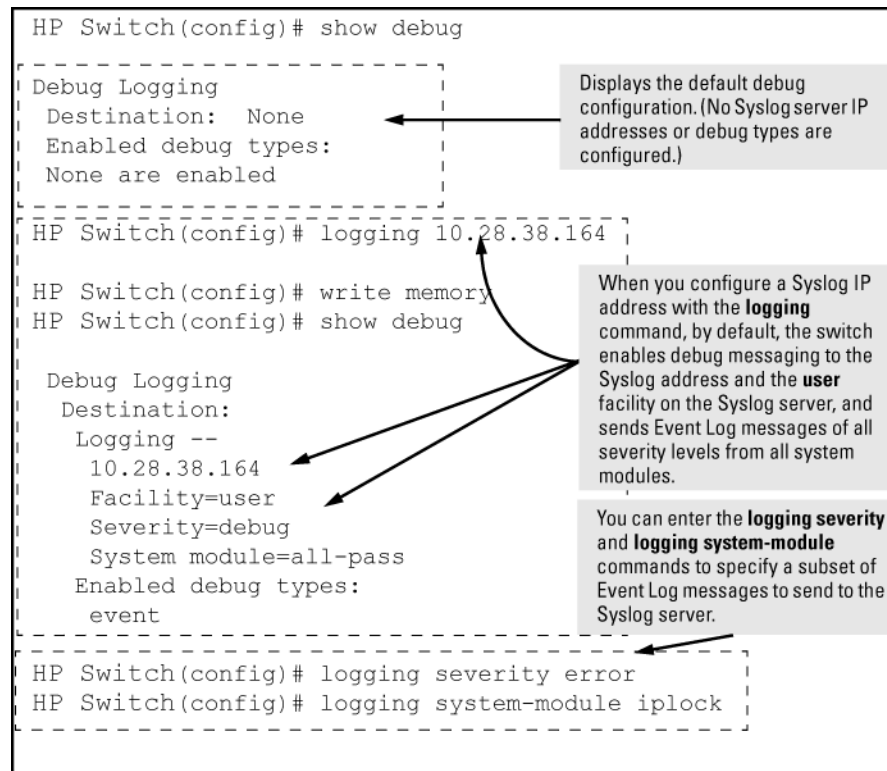
Debug Logging
Destination:
Logging --
 10.28.38.164
Facility=kern
Severity=warning
System module=all-pass
Enabled debug types:
event
```

Example:

In the following Example:, no syslog servers are configured on the switch (default setting). When you configure a syslog server, debug logging is enabled to send Event Log messages to the server. To limit the Event Log

messages sent to the syslog server, specify a set of messages by entering the `logging severity` and `logging system-module` commands.

Figure 63: Syslog configuration to receive event log messages from specified system module and severity levels



As shown at the top of **Figure 63: Syslog configuration to receive event log messages from specified system module and severity levels** on page 337, if you enter the `show debug` command when no syslog server IP address is configured, the configuration settings for syslog server facility, Event Log severity level, and system module are not displayed. However, after you configure a syslog server address and enable syslog logging, all debug and logging settings are displayed with the `show debug` command.

If you do not want Event Log messages sent to syslog servers, you can block the messages from being sent by entering the `no debug event` command. (There is no effect on the normal logging of messages in the switch's Event Log.)

Example:

The next Example: shows how to configure:

- Debug logging of ACL and IP-OSPF packet messages on a syslog server at 18.38.64.164 (with user as the default logging facility).
- Display of these messages in the CLI session of your terminal device's management access to the switch.
- Blocking Event Log messages from being sent from the switch to the syslog server and a CLI session.

To configure syslog operation in these ways with the debug/syslog feature disabled on the switch, enter the commands shown in **Figure 64: Debug/syslog configuration for multiple debug types and multiple destinations** on page 338.

Figure 64: Debug/syslog configuration for multiple debug types and multiple destinations

```

HP Switch# config
HP Switch(config)# logging 10.38.64.164
HP Switch(config)# show debug

Debug Logging
Destination:
Logging --
 10.38.64.164
Facility=user
Severity=debug
System module=all-pass
Enabled debug types:
event

HP Switch(config)# no debug event
HP Switch(config)# debug acl
HP Switch(config)# debug ip ospf packet
HP Switch(config)# debug destination session
HP Switch(config)# show debug

Debug Logging
Destination:
Logging --
 10.38.64.164
Facility=user
Severity=debug
System module=all-pass
Session
Enabled debug types:
acl log
ip ospf packet

```

Configure a Syslog server IP address. (No other Syslog servers are configured on the switch.) The server address serves as an active debug destination for any configured debug types.)

Display the new debug configuration. (Default debug settings - facility, severity, system module, and debug types- are displayed.)

Remove the unwanted event message logging to debug destinations.

Configure the debug messages types that you want to send to the Syslog server and CLI session.

Configure the CLI session as a debug destination.

Display the final debug and Syslog server configuration.

Debug command

At the manager level, use the `debug` command to perform two main functions:

- Specify the types of event messages to be sent to an external destination.
- Specify the destinations to which selected message types are sent.

By default, no debug destination is enabled and only Event Log messages are enabled to be sent.



To configure a syslog server, use the `logging <syslog-ip-addr>` command. For more information, see **Configuring a syslog server** on page 341.

Debug messages

Syntax:

[no] debug <debug-type>



all	Configures the switch to send all debug message types to configured debug destinations. (Default: Disabled—No debug messages are sent.)
cdp	Sends CDP information to configured debug destinations.
destination	<p>logging—Disables or re-enables syslog logging on one or more syslog servers configured with the <code>logging <syslog-ip-addr></code> command.</p> <p>session—Assigns or re-assigns destination status to the terminal device that was most recently used to request debug output.</p> <p>buffer—Enables syslog logging to send the debug message types specified by the <code>debug <debug-type></code> command to a buffer in switch memory.</p>
event	<p>Configures the switch to send Event Log messages to configured debug destinations.</p> <div style="border: 1px solid black; padding: 5px; margin: 10px 0;">  <p>NOTE This value does not affect the reception of event notification messages in the Event Log on the switch.</p> </div> <p>Event Log messages are automatically enabled to be sent to debug destinations in these conditions:</p> <ul style="list-style-type: none"> • If no syslog server address is configured and you enter the <code>logging <syslog-ip-addr></code> command to configure a destination address. • If at least one syslog server address is configured in the startup configuration, and the switch is rebooted or reset. <p>Event log messages are the default type of debug message sent to configured debug destinations.</p>
ip [fib packet]	Sends IP messages to configured destinations.
ip [fib [events]]	For the configured debug destinations: <code>events</code> —Sends IP forwarding information base events.
ip [packet]	Enables the specified PIM message type.
ipv6 [dhcpv6-client nd packet]	<div style="border: 1px solid black; padding: 5px; margin: 10px 0;">  <p>NOTE See the "IPv6 Diagnostic and Troubleshooting" in the IPv6 configuration guide for your switch for more detailed IPv6 debug options.</p> </div> <p>When no debug options are included, displays debug messages for all IPv6 debug options.</p> <p><code>dhcpv6-client [events packet]</code>—Displays DHCPv6 client event and packet data.</p> <p><code>nd</code>—Displays debug messages for IPv6 neighbor discovery.</p> <p><code>packet</code>—Displays IPv6 packet messages.</p>

Table Continued


lldp	Enables all LLDP message types for the configured destinations.
security [port-access port-security radius-server ssh tacacs-server user-profile-mib]	<p>port-access—Sends port-access debug messages to the debug destination.</p> <p>radius-server—Sends RADIUS debug messages to the debug destination.</p> <p>ssh—Sends SSH debug messages at the specified level to the debug destination. The levels are fatal, error, info, verbose, debug, debug2, and debug3.</p> <p>tacacs-server—Sends TACACS debug messages to the debug destination.</p> <p>user-profile-mib—Sends user profile MIB debug messages to the debug destination.</p>
snmp <pdu>	Displays the SNMP debug messages.pdu—Displays SNMP pdu debug messages.

Debug destinations

Use the `debug destination` command to enable (and disable) syslog messaging on a syslog server or to a CLI session for specified types of debug and Event Log messages.

Syntax:

[no] debug destination {<logging | session | buffer>}

logging	<p>Enables syslog logging to configured syslog servers so that the debug message types specified by the <code>debug <debug-type></code> command (see Debug messages on page 338) are sent. (Default: Logging disabled) To configure a syslog server IP address, see Configuring a syslog server on page 341.</p> <hr/> <p> NOTE Debug messages from the switches covered in this guide have a debug severity level. Because the default configuration of some syslog servers ignores syslog messages with the debug severity level, ensure that the syslog servers you want to use to receive debug messages are configured to accept the debug level. For more information, see Operating notes for debug and Syslog on page 350.</p>
session	<p>Enables transmission of event notification messages to the CLI session that most recently executed this command. The session can be on any one terminal emulation device with serial, Telnet, or SSH access to the CLI at the Manager level prompt (<code>switch#_</code>). If more than one terminal device has a console session with the CLI, you can redirect the destination from the current device to another device. Do so by executing <code>debug destination session</code> in the CLI on the terminal device on which you now want to display event messages. Event message types received on the selected CLI session are configured with the <code>debug <debug-type></code> command.</p>
buffer	<p>Enables syslog logging to send the debug message types specified by the <code>debug <debug-type></code> command to a buffer in switch memory. To view the debug messages stored in the switch buffer, enter the <code>show debug buffer</code> command.</p>

Logging command

At the global configuration level, the `logging` command allows you to enable debug logging on specified syslog servers and select a subset of Event Log messages to send for debugging purposes according to:

- Severity level
- System module

By specifying both a severity level and system module, you can use both configured settings to filter the Event Log messages you want to use to troubleshoot switch or network error conditions.



After you configure a syslog server and a severity level and/or system module to filter the Event Log messages that are sent, if you save these settings to the startup configuration file by entering the `write memory` command, these debug and logging settings are automatically re-activated after a switch reboot or power recycle. The debug settings and destinations configured in your previous troubleshooting session will then be applied to the current session, which may not be desirable.

After a reboot, messages remain in the Event Log and are not deleted. However, after a power recycle, all Event Log messages are deleted.

If you configure a severity level, system module, or both to temporarily filter Event Log messages, be sure to reset the values to their default settings by entering the `no` form of the following commands to ensure that Event Log messages of all severity levels and from all system modules are sent to configured syslog servers:

```
switch(config)# no logging severity <debug | major | error | warning | info>
switch(config)# no logging system-module <system-module>
```

Configuring a syslog server

Syslog is a client-server logging tool that allows a client switch to send event notification messages to a networked device operating with syslog server software. Messages sent to a syslog server can be stored to a file for later debugging analysis.

To use the syslog feature, you must install and configure a syslog server application on a networked host accessible to the switch. For instructions, see the documentation for the syslog server application.

To configure a syslog service, use the `logging <syslog-ip-addr>` command as shown below.

When you configure a syslog server, Event Log messages are automatically enabled to be sent to the server. To reconfigure this setting, use the following commands:

- `debug`
Specifies additional debug message types (see [Debug messages](#) on page 338).
- `logging`
Configures the system module or severity level used to filter the Event Log messages sent to configured syslog servers. (See [Configuring the severity level for Event Log messages sent to a syslog server](#) on page 349 and [Configuring the system module used to select the Event Log messages sent to a syslog server](#) on page 350.)

To display the currently configured syslog servers as well as the types of debug messages and the severity-level and system-module filters used to specify the Event Log messages that are sent, enter the `show debug` command (See [Debug/syslog configuration commands](#) on page 333).

Syntax:

```
[no] logging <syslog-ip-addr>
```

Enables or disables syslog messaging to the specified IP address. You can configure up to six addresses. If you configure an address when none are already configured, this command enables destination logging (syslog) and the Event debug type. Therefore, at a minimum, the switch begins sending Event Log messages to configured syslog servers. The ACL, IP-OSPF, and/or IP-RIP message types are also sent to the syslog servers if they are currently enabled as debug types. (See **Debug messages** on page 338.)

<code>no logging</code>	Removes all currently configured syslog logging destinations from the running configuration. Using this form of the command to delete the only remaining syslog server address disables debug destination logging on the switch, but the default Event debug type does not change.
<code>no logging <syslog-ip-address></code>	Removes only the specified syslog logging destination from the running configuration. Removing all configured syslog destinations with the <code>no logging</code> command (or a specified syslog server destination with the <code>no logging <syslog-ip-address></code> command) does not delete the syslog server IP addresses stored in the startup configuration.

Deleting syslog addresses in the startup configuration

Enter a `no logging` command followed by the `write memory` command.

Verifying the deletion of a syslog server address

Display the startup configuration by entering the `show config` command.

Blocking the messages sent to configured syslog servers from the currently configured debug message type

Enter the `no debug <debug-type>` command. (See **Debug messages** on page 338.)

Disabling syslog logging on the switch without deleting configured server addresses

Enter the `no debug destination logging` command. Note that, unlike the case in which no syslog servers are configured, if one or more syslog servers are already configured and syslog messaging is disabled, configuring a new server address does not re-enable syslog messaging. To re-enable syslog messaging, you must enter the `debug destination logging` command.

Sending logging messages using TCP

Syntax:

```
[no] logging <ip-addr> [udp 1024-49151 | tcp 1024-49151]
```

Allows the configuration of the UDP or TCP transport protocol for the transmission of logging messages to a syslog server.

Specifying a destination port with UDP or TCP is optional.

Default ports: UDP port is 514

TCP port is 1470

Default Transport Protocol: UDP

Because TCP is a connection-oriented protocol, a connection must be present before the logging information is sent. This helps ensure that the logging message will reach the syslog server. Each configured syslog server needs its own connection. You can configure the destination port that is used for the transmission of the logging messages.

Configuring TCP for logging message transmission using the default port

```
switch(config)# logging 192.123.4.5 tcp
```

(Default TCP port 1470 is used.)

Configuring TCP for logging message transmission using a specified port

```
switch(config)# logging 192.123.4.5 9514
```

(TCP port 9514 is used.)

Configuring UDP for logging message transmission using the default port

```
switch(config)# logging 192.123.4.5 udp
```

(Default UDP port 514 is used.)

Configuring UDP for logging message transmission using a specified port

```
switch(config)# logging 192.123.4.5 9512
```

(UDP port 9512 is used.)

Syntax:

[no] logging facility <facility-name>

The logging facility specifies the destination subsystem used in a configured syslog server. (All configured syslog servers must use the same subsystem.) Hewlett Packard Enterprise recommends the default (user) subsystem unless your application specifically requires another subsystem. Options include:

user	(default) Random user-level messages
kern	Kernel messages
mail	Mail system
daemon	System daemons
auth	Security/authorization messages
syslog	Messages generated internally by syslog
lpr	Line-printer subsystem
news	Netnews subsystem
uucp	uucp subsystem

Table Continued

<code>cron</code>	cron/at subsystem
<code>sys9</code>	cron/at subsystem
<code>sys10 - sys14</code>	Reserved for system use
<code>local10 - local17</code>	Reserved for system use

Use the `no` form of the command to remove the configured facility and reconfigure the default (user) value.

Disable LinkUp/Down Syslog messages based on port

This feature provides a per-port basis filter that can restrict the logging of events that are associated with a link status change. Unimportant linkup/linkdown events can be filtered out, avoiding unwanted messages in the event log and reducing troubleshooting time.

The specific port-based events to be controlled are:

RMON_PMGR_PORT_UP—Indicates that the port has changed from and off-line to an on-line state. To be online the port must be both connected to the LAN and enabled through configuration.

RMON_PMGR_PORT_DOWN—Indicates that the port has changed from an on-line state to an offline state. For this state to occur, the port is physically disconnected from the LAN, disabled through the configuration, or both.

The following rules apply:

- Only one filter can be enabled at a time.
- The maximum number of configured filters is 10.
- A filter is identified by a unique name of up to 16 printable ASCII characters.
- Filters can be dynamically replaced; the newly enabled filter automatically disables the previous filter.
- A filter always contains a default sub-filter that functions as the filtering rules terminator.
- To apply filtering to an event logging process, the filter must be explicitly enabled from the CLI.
- Enabled filter modules can be dynamically modified; the changes will take effect immediately.

A filter module may include up to 19 option sub-filters and a default sub-filter. The sub-filter types are:

- **Severity**—checks the severity level of the event log message. The severity values are:
 - major
 - warning
 - error
 - info
 - debug
- **Event number**—Checks the event number of the event log message.
- **Regular expression**—Checks everything beyond the date/time portion of the event log message.

A sub-filter has a sequence number, criteria to be matched, and a resulting action when a match occurs. All of the parameters must be specified in order to create the sub-filter.

- **Sequence number:** Used for the ordering of sub-filters. Range 1-98.
- **Matching criteria:** Can be the severity level, event number, or a regular expression.
- **Action to execute:** When a match occurs, the resulting action is either permit the logging of the event, or deny the logging of the event.

The following sub-filter rules apply:

- Up to 19 optional sub-filters and a default sub-filter are allowed in a filter module.
- Sub-filters in the filter module can be of the same or different types.
- Sub-filter entries can be modified with new criteria and action definitions.

- Sub-filters are executed from the lowest sequence number to the highest. As soon as a match is found the log event is immediately accepted or rejected and no further matching operation is performed.
- The default sub-filter must always be the last entry in a filter module. It functions as the rules terminator when the criteria matching performed by the prior sub-filters in a filter does not produce an action.
- The default sub-filter cannot be deleted, re-ordered, or changed. The only parameter that can be modified is the action parameter of permit or deny. The default is permit.

Creating a filter

Syntax:

```
[no] logging filter <name> <sequence> [severity <severity>|event-num <num>|<regexp>] [permit|deny]
```

Creates a logging filter to restrict which events are logged. The no form of the command removes the logging filter.

<name>: The name that identifies the filter.

severity <severity>: Specifies the severity of an event—major, warning, error, info, or debug.

event-num <num>: Specifies an event number to match.

deny: If the log entry matches the specified criteria, do not log the event message. No further criteria are evaluated for a match.

permit: If the log entry matches the specified criteria, log the event message. No further criteria are evaluated for a match.

Enabling a Filter after Creation

Syntax:

```
[no] logging filter <name> enable | disable
```

Enables a log filter. Only one filter can be enabled at a time. An enabled filter automatically disables a previously enabled filter.

<name>: The name that identifies the filter.

Clearing a Filter

Syntax:

```
[no] clear logging filter <name|all>
```

Clears statistics counters for the named logging filter or for all filters.

Viewing Filter Configuration Information

Syntax:

```
show logging filter name
```

Displays the logging filter's configuration information. The Matches column indicates the number of times that criteria has matched.

Specifying the criteria for a filter and then enabling the filter

```
switch(config)# logging filter SevWarnFatal 10 severity warning permit
switch(config)# logging filter SevWarnFatal 20 severity major permit
switch(config)# logging filter SevWarnFatal default deny

switch(config)# logging filter SevWarnFatal enable
```

1. The filter named SevWarnFatal adds a sub-filter of the severity type, with a sequence number of 10. The sub-filter specifies that a match for an event log message with a severity of “warning” will be logged.
2. The second sub-filter has a sequence number of 20 and a severity type of major. The sub-filter specifies that a match for an event log message with a severity of “major” will be logged.
3. The default sub-filter, which is created automatically at the time of filter creation, is always the last entry in the filter module. It matches “anything” and cannot be changed. You can change the actions to either permit or deny. This example specifies that any message that did not meet the prior matching criteria will not be logged.
4. The last command enables the filter named SevWarnFatal. If there was another filter enabled, this filter automatically replaces it and the other filter is disabled.

Specifying the criteria for a filter named noUpDownEvents and then enabling the filter

```
switch(config)# logging filter noUpDownEvent 10 event-num 76 deny
switch(config)# logging filter noUpDownEvent 20 event-num 77 deny
switch(config)# logging filter noUpDownEvent default permit

switch(config)# logging filter noUpDownEvent enable
```

1. The filter named noUpDownEvents adds a sub-filter with a type of event-num, and a sequence number of 10. The sub-filter specifies that a match for an event log message with an event number of “76” will not be logged.
2. The second sub-filter has a sequence number of 20 and a type of event-num. The sub-filter specifies that a match for an event log message with an event number of “77” will not be logged.
3. The default sub-filter, which is created automatically at the time of filter creation, is always the last entry in the filter module. It matches “anything” and cannot be changed. You can change the actions to either permit or deny. This example specifies that any message that did not meet the prior matching criteria will be logged.
4. The last command enables the filter named noUpDownEvents. If there was another filter enabled, this filter automatically replaces it and the other filter is disabled.

Specifying the criteria for a match using a regular expression and then enabling the filter

```
switch(config)# logging filter noUpPorts 10 "(A10|A22|B5) is now on-line" deny
switch(config)# logging filter noUpPorts default permit

switch(config)# logging filter noUpPorts enable
```

This example denies logging of the matching regular expression “port <port-num> is now on-line” for ports A10, A22, and B5.

1. The filter named noUpPorts adds a sub-filter with a type of regular expression for ports A10, A22, and B5. The sub-filter specifies the matching criteria for the regular expression and if there is a match, the event log message is not logged.
2. The default sub-filter specifies that any message that did not meet the prior matching criteria will be logged.
3. The last command enables the filter named noUpPorts.

Specifying the criteria for a match using a regular expression for specific ports

```
switch(config)# logging filter noStpBlockPorts 10 "(A[1-9]|A10|B[1-4])
.*Blocked by STP" permit
```

```
switch(config)# logging filter noStpBlockPorts 20 ".*Blocked by STP" deny
switch(config)# logging filter noStpBlockPorts default permit
```

```
switch(config)# logging filter noStpBlockPorts enable
```

1. The filter named noStpBlockPorts adds a sub-filter with a type of regular expression with a sequence number of 10. This rule specifies that event messages from ports A1-A10, and B1-B4 with the “.*Blocked by STP” expression pattern in the message body are logged.
2. The second command adds a sub-filter with a type of regular expression and a sequence number of 20. This rule specifies that event messages generated from any ports with the “.*Blocked by STP” expression pattern in the message body are not logged.
3. The default sub-filter specifies that any message that did not meet the prior matching criteria will be logged.
4. The last command enables the filter named noStpBlockPorts.

Output examples:

The configured logging filters

```
HP Switch# show logging filter
```

```
Status and Counters - Log Filters Information
```

Name	Enabled
noUpPorts	No
SevWarnFatal	No
noUpDownEvents	No
noStpBlockPorts	Yes

Output for specified logging filters

```
HP Switch# show logging filter sevWarnFatal
```

```
Status and Counters - Log Filters Information
```

```
Name           : Enabled
Enabled        : Yes
Messages Dropped : 0
```

Seq	Type	Value	Action	Matches
10	Severity	warning	Permit	2
20	Severity	major	Permit	2
def		(any)	Deny	0

```
switch(config)# show logging filter noStpBlockPorts
```

```
Status and Counters - Log Filters Information
```

```
Name           : noStpBlockPorts
Enabled        : Yes
Messages Dropped : 0
```

Seq	Type	Value	Action	Matches
10	RegExp	(A[1-9] A10 B[1-4]).*Blocked by STP	Permit	2
20	RegExp	.*Blocked by STP	Deny	2
def		(any)	Permit	0

Output of running-config file

```
HP Switch# show running-config

Running configuration:

; J9470A Configuration Editor; Created on release #XX.15.13.0000x
; Ver #04:0f.ff.3f.ef:24
hostname "HP Switch"
module 1 type j94dda
logging filter "noUpPorts" 10 "(A10|A22|B5) is now on-line" deny
logging filter "noUpPorts" default permit
logging filter "SevWarnFatal" 10 severity warning permit
logging filter "SevWarnFatal" 20 severity major permit
logging filter "SevWarnFatal" default deny
logging filter "noUpDownEvent" 10 event-num 76 deny
logging filter "noUpDownEvent" 20 event-num 77 deny
logging filter "noUpDownEvent" default permit
logging filter "noStpBlockPorts" 10 "(A[1-9]|A10|B[1-4]) .*Blocked by STP" permit
logging filter "noStpBlockPorts" 20 " .*Blocked by STP" deny
logging filter "noStpBlockPorts" default permit
logging filter "noStpBlockPorts" enable
snmp-server community "public" unrestricted
snmp-server host 15.255.133.156 community "public"
snmp-server host 15.255.133.146 community "public"
vlan 1
.
.
.
```

Adding a description for a Syslog server

You can associate a user-friendly description with each of the IP addresses (IPv4 only) configured for syslog using the CLI or SNMP.



The Hewlett Packard Enterprise MIB `hpicfSyslog.mib` allows the configuration and monitoring of syslog for SNMP (RFC 3164 supported).



Entering the `no logging` command removes ALL the syslog server addresses without a verification prompt.

The CLI command is:

Syntax:

```
logging <ip-addr> [control-descr ZZZZTRISHZZZZ <text_string>]
no logging <ip-addr> [control-descr]
```

An optional user-friendly description that can be associated with a server IP address. If no description is entered, this is blank. If `<text_string>` contains white space, use quotes around the string. IPv4 addresses only.

Use the `no` form of the command to remove the description. Limit: 255 characters



To remove the description using SNMP, set the description to an empty string.

The logging command with a control description

```
switch(config)# logging 10.10.10.2 control-descr syslog_one
```

Adding a priority description

This description can be added with the CLI or SNMP. The CLI command is:

Syntax:

```
logging priority-descr <text_string>  
no logging priority-descr
```

Provides a user-friendly description for the combined filter values of `severity` and `system module`. If no description is entered, this is blank.

If `text_string` contains white space, use quotes around the string.

Use the `no` form of the command to remove the description.

Limit: 255 characters

The logging command with a priority description

```
switch(config)# logging priority-descr severe-pri
```



A notification is sent to the SNMP agent if there are any changes to the syslog parameters, either through the CLI or with SNMP.

Configuring the severity level for Event Log messages sent to a syslog server

Event Log messages are entered with one of the following severity levels (from highest to lowest):

Major	A fatal error condition has occurred on the switch.
Error	An error condition has occurred on the switch.
Warning	A switch service has behaved unexpectedly.
Information	Information on a normal switch event.
Debug	Reserved for HPE switch internal diagnostic information.

Using the `logging severity` command, you can select a set of Event Log messages according to their severity level and send them to a syslog server. Messages of the selected and higher severity will be sent. To configure a syslog server, see [Configuring a syslog server](#) on page 341.

Syntax:

```
[no] logging severity {< major | error | warning | info | debug >}
```

Configures the switch to send all Event Log messages with a severity level equal to or higher than the specified value to all configured Syslog servers.

Default: `debug` (Reports messages of all severity levels.)

Use the `no` form of the command to remove the configured severity level and reconfigure the default value, which sends Event Log messages of all severity levels to syslog servers.



The severity setting does not affect event notification messages that the switch normally sends to the Event Log. All messages remain recorded in the Event Log.

Configuring the system module used to select the Event Log messages sent to a syslog server

Event Log messages contain the name of the system module that reported the event. Using the `logging system-module` command, you can select a set of Event Log messages according to the originating system module and send them to a syslog server.

Syntax:

```
[no] logging system-module <system-module>
```

Configures the switch to send all Event Log messages being logged from the specified system module to configured syslog servers. (To configure a syslog server, see [Configuring a syslog server](#).)

See [Event Log system modules](#) for the correct value to enter for each system module.

Default: `all-pass` (Reports all Event Log messages.)

Use the `no` form of the command to remove the configured system module value and reconfigure the default value, which sends Event Log messages from all system modules to syslog servers.

You can select messages from only one system module to be sent to a syslog server; you cannot configure messages from multiple system modules to be sent. If you re-enter the command with a different system module name, the currently configured value is replaced with the new one.



This setting has no effect on event notification messages that the switch normally sends to the Event Log.

Operating notes for debug and Syslog

- Rebooting the switch or pressing the `Reset` button resets the debug configuration.

Debug option	Effect of a reboot or reset
logging (debug destination)	If syslog server IP addresses are stored in the startup-config file, they are saved across a reboot and the logging destination option remains enabled. Otherwise, the logging destination is disabled.
session (debug destination)	Disabled.
ACL (debug type)	Disabled.
All (debug type)	Disabled.

Table Continued

Debug option	Effect of a reboot or reset
event (debug type)	<p>If a syslog server IP address is configured in the startup-config file, the sending of Event Log messages is reset to <code>enabled</code>, regardless of the last active setting.</p> <p>If no syslog server is configured, the sending of Event Log messages is <code>disabled</code>.</p>
IP (debug type)	Disabled.

- Debug commands do not affect normal message output to the Event Log.

Using the `debug event` command, you can specify that Event Log messages are sent to the debug destinations you configure (CLI session, syslog servers, or both) in addition to the Event Log.

- Ensure that your syslog servers accept debug messages.

All syslog messages resulting from a debug operation have a "debug" severity level. If you configure the switch to send debug messages to a syslog server, ensure that the server's syslog application is configured to accept the "debug" severity level. (The default configuration for some syslog applications ignores the "debug" severity level.)

- Duplicate IP addresses are not stored in the list of syslog servers.
- If the default severity value is in effect, all messages that have severities greater than the default value are passed to syslog. For example, if the default severity is "debug," all messages that have severities greater than debug are passed to syslog.
- There is a limit of six syslog servers. All syslog servers are sent the same messages using the same filter parameters. An error is generated for an attempt to add more than six syslog servers.

Diagnostic tools

Port auto-negotiation

When a link LED does not light (indicating loss of link between two devices), the most common reason is a failure of port auto-negotiation between the connecting ports. If a link LED fails to light when you connect the switch to a port on another device, do the following:

Procedure

1. Ensure that the switch port and the port on the attached end-node are both set to `Auto` mode.
2. If the attached end-node does not have an `Auto` mode setting, you must manually configure the switch port to the same setting as the end-node port. See [Port Status and Configuration](#) on page 53.

Ping and link tests

The ping test and the link test are point-to-point tests between your switch and another IEEE 802.3-compliant device on your network. These tests can tell you whether the switch is communicating properly with another device.



To respond to a ping test or a link test, the device you are trying to reach must be IEEE 802.3-compliant.

Ping test

A test of the path between the switch and another device on the same or another IP network that can respond to IP packets (ICMP Echo Requests). To use the `ping` (or `tracert`) command with host names or fully qualified domain names, see [DNS resolver](#) on page 365.

Link test

A test of the connection between the switch and a designated network device on the same LAN (or VLAN, if configured). During the link test, IEEE 802.2 test packets are sent to the designated network device in the same VLAN or broadcast domain. The remote device must be able to respond with an 802.2 Test Response Packet.

Executing ping or link tests (WebAgent)

To start a ping or link test in the WebAgent:

1. In the navigation pane, click **Troubleshooting**.
2. Click **Ping/Link Test**.
3. Click **Start**.
4. To halt a link or ping test before it concludes, click **Stop**.

For an Example: of the text screens, see [Figure 65: Ping test and link test screen on the WebAgent](#) on page 352.

Figure 65: Ping test and link test screen on the WebAgent

The screenshot displays two configuration windows from the WebAgent. The top window is titled 'Ping Test' and contains a 'Ping Status' section with three input fields: 'Destination IP Address' (empty), 'Number of Packets' (set to 5), and 'Time Out in Seconds' (set to 1). The bottom window is titled 'Link Test' and contains a 'Link Status' section with four input fields: 'Destination MAC Address' (empty), 'VLAN' (a dropdown menu), 'Number of Packets' (set to 5), and 'Time Out in Seconds' (set to 1). Both windows have 'Start', 'Stop', and help icons in the top right corner.

Destination IP Address is the network address of the target, or destination, device to which you want to test a connection with the switch. An IP address is in the X.X.X.X format where X is a decimal number between 0 and 255.

Number of Packets to Send is the number of times you want the switch to attempt to test a connection.

Timeout in Seconds is the number of seconds to allow per attempt to test a connection before determining that the current attempt has failed.

Testing the path between the switch and another device on an IP network

The ping test uses ICMP echo requests and ICMP echo replies to determine if another device is alive. It also measures the amount of time it takes to receive a reply from the specified destination. The `ping` command has several extended commands that allow advanced checking of destination availability.

Syntax:

```
ping {<ip-address | hostname | switch-num>} [repetitions <1-10000>] [timeout <1-60>] [{source <ip-address> | <vlan-id>}] [data-size <0-65471>] [data-fill <0-1024>]
```

```
ping6 {<ip-address | hostname | [switch-num>]} [repetitions <1-10000>] [timeout <1-60>] [{source <ip-address> | <vlan-id>}] [data-size <0-65471>] [data-fill <0-1024>]
```

Sends ICMP echo requests to determine if another device is alive.

<code>{< ip-address hostname >}</code>	Target IP address or hostname of the destination node being pinged
<code>repetitions <1-10000></code>	Number of ping packets sent to the destination address. Default: 1
<code>timeout <1-60></code>	Timeout interval in seconds; the ECHO REPLY must be received before this time interval expires for the ping to be successful. Default: 5
<code>source {<ip-addr hostname>}</code>	Source IP address or hostname. The source IP address must be owned by the router. If a VLAN is specified, the IP address associated with the specified VLAN is used.
<code>data-size <0-65471></code>	Size of packet sent. Default: 0 (zero)
<code>data-fill <0-1024></code>	The data pattern in the packet. Default: Zero length string

Ping tests

```
switch# ping 10.10.10.10
10.10.10.10 is alive, time = 15 ms

switch# ping 10.10.10.10 repetitions 3
10.10.10.10 is alive, iteration 1, time = 15 ms
10.10.10.10 is alive, iteration 1, time = 15 ms
10.10.10.10 is alive, iteration 1, time = 15 ms

switch# ping 10.10.10.10 timeout 2
10.10.10.10 is alive, time = 10 ms

switch# ping 10.11.12.13
The destination address is unreachable.
```

Halting a ping test

To halt a ping test before it concludes, press **[Ctrl] [C]**.



To use the `ping` (or `traceroute`) command with host names or fully qualified domain names, see [DNS resolver](#) on page 365.

Issuing single or multiple link tests

Single or multiple link tests can have varying repetitions and timeout periods. The defaults are:

- Repetitions: 1 (1 to 999)
- Timeout: 5 seconds (1 to 256 seconds)

Syntax:

```
link <mac-address> [repetitions <1-999>] [timeout <1-256>] [vlan <vlan-id >]
```

Example:

Figure 66: *Link tests*

Basic Link Test	HP Switch# link 0030c1-7fcc40 Link-test passed.
Link Test with Repetitions	HP Switch# link 0030c1-7fcc40 repetitions 3 802.2 TEST packets sent: 3, responses received: 3
Link Test with Repetitions and Timeout	HP Switch# link 0030c1-7fcc40 repetitions 3 timeout 1 802.2 TEST packets sent: 3, responses received: 3
Link Test Over a Specific VLAN	HP Switch# link 0030c1-7fcc40 repetitions 3 timeout 1 vlan 1 802.2 TEST packets sent: 3, responses received: 3
Link Test Over a Specific VLAN; Test Fail	HP Switch# link 0030c1-7fcc40 repetitions 3 timeout 1 vlan 222 802.2 TEST packets sent: 3, responses received: 0

Tracing the route from the switch to a host address

The `traceroute` command enables you to trace the route from the switch to a host address.

This command outputs information for each (router) hop between the switch and the destination address. Note that every time you execute `traceroute`, it uses the same default settings unless you specify otherwise for that instance of the command.

Syntax:

```
traceroute {<ip-address | hostname>}
```

```
traceroute6 {<ip-address | hostname>}
```

Lists the IP address or hostname of each hop in the route, plus the time in microseconds for the `traceroute` packet reply to the switch for each hop.

{< ip-address hostname >}	The IP address or hostname of the device to which to send the traceroute.
[minttl < 1-255 >]	<p>For the current instance of <code>traceroute</code>, changes the minimum number of hops allowed for each probe packet sent along the route.</p> <ul style="list-style-type: none"> • If <code>minttl</code> is greater than the actual number of hops, the output includes only the hops at and above the <code>minttl</code> threshold. (The hops below the threshold are not listed.) • If <code>minttl</code> matches the actual number of hops, only that hop is shown in the output. • If <code>minttl</code> is less than the actual number of hops, all hops are listed. <p>For any instance of <code>traceroute</code>, if you want a <code>minttl</code> value other than the default, you must specify that value.(Default: 1)</p>
[maxttl < 1-255 >]	<p>For the current instance of <code>traceroute</code>, changes the maximum number of hops allowed for each probe packet sent along the route.If the destination address is further from the switch than <code>maxttl</code> allows, <code>traceroute</code> lists the IP addresses for all hops it detects up to the <code>maxttl</code> limit.For any instance of <code>traceroute</code>, if you want a <code>maxttl</code> value other than the default, you must specify that value.(Default: 30)</p>
[timeout < 1-120 >]	<p>For the current instance of <code>traceroute</code>, changes the timeout period the switch waits for each probe of a hop in the route. For any instance of <code>traceroute</code>, if you want a <code>timeout</code> value other than the default, you must specify that value.Default: 5 seconds</p>
[probes < 1-5 >]	<p>For the current instance of <code>traceroute</code>, changes the number of queries the switch sends for each hop in the route.For any instance of <code>traceroute</code>, if you want a <code>probes</code> value other than the default, you must specify that value.(Default: 3)</p>
[[source <ip-addr> [vlan-id>]]	<p>The source IP address or VLAN. The source IP address must be owned by the router. If a VLAN is specified, the IP address associated with the specified VLAN is used.</p>



For information about `traceroute6`, see the IPv6 configuration guide for your switch.

Halting an ongoing traceroute search

Press the **[Ctrl] [C]** keys.

A low maxttl causes traceroute to halt before reaching the destination address

Executing `traceroute` with its default values for a destination IP address that is four hops away produces a result similar to this:

Figure 67: A completed traceroute enquiry

```

HP Switch# traceroute 125.25.24.35
traceroute to 125.25.24.35 ,
          1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2      0 ms      0 ms      0 ms
 2 10.71.217.2      7 ms      3 ms      0 ms
 3 10.243.170.1     0 ms      1 ms      0 ms
 4 125.25.24.35    3 ms      3 ms      0 ms
  
```

Intermediate router hops with the time taken for the switch to receive acknowledgement of each probe reaching each router.

Destination IP Address

Continuing from the previous Example: (**Figure 67: A completed traceroute enquiry** on page 356), executing `traceroute` with an insufficient `maxttl` for the actual hop count produces an output similar to this:

Figure 68: Incomplete traceroute because of low maxttl setting

```

Switch# traceroute 125.25.24.35(maxttl 3)
traceroute to 125.25.24.35 ,
          1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2      0 ms      0 ms      0 ms
 2 10.71.217.2      0 ms      0 ms      0 ms
 3 10.243.170.1     0 ms *    0 ms
  
```

Traceroute does not reach destination IP address because of low maxttl setting.

The asterisk indicates there was a timeout on the second probe to the third hop.

If a network condition prevents traceroute from reaching the destination

Common reasons for `traceroute` failing to reach a destination include:

- Timeouts (indicated by one asterisk per probe, per hop)
- Unreachable hosts
- Unreachable networks
- Interference from firewalls
- Hosts configured to avoid responding

Executing `traceroute` where the route becomes blocked or otherwise fails results in an output marked by timeouts for all probes beyond the last detected hop. For example, with a maximum hop count of 7 (`maxttl = 7`), where the route becomes blocked or otherwise fails, the output appears similar to this:

Figure 69: Traceroute failing to reach the destination address

```

HP Switch# traceroute 125.25.24.35 maxttl 7
traceroute to 107.64.197.100 ,
          1 hop min, 7 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2      0 ms      0 ms      0 ms
 2 10.71.217.2      0 ms      0 ms      0 ms
 3 * 10.243.170.1    0 ms *
 4 * * * *
 5 * * * *
 6 * * * *
 7 * * * *
  
```

At hop 3, the first and third probes timed out but the second probe reached the router. All further probes within the maxttl timed-out without finding a router or the destination IP address.

An asterisk indicates a timeout without finding the next hop.

Viewing switch configuration and operation

In some troubleshooting scenarios, you may need to view the switch configuration to diagnose a problem. The complete switch configuration is contained in a file that you can browse from the CLI using the commands described in this section.

Viewing the startup or running configuration file

Syntax:

```
write terminal
```

Displays the running configuration.

<code>show config</code>	Displays the startup configuration.
<code>show running-config</code>	Displays the running-config file.

For more information and examples of how to use these commands, see “Switch Memory and Configuration” in the basic operation guide.

Viewing the configuration file (WebAgent)

To display the running configuration using the WebAgent:

1. In the navigation pane, click **Troubleshooting**.
2. Click **Configuration Report**.
3. Use the right-side scroll bar to scroll through the configuration listing.

Viewing a summary of switch operational data

Syntax:

```
show tech
```

By default, the `show tech` command displays a single output of switch operating and running-configuration data from several internal switch sources, including:

- Image stamp (software version data)
- Running configuration
- Event Log listing
- Boot history
- Port settings
- Status and counters — port status
- IP routes
- Status and counters — VLAN information
- GVRP support
- Load balancing (trunk and LACP)

The `show tech` command on page 358 shows sample output from the `show tech` command.

The show tech command

```
switch# show tech

show system

Status and Counters - General System Information

System Name       : Switch
System Contact    :
System Location   :

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None

Software revision : XX.14.xx      Base MAC Addr  : 001871-c42f00
ROM Version       : XX.12.12     Serial Number  : SG641SU00L

Up Time           : 23 hours     Memory - Total :
CPU Util (%)      : 10           Free          :

IP Mgmt - Pkts Rx : 759          Packet - Total : 6750
                Pkts Tx : 2          Buffers Free  : 5086
                                           Lowest        : 4961
                                           Missed        : 0

show flash
Image      Size(Bytes)  Date  Version
-----
-----
```

To specify the data displayed by the `show tech` command, use the `copy show tech` command as described in .

Saving show tech command output to a text file

When you enter the `show tech` command, a summary of switch operational data is sent to your terminal emulator. You can use your terminal emulator's text capture features to save the `show tech` data to a text file for viewing, printing, or sending to an associate to diagnose a problem.


For example, if your terminal emulator is the Hyperterminal application available with Microsoft® Windows® software, you can copy the `show tech` output to a file and then use either Microsoft Word or Notepad to display the data. (In this case, Microsoft Word provides the data in an easier-to-read format.)

The following example uses the Microsoft Windows terminal emulator. If you are using a different terminal emulator application, see the documentation provided with the application.

Procedure


1. In Hyperterminal, click on `Transfer|Capture Text...`

Figure 70: Capture text window of the Hyperterminal application



2. In the `File` field, enter the path and file name in which you want to store the `show tech` output.

Figure 71: Entering a path and filename for saving `show tech` output



3. Click **[Start]** to create and open the text file.
4. From the global configuration context, enter the `show tech` command:

```
switch# show tech
```

The `show tech` command output is copied into the text file and displayed on the terminal emulator screen. When the command output stops and displays `-- MORE --`, press the Space bar to display and copy more information. The CLI prompt appears when the command output finishes.

5. Click on `Transfer|Capture Text|Stop` in HyperTerminal to stop copying data and save the text file.

If you do not stop HyperTerminal from copying command output into the text file, additional unwanted data can be copied from the HyperTerminal screen.

6. To access the file, open it in Microsoft Word, Notepad, or a similar text editor.

Viewing more information on switch operation

Use the following commands to display additional information on switch operation for troubleshooting purposes.

Syntax:

```
show boot-history
```

Displays the crash information saved for each management module on the switch.

```
show history
```

Displays the current command history. This command output is used for reference or when you want to repeat a command (See [Displaying the information you need to diagnose problems](#) on page 362).

`show system-information`

Displays globally configured parameters and information on switch operation.

`show version`

Displays the software version currently running on the switch and the flash image from which the switch booted (primary or secondary). For more information, see "Displaying Management Information" in the "Redundancy (Switch 8212zl)".

`show interfaces`

Displays information on the activity on all switch ports (see "Viewing Port Status and Configuring Port Parameters" in the "Port Status and Configuration").

`show interfaces-display`

Displays the same information as the `show interfaces` command and dynamically updates the output every three seconds. Press **Ctrl + C** to stop the dynamic updates of system information. Use the Arrow keys to view information that is off the screen.

Searching for text using pattern matching with show command

Selected portions of the output are displayed, depending on the parameters chosen.

Syntax:

```
show {< command option > | < include | exclude | begin >} <regular expression>
```

Uses matching pattern searches to display selected portions of the output from a `show` command. There is no limit to the number of characters that can be matched. Only regular expressions are permitted; symbols such as the asterisk cannot be substituted to perform more general matching.

<code>include</code>	Only the lines that contain the matching pattern are displayed in the output.
<code>exclude</code>	Only the lines that contain the matching pattern are not displayed in the output.
<code>begin</code>	The display of the output begins with the line that contains the matching pattern.



Pattern matching is case-sensitive.

Following are examples of what portions of the running config file display depending on the option chosen.

Pattern matching with include option

```
switch(config)# show run | include ipv6 1
  ipv6 enable
```



```
    ipv6 enable
ipv6 access-list "EH-01"
switch(config)#
```

- ¹Displays only lines that contain “ipv6”.

Pattern matching with exclude option

```
switch(config)# show run | exclude ipv6 1

Running configuration:

; J9299A Configuration Editor; Created on release #YA.15.XX
; Ver #01:01:00

hostname "HP Switch"
mirror-port 4
qos dscp-map 000000 priority 0
qos dscp-map 001000 priority 1
...
interface 1
    lacp Active
exit
interface 6
    name "Print_Server"
exit
...
vlan 1
    name "DEFAULT_VLAN"
    untagged 1-24
    ip address dhcp-bootp
    exit
...
port-security 17 learn-mode static address-limit 3 action send-alarm
power-over-ethernet pre-std-detect
no ip ssh cipher 3des-cbc
ip timep dhcp
snmp-server community "public" unrestricted
snmp-server community "public" unrestricted
...
vlan 1
    exit
```

- ¹Displays all lines that do not contain “ipv6”.

Pattern matching with begin option

```
switch(config)# show run | begin ipv6 1
    ipv6 enable
    no untagged 21-24
    exit
vlan 20
    name "VLAN20"
    untagged 21-24
    ipv6 enable
    no ip address
    exit
policy qos "michael"
    exit
ipv6 access-list "EH-01"
    sequence 10 deny tcp 2001:db8:255::/48 2001:db8:125::/48
    exit
```

```
no autorun
password manager
```

- ¹Displays the running config beginning at the first line that contains “ipv6”.

The following is an Example: of the `show arp` command output, and then the output displayed when the `include` option has the IP address of `15.255.128.1` as the regular expression.

The show arp command and pattern matching with the include option

```
switch(config)# show arp
```

```
IP ARP table
```

IP Address	MAC Address	Type	Port
15.255.128.1	00000c-07ac00	dynamic	B1
15.255.131.19	00a0c9-b1503d	dynamic	
15.255.133.150	000bcd-3cbeec	dynamic	B1

```
switch(config)# show arp | include 15.255.128.1
15.255.128.1    00000c-07ac00    dynamic B1
```

Displaying the information you need to diagnose problems

Use the following commands in a troubleshooting session to more accurately display the information you need to diagnose a problem.

Syntax:

```
alias
```

Creates a shortcut alias name for commonly used commands and command options.

Syntax:

```
kill
```

Terminates a currently running, remote troubleshooting session. Use the `show ip ssh` command to list the current management sessions.

Syntax:

```
[no] page
```

Toggles the paging mode for `show` commands between continuous listing and per-page listing.

Syntax:

```
repeat
```

Repeatedly executes one or more commands so that you can see the results of multiple commands displayed over a period of time. To halt the command execution, press any key on the keyboard.

Syntax:

```
setup
```

Displays the Switch Setup screen from the menu interface.

Restoring the factory-default configuration

As part of your troubleshooting process, it may become necessary to return the switch configuration to the factory default settings. This process:

- Momentarily interrupts the switch operation
- Clears any passwords
- Clears the console Event Log
- Resets the network counters to zero
- Performs a complete self test
- Reboots the switch into its factory default configuration, including deleting an IP address

There are two methods for resetting to the factory-default configuration:

- CLI
- `Clear/Reset` button combination



Hewlett Packard Enterprise recommends that you save your configuration to a TFTP server before resetting the switch to its factory-default configuration. You can also save your configuration via Xmodem to a directly connected PC.

Resetting to the factory-default configuration

Using the CLI

This command operates at any level **except** the Operator level.

Syntax:

```
erase startup-configuration
```

Deletes the startup-config file in flash so that the switch will reboot with its factory-default configuration.



The `erase startup-config` command does not clear passwords unless `include-credentials` has been set, at which time this command does erase username/password information and any other credentials stored in the config file. For more information, see the section on "Saving Security Credentials in a Config File" in the access security guide for your switch.

Using Clear/Reset

Procedure

1. Using pointed objects, simultaneously press both the `Reset` and `Clear` buttons on the front of the switch.
2. Continue to press the `Clear` button while releasing the `Reset` button.
3. When the Self Test LED begins to flash, release the `Clear` button.

The switch then completes its self test and begins operating with the configuration restored to the factory default settings.

Restoring a flash image

The switch can lose its operating system if either the primary or secondary flash image location is empty or contains a corrupted OS file and an operator uses the `erase flash` command to erase a good OS image file from the opposite flash location.

Recovering from an empty or corrupted flash state

Use the switch's console serial port to connect to a workstation or laptop computer that has the following:

- A terminal emulator program with Xmodem capability, such as the HyperTerminal program included in Windows PC software.
- A copy of a good OS image file for the switch



The following procedure requires the use of Xmodem and copies an OS image into primary flash only.

This procedure assumes you are using HyperTerminal as your terminal emulator. If you use a different terminal emulator, you may need to adapt this procedure to the operation of your particular emulator.

1. Start the terminal emulator program.

Ensure that the terminal program is configured as follows:

- Baud rate: 9600
- No parity
- 8 Bits
- 1 stop bit
- No flow control

2. Use the `Reset` button to reset the switch.

The following prompt should then appear in the terminal emulator:

```
Enter h or ? for help.
```

```
=>
```

3. Because the OS file is large, you can increase the speed of the download by changing the switch console and terminal emulator baud rates to a high speed. For Example:

- a. Change the switch baud rate to 115,200 Bps.

```
=> sp 115200
```

- b. Change the terminal emulator baud rate to match the switch speed:

- I. In HyperTerminal, select **Call|Disconnect**.
- II. Select **File|Properties**.
- III. Click on **Configure**.
- IV. Change the baud rate to **115200**.
- V. Click on **[OK]**, then in the next window, click on **[OK]** again.
- VI. Select **Call|Connect**.
- VII. Press **[Enter]** one or more times to display the => prompt.

4. Start the Console Download utility by entering `do` at the => prompt and pressing **[Enter]**:

```
=> do
```

5. You then see this prompt:

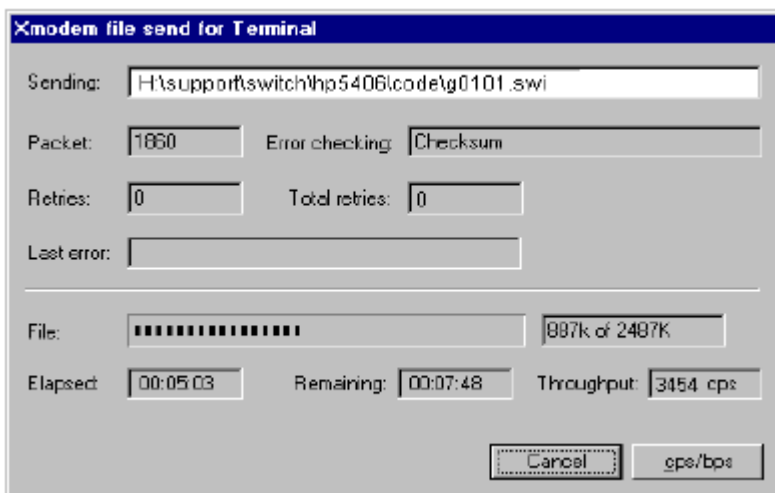
```
You have invoked the console download utility.  
Do you wish to continue? (Y/N)>_
```

6. At the above prompt:

- a. Enter **y** (for Yes)
- b. Select **Transfer|File** in HyperTerminal.
- c. Enter the appropriate filename and path for the OS image.
- d. Select the **Xmodem** protocol (and not the 1k Xmodem protocol).
- e. Click on **[Send]**.

If you are using HyperTerminal, you will see a screen similar to the following to indicate that the download is in progress:

Figure 72: Example: of Xmodem download in progress



When the download completes, the switch reboots from primary flash using the OS image you downloaded in the preceding steps, plus the most recent startup-config file.

DNS resolver

The domain name system (DNS) resolver is designed for use in local network domains, where it enables the use of a host name or fully qualified domain name with DNS-compatible switch CLI commands.

DNS operation supports both IPv4 and IPv6 DNS resolution and multiple, prioritized DNS servers. (For information on IPv6 DNS resolution, see the latest IPv6 configuration guide for your switch.)

Basic operation

- When the switch is configured with only the IP address of a DNS server available to the switch, a DNS-compatible command, executed with a fully qualified domain name, can reach a device found in any domain accessible through the configured DNS server.
- When the switch is configured with both of the following:
 - The IP address of a DNS server available to the switch
 - The domain suffix of a domain available to the configured DNS server then:
 - A DNS-compatible command that includes the host name of a device in the same domain as the configured domain suffix can reach that device.
 - A DNS-compatible command that includes a fully qualified domain name can reach a device in any domain that is available to the configured DNS server.

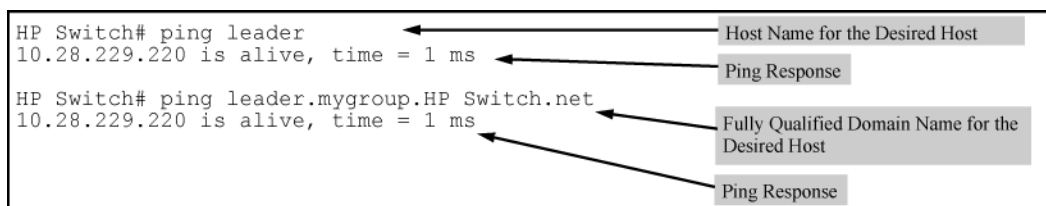
Example:

Suppose the switch is configured with the domain suffix `mygroup.HP Switch.net` and the IP address for an accessible DNS server. If an operator wants to use the switch to ping a target host in this domain by using the DNS name "leader" (assigned by a DNS server to an IP address used in that domain), the operator can use either of the following commands:

Figure 73: Example: of using either a host name or a fully qualified domain name

```
HP Switch# ping leader
10.28.229.220 is alive, time = 1 ms

HP Switch# ping leader.mygroup.HP Switch.net
10.28.229.220 is alive, time = 1 ms
```



The diagram shows two terminal sessions. The first session shows the command `ping leader` being executed, with an arrow pointing to `leader` labeled "Host Name for the Desired Host" and another arrow pointing to the output `10.28.229.220 is alive, time = 1 ms` labeled "Ping Response". The second session shows the command `ping leader.mygroup.HP Switch.net` being executed, with an arrow pointing to `leader.mygroup.HP Switch.net` labeled "Fully Qualified Domain Name for the Desired Host" and another arrow pointing to the output `10.28.229.220 is alive, time = 1 ms` labeled "Ping Response".

In the preceding Example:, if the DNS server's IP address is configured on the switch, but a domain suffix is either not configured or is configured for a different domain than the target host, the fully qualified domain name **must** be used.

Note that if the target host is in a domain **other than** the domain configured on the switch:

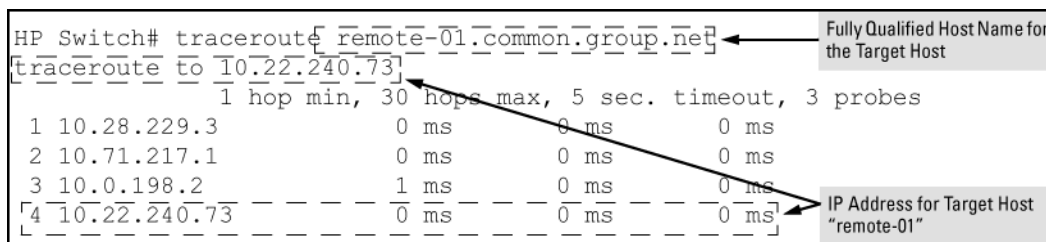
- The host's domain must be reachable from the switch. This requires that the DNS server for the switch must be able to communicate with the DNS servers in the path to the domain in which the target host operates.
- The fully qualified domain name must be used, and the domain suffix must correspond to the domain in which the target host operates, regardless of the domain suffix configured in the switch.

Example:

Suppose the switch is configured with the domain suffix `mygroup.HP Switch.net` and the IP address for an accessible DNS server in this same domain. This time, the operator wants to use the switch to trace the route to a host named "remote-01" in a different domain named `common.group.net`. Assuming this second domain is accessible to the DNS server already configured on the switch, a `traceroute` command using the target's fully qualified DNS name should succeed.

Figure 74: Example: using the fully qualified domain name for an accessible target in another domain

```
HP Switch# traceroute remote-01.common.group.net
[traceroute to 10.22.240.73]
1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.229.3          0 ms    0 ms    0 ms
 2 10.71.217.1         0 ms    0 ms    0 ms
 3 10.0.198.2          1 ms    0 ms    0 ms
[4 10.22.240.73       0 ms    0 ms    0 ms]
```



The diagram shows a terminal session where the command `traceroute remote-01.common.group.net` is executed. An arrow points to `remote-01.common.group.net` labeled "Fully Qualified Host Name for the Target Host". Another arrow points to the IP address `10.22.240.73` in the output, labeled "IP Address for Target Host 'remote-01'". The output shows a successful traceroute with four hops, all with 0 ms delay.

Configuring and using DNS resolution with DNS-compatible commands

The DNS-compatible commands include `ping` and `traceroute`.)

Procedure

1. Determine the following:
 - a. The IP address for a DNS server operating in a domain in your network.
 - b. The priority (1 to 3) of the selected server, relative to other DNS servers in the domain.
 - c. The domain name for an accessible domain in which there are hosts you want to reach with a DNS-compatible command. (This is the domain suffix in the fully qualified domain name for a given host

operating in the selected domain. See **Basic operation** on page 365.) Note that if a domain suffix is not configured, fully qualified domain names can be used to resolve DNS-compatible commands.

- d. The host names assigned to target IP addresses in the DNS server for the specified domain.
2. Use the data from the first three bullets in step1 to configure the DNS entry on the switch.
3. Use a DNS-compatible command with the host name to reach the target devices.

Configuring a DNS entry

The switch allows up to two DNS server entries (IP addresses for DNS servers). One domain suffix can also be configured to support resolution of DNS names in that domain by using a host name only. Including the domain suffix enables the use of DNS-compatible commands with a target's host name instead of the target's fully qualified domain name.

Syntax:

```
[no] ip dns server-address priority <1-3> <ip-addr>
```

Configures the access priority and IP address of a DNS server accessible to the switch. These settings specify:

- The relative priority of the DNS server when multiple servers are configured
- The IP address of the DNS server

These settings must be configured before a DNS-compatible command can be executed with host name criteria.

The switch supports two prioritized DNS server entries. Configuring another IP address for a priority that has already been assigned to an IP address is not allowed.

To replace one IP address at a given priority level with another address having the same priority, you must first use the `no` form of the command to remove the unwanted address. Also, only one instance of a given server address is allowed in the server list. Attempting to enter a duplicate of an existing entry at a different priority level is not allowed .

To change the priority of an existing server address, use the `no` form of the command to remove the entry, then re-enter the address with the new priority.

The `no` form of the command replaces the configured IP address with the null setting. (Default: null)

Syntax:

```
[no] ip dns domain-name <domain-name-suffix>
```

This optional DNS command configures the domain suffix that is automatically appended to the host name entered with a DNS-compatible command. When the domain suffix and the IP address for a DNS server that can access that domain are both configured on the switch, you can execute a DNS-compatible command using only the host name of the desired target. (For an Example:, see **Example: of using either a host name or a fully qualified domain name**.) In either of the following two instances, you must manually provide the domain identification by using a fully qualified DNS name with a DNS-compatible command:

- If the DNS server IP address is configured on the switch, but the domain suffix is not configured (null).
- The domain suffix configured on the switch is not the domain in which the target host exists.

The switch supports one domain suffix entry and three DNS server IP address entries. (See the preceding command description.)

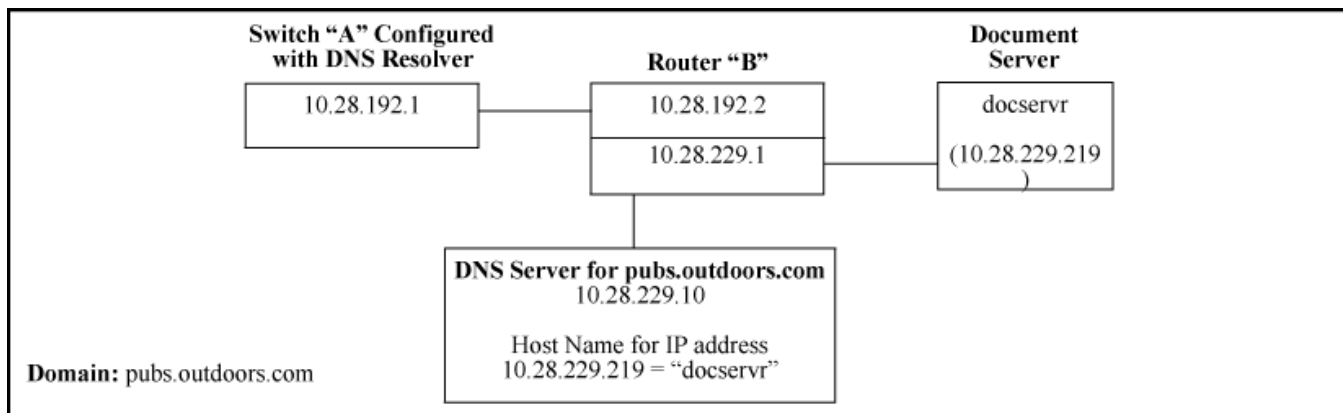
The `no` form of the command replaces the configured domain suffix with the null setting. (Default: null)

Using DNS names with ping and traceroute: Example:

In the network illustrated in **Figure 75: Example: network domain** on page 368, the switch at 10.28.192.1 is configured to use DNS names for DNS-compatible commands in the **pubs.outdoors.com** domain. The DNS

server has been configured to assign the host name **docservr** to the IP address used by the document server (10.28.229.219).

Figure 75: Example: network domain



Configuring switch "A" with the domain name and the IP address of a DNS server for the domain enables the switch to use host names assigned to IP addresses in the domain to perform `ping` and `traceroute` actions on the devices in the domain. To summarize:

Entity	Identity
DNS server IP address	10.28.229.10
Domain name (and domain suffix for hosts in the domain)	pubs.outdoors.com
Host name assigned to 10.28.229.219 by the DNS server	docservr
Fully qualified domain name for the IP address used by the document server (10.28.229.219)	docservr.pubs.outdoors.com
Switch IP address	10.28.192.1
Document server IP address	10.28.229.219

With the above already configured, the following commands enable a DNS-compatible command with the host name `docservr` to reach the document server at 10.28.229.219.

Configuring switch "A" in Example: network domain to support DNS resolution

```
switch(config)# ip dns server-address 10.28.229.10
switch(config)# ip dns domain-name pubs.outdoors.com
```

Ping and traceroute execution for the network in Example: network domain

```
switch(config)# ping docservr
10.28.229.219 is alive, time = 1 ms

switch# traceroute docservr
traceroute to 10.28.229.219
 1 hop min, 30 hops max, 5 sec. timeout, 3 probes
```



```

1 10.28.192.2 1 1 ms 0 ms 0 ms
2 10.28.229.219 2 0 ms 0 ms 0 ms

```

- ¹First-Hop Router (“B”)
- ²Traceroute Target

As mentioned under the following example, if the DNS entry configured in the switch does not include the domain suffix for the desired target, you must use the target host's fully qualified domain name with DNS-compatible commands. For example, using the document server in **Figure 75: Example: network domain** on page 368 as a target:


Figure 76: Example: of ping and traceroute execution when only the DNS server IP address is configured

```

HP Switch# ping [docsrvr.pubs.outdoors.com]
10.28.229.219 is alive, time = 1 ms

HP Switch# traceroute [docsrvr.pubs.outdoors.com]
traceroute to 10.28.229.219
          1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.192.2          1 ms          0 ms          0 ms
 2 10.28.229.219       0 ms          0 ms          0 ms

```



Viewing the current DNS configuration

The `show ip` command displays the current domain suffix and the IP address of the highest priority DNS server configured on the switch, along with other IP configuration information. If the switch configuration currently includes a non-default (non-null) DNS entry, it will also appear in the `show run` command output.

Figure 77: Example: of viewing the current DNS configuration

```

HP Switch# show ip

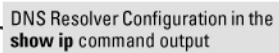
Internet (IP) Service

  IP Routing : Disabled

  Default Gateway : 10.28.192.2
  Default TTL     : 64
  Arp Age        : 20
  Domain Suffix  : pubs.outdoors.com
  DNS server     : 10.28.229.10

  VLAN          | IP Config | IP Address | Subnet Mask
  -----+-----
  DEFAULT_VLAN | Manual    | 10.28.192.1 | 255.255.255.0

```



Operating notes

- Configuring another IP address for a priority that has already been assigned to an IP address is not allowed. To replace one IP address at a given priority level with another address having the same priority, you must first use the `no` form of the command to remove the unwanted address. Also, only one instance of a given server address is allowed in the server list. Attempting to enter a duplicate of an existing entry at a different priority level is not allowed. To change the priority of an existing server address, use the `no` form of the command to remove the entry, then re-enter the address with the new priority.
- To change the position of an address already configured with priority `x`, you must first use `no ip dns server-address priority x <ip-addr>` to remove the address from the configuration, then use `ip dns server-address priority <ip-addr>` to reconfigure the address with the new priority. Also, if the

priority to which you want to move an address is already used in the configuration for another address, you must first use the `no` form of the command to remove the current address from the target priority.

- The DNS servers and domain configured on the switch must be accessible to the switch, but it is not necessary for any intermediate devices between the switch and the DNS server to be configured to support DNS operation.
- When multiple DNS servers are configured on the switch, they can reside in the same domain or different domains.
- A DNS configuration must include the IP address for a DNS server that is able to resolve host names for the desired domain. If a DNS server has limited knowledge of other domains, its ability to resolve DNS-compatible command requests is also limited.
- If the DNS configuration includes a DNS server IP address but does not also include a domain suffix, then any DNS-compatible commands should include the target host's fully qualified domain name.
- Switch-Initiated DNS packets go out through the VLAN having the best route to the DNS server, even if a Management VLAN has been configured.
- The DNS server address must be manually input. It is not automatically determined via DHCP.

Event Log messages

Please see the *Event Log Message Reference Guide* for information about Event Log messages.

Overview

The switch assigns MAC addresses in these areas:

- For management functions, one Base MAC address is assigned to the default VLAN (VID = 1). (All VLANs on the switches covered in this guide use the same MAC address.)
- For internal switch operations: One MAC address per port (see [Viewing the port and VLAN MAC addresses](#) on page 373).

MAC addresses are assigned at the factory. The switch automatically implements these addresses for VLANs and ports as they are added to the switch.



The switch's base MAC address is also printed on a label affixed to the switch.

Determining MAC addresses

Use the CLI to view the switch's port MAC addresses in hexadecimal format.

Use the menu interface to view the switch's base MAC address and the MAC address assigned to any VLAN you have configured on the switch. (The same MAC address is assigned to VLAN1 and all other VLANs configured on the switch.)



The switch's base MAC address is used for the default VLAN (VID =1) that is always available on the switch. This is true for dynamic VLANs as well; the base MAC address is the same across all VLANs.

Viewing the MAC addresses of connected devices

Syntax:

```
show mac-address [port-list | mac-addr | vlan < vid>]
```

Lists the MAC addresses of the devices the switch has detected, along with the number of the specific port on which each MAC address was detected.

<code>[port-list]</code>	Lists the MAC addresses of the devices the switch has detected, on the specified ports.
<code>[mac-addr]</code>	<p>Lists the port on which the switch detects the specified MAC address. Returns the following message if the specified MAC address is not detected on any port in the switch:</p> <pre>MAC address <mac-addr> not found.</pre>
<code>[vlan <vid>]</code>	Lists the MAC addresses of the devices the switch has detected on ports belonging to the specified VLAN, along with the number of the specific port on which each MAC address was detected.

Viewing the switch's MAC address assignments for VLANs configured on the switch

The Management Address Information screen lists the MAC addresses for:

- Base switch (default VLAN; VID=1)
- Any additional VLANs configured on the switch.

Also, the Base MAC address appears on a label on the back of the switch.




The Base MAC address is used by the first (default) VLAN in the switch. This is usually the VLAN named "DEFAULT_VLAN" unless the name has been changed (by using the VLAN Names screen). On the switches covered in this guide, the VID (VLAN identification number) for the default VLAN is always "1," **and cannot be changed**.

- From the Main Menu, select
 1. **Status and Counters**
 2. **Switch Management Address Information**

If the switch has only the default VLAN, the following screen appears. If the switch has multiple static VLANs, each is listed with its address data.

Figure 78: Example: of the Management Address Information screen



Viewing the port and VLAN MAC addresses

The MAC address assigned to each switch port is used internally by such features as Flow Control and the spanning-tree protocol. Using the `walkmib` command to determine the MAC address assignments for individual ports can sometimes be useful when diagnosing switch operation.



This procedure displays the MAC addresses for all ports and existing VLANs in the switch, regardless of which VLAN you select.

Procedure

1. If the switch is at the CLI Operator level, use the `enable` command to enter the Manager level of the CLI.
2. Enter the following command to display the MAC address for each port on the switch:

```
switch# walkmib ifPhysAddress
```

(The above command is not case-sensitive.)

Example:

A switch with the following module configuration shows MAC address assignments similar to those shown in the example below:

- A 4-port module in slot A, a 24-port module in slot C, and no modules in slots B and D
- Two non-default VLANs configured

Figure 79: Example: of Port MAC address assignments on a switch

<pre> HP Switch# walkmib ifphysaddress ifPhysAddress.1 = 00 12 79 88 b1 ff ifPhysAddress.2 = 00 12 79 88 b1 fe ifPhysAddress.3 = 00 12 79 88 b1 fd ifPhysAddress.4 = 00 12 79 88 b1 fc ifPhysAddress.49 = 00 12 79 88 b1 cf ifPhysAddress.50 = 00 12 79 88 b1 ce ifPhysAddress.51 = 00 12 79 88 b1 cd ifPhysAddress.52 = 00 12 79 88 b1 cc ifPhysAddress.53 = 00 12 79 88 b1 cb ifPhysAddress.54 = 00 12 79 88 b1 ca ifPhysAddress.55 = 00 12 79 88 b1 c9 ifPhysAddress.56 = 00 12 79 88 b1 c8 ifPhysAddress.57 = 00 12 79 88 b1 c7 ifPhysAddress.58 = 00 12 79 88 b1 c6 ifPhysAddress.59 = 00 12 79 88 b1 c5 ifPhysAddress.60 = 00 12 79 88 b1 c4 ifPhysAddress.61 = 00 12 79 88 b1 c3 ifPhysAddress.62 = 00 12 79 88 b1 c2 ifPhysAddress.63 = 00 12 79 88 b1 c1 ifPhysAddress.64 = 00 12 79 88 b1 c0 ifPhysAddress.65 = 00 12 79 88 b1 bf ifPhysAddress.66 = 00 12 79 88 b1 be ifPhysAddress.67 = 00 12 79 88 b1 bd ifPhysAddress.68 = 00 12 79 88 b1 bc ifPhysAddress.69 = 00 12 79 88 b1 bb ifPhysAddress.70 = 00 12 79 88 b1 ba ifPhysAddress.71 = 00 12 79 88 b1 b9 ifPhysAddress.72 = 00 12 79 88 b1 b8 ifPhysAddress.362 = 00 12 79 88 a1 00 ifPhysAddress.461 = 00 12 79 88 a1 00 ifPhysAddress.488 = 00 12 79 88 a1 00 ifPhysAddress.4456 = </pre>	<p>ifPhysAddress.1 - 4: Ports A1 - A4 in Slot A (Addresses 5 - 24 in slot A are unused.)</p> <p>ifPhysAddress.49 - 72: Ports C1 - C24 in Slot C (In this example, there is no module in slot B.)</p> <p>ifPhysAddress.362 Base MAC Address (MAC Address for default VLAN; VID = 1)</p> <p>ifPhysAddress.461 and 488 Physical addresses for non-default VLANs configured on the switch. On the switches covered by this manual, all VLANs use the same MAC address as the Default VLAN. Refer to "Multiple VLAN Considerations" in the "Static LANs (VLANs)" chapter of the <i>Advanced Traffic Management Guide</i> for your switch.</p> <p>Virtual</p>
--	---

Configuring the savepower LED option

The `savepower led` command provides the ability to turn off port LEDs even when a link exists. If power-saving is enabled, it can be temporarily overridden by the LED Mode button on the front panel. If the LED Mode button is pressed, the LEDs will behave normally (turn on) for a period of 10 minutes, and then turn off again.

Syntax:

```
[no] savepower led
```

Turns power-saving option on or off for the LEDs.

The savepower led command

```
switch(config)# savepower led
```

The `no` form of the `savepower led` command cancels power saving mode and the LEDs are returned to their original state.

To display the configured status of the LED power-saving option, use the `show savepower led` command.

Output of the show savepower led command

```
switch(config)# show savepower led
```

```
LED Save Power Information
```

```
Configuration Status : Enabled
```

Configuring the savepower port-low-pwr option

The `port-low-pwr` option puts all the ports on the switch into auto low power mode if they are not linked.

Syntax:

```
[no] savepower port-low-pwr
```

Puts ports in low power mode.

When a link is detected, the ports return to normal power mode.

The `no` form of the command puts the ports into normal power mode.

The savepower port-low-power command

```
switch(config)# savepower port-low-pwr
```

To display the status of the power-down feature, use the `show savepower portlow-pwr` command. The output shows if the feature is enabled or not enabled.

Output for the show savepower port-low-pwr command

```
switch(config)# show savepower port-low-pwr
```

```
Port Save Power Status: Enabled
```


Job Scheduler

The Job Scheduler feature enables the user to schedule commands or jobs on the switch for one time or multiple times. This is similar in concept to the UNIX 'cron' utility. The user can schedule any CLI command that the user would otherwise enter interactively. This includes commands to enable or disable ports, LEDs, and Power-Over-Ethernet. Jobs can also be scheduled to be triggered by certain pre-defined events such as switch reboot. The only major restriction on commands scheduled is that, it should not prompt/ask for any user inputs.

Commands

Job at | delay | enable | disable

Set schedule jobs using the options and set the count for the number of times the job is repeated.

Syntax

```
job JOB NAME at | delay | enable | disable
```

Description

Schedule a command to run automatically. Jobs can be scheduled to run once, multiple times on a recurring basis, or after certain events such as reboots. All commands run with manager privilege in configuration context.

The [no] form of the command deletes a scheduled job.

By default, jobs will be repeated an infinite number of times.

Restrictions

Jobs scheduled at any event will not be counted.

Jobs that are scheduled at the event "reboot" will not work in some multi management switches.

Range

- <1-1000>: is the value range for the `count` option.
- ([[DD:]HH:]MM): is the format used for the specific delay.

Options

count	Specify the number of times the job should run.
delay	Specify the delay before running the job.
enable	Enable a job that is disabled or expired.
disable	Disable a job. By default, a job is enabled.

Usage

```
job <JOB NAME> at <([DD:]HH:]MM on <WEEKDAY-LIST>)> config-save <COMMAND> count <1-1000>
```

```
job <JOB NAME> at <[HH:]MM on [MM/]DD> config-save <COMMAND> count <1-1000>
```

```
job <JOB NAME> at <EVENT> config-save <COMMAND>
```

```
job <JOB NAME> delay <([DD:]HH:]MM> config-save <COMMAND> count <1-1000>
```

```
job <JOB NAME> enable | disable
```

```
[no] job <JOB NAME>
```

Show job

Syntax

```
show job
```

Description

Show the jobs scheduled.

Show job

```
switch# show job
```

```
Job Scheduler Status and Configuration
```

```
Scheduler Status : Waiting for the system time to be set
```

Name	Event or Time	Repeat Count	Save Cfg	Command
Burrrrrrrrrrrrr...	reboot	--	Yes	chassislocate blink
baz	reboot	--	No	show time
foo	17:00 SxTWTxS	--	No	savepower led
a1	12:00	2	Yes	sh time
a2	Every 2:14:30 days	75	Yes	vlan 3
a3	Every 00:00:25 days	1	No	vlan 4



Caution

The scheduler does not run until the system time is set.

Show job <Name>

Syntax

```
show job JOB NAME
```

Description

Show the job by name.

Show job <JOB NAME>

```
switch# show job a1
```

```
Job Information
```

```
Job Name      : a1
Runs At       : 01:24
Config Save   : No
Repeat Count  : --
Job Status    : Enabled
Run Count     : 1
Error Count   : 0
Command       : show time
Job Status    : Enabled
```

```
Output from Last Run
```

```
-----
Tue Dec 15 01:24:00 2015
```

```
switch# show job a2
```

```
Job Information
```

```
Job Name      : a2
Runs At       : Every 2:14:30 days
Config Save   : Yes
Repeat Count  : 75
Run Count     : 0
Error Count   : 0
Command       : vlan 3
Job Status    : Disabled
```

```
switch# show job foo
```

```
Job Information
```

```
Job Name      : foo
Runs At       : 17:00 SxTWTxS
Config Save   : Yes
Repeat Count  : --
Run Count     : 0
Error Count   : 0
Command       : savepower led
Job Status    : Enabled
```

HPE's Virtual Technician is a set of tools aimed at aiding network switch administrators in diagnosing and caring for their networks. VT provides tools for switch diagnoses when faced with unforeseen issues.

To improve the Virtual Technician features of our devices, HPE has added the following tools:

- Cisco Discovery Protocol
- Enabling Debug tracing for MOCANA code
- User diagnostic crash via front panel security button
- User diagnostic crash via the serial console

Cisco Discovery Protocol (CDP)

Show cdp traffic

Syntax

```
show cdp traffic
```

Description

Displays the number of Cisco Discovery Protocol (CDP) packets transmitted, received and dropped.

CDP frame Statistics

Port No	Transmitted Frames	Received Frames	Discarded Frames	Error Frames
A1	46	26	6	7
A2	30	35	7	9
A3	120	420	670	670

Clear cdp counters

Syntax

```
clear cdp counters
```

Description

Allows a user to clear CDP statistics.

Clear cdp counters

Port No	Transmitted Frames	Received Frames	Discarded Frames	Error Frames
A1	46	26	6	7
A2	30	35	7	9
A3	120	420	670	670

Enable/Disable debug tracing for MOCANA code

Debug security

Syntax

```
debug security ssl
```

Description

Enables the debug tracing for MOCANA code.

Use the [no] parameter to disable debug tracing.

ssl Display all SSL messages.

User diagnostic crash via Front Panel Security (FPS) button

Allows the switch's front panel **Clear** button to manually initiate a diagnostic reset. In the case of an application hang, this feature allows you to perform reliable diagnostics by debugging via the front panel **Clear** button. Diagnostic reset is controlled via Front Panel Security (FPS) options.

Front panel security password-clear

From the configure context:

Syntax

```
[no] front-panel-security password-clear <RESET-ON-CLEAR> | factory-reset | password-recovery |  
diagnostic-reset <CLEAR-BUTTON | SERIAL-CONSOLE>
```

Description

Enable the ability to clear the password(s) and/or configuration via the front panel buttons.

[no] disables the password clear option.

Parameters

- If `password-clear` is disabled, the password(s) cannot be reset using the clear button on the front panel of the device.
- If `factory-reset` is disabled, the configuration/password(s) can not be reset using the clear and reset button combination at boot time.
- When `password-recovery` is enabled (and the front panel buttons disabled), a lost password can be recovered by contacting HPE customer support.
- When `password-recovery` is disabled, there is no way to access a device after losing a password with the front panel buttons disabled.
- If `diagnostic-reset` is disabled, the user cannot perform a diagnostic switch reset on those rare events where the switch becomes unresponsive to user input because of unknown reason(s).
- If `diagnostic-reset` is enabled, the user can perform a diagnostic hard reset which will capture valuable diagnostic data and reset the switch.

Options

factory-reset	Enable/Disable factory-reset ability.
password-clear	Enable/Disable password clear.
password-recovery	Enable/Disable password recovery.
diagnostic-reset	Enable/Disable diagnostic reset.

Front-panel-security diagnostic-reset

From the configure context:

Syntax

```
front-panel-security diagnostic-reset <CLEAR-BUTTON | SERIAL-CONSOLE>
```

Description

Enables the diagnostic reset so that the switch can capture diagnostic data.

- To initiate diagnostic reset via the clear button, press the clear button for at least 30 seconds but not more than 40 seconds.
- To initiate diagnostic switch reset via the serial console, enter the diagnostic reset sequence on the serial console.

Options

Clear button	Enables the diagnostics by choosing the clear button option.
Serial console	Enables the diagnostics by choosing the serial console option.

[no] front-panel-security diagnostic-reset

From the configure context:

Syntax

```
[no] front-panel-security diagnostic-reset
```

Description

Disables the diagnostic reset feature so that the user is prevented from capturing diagnostic data and performing a diagnostic reset on the switch. Both the sub-options `reset-via-serial-console` and `reset-via-clear-button` will be disabled. This is necessary if the switch becomes unresponsive (hangs) for unknown reasons.

No front-panel-security diagnostic-reset

```
no front-panel-security diagnostic-reset
```

```
Clear Password          - Enabled
Reset-on-clear         - Disabled
Factory Reset          - Enabled
Password Recovery      - Enabled
Diagnostic Reset       - Disabled
```



Disabling the diagnostic reset prevents the switch from capturing diagnostic data on those rare events where the switch becomes unresponsive to user input because of unknown reasons. Ensure that you are familiar with the front panel security options before proceeding.

Front-panel-security diagnostic-reset clear-button

From the configure context:

Syntax

```
front-panel-security diagnostic-reset clear-button
```

Description

This command will enable diagnostic-reset via clear button. The user will be allowed to perform diagnostic reset by depressing the clear button for 30 seconds and not more than 40 seconds.

Front-panel-security diagnostic-rest clear-button

```
front-panel-security diagnostic-rest clear-button
```

Diagnostic Reset	- Enabled
clear-button	- Enabled
serial-console	-Disabled



Disabling the diagnostic reset prevents the switch from capturing diagnostic data on those rare events where the switch becomes unresponsive to user input because of unknown reasons. Ensure that you are familiar with the front panel security options before proceeding.

[No] front-panel-security diagnostic-reset clear-button

From the configure context:

Syntax

```
[no] front-panel-security diagnostic-reset clear-button
```

Description

Disables the diagnostic-reset via clear button.



Disabling the diagnostic reset prevents the switch from capturing diagnostic data on those rare events where the switch becomes unresponsive to user input because of unknown reasons. Ensure that you are familiar with the front panel security options before proceeding.

Show front-panel-security

Syntax

```
show front-panel-security
```

Options

Show front-panel-security

Clear Password - Enabled
Reset -on-clear - Disabled
Factory Reset - Enabled
Password Recovery - Enabled
Diagnostic Reset - Enabled



By default, user initiated diagnostic reset is enabled.

Diagnostic table

Validation rules

Validation	Error
Extra 'token' passed after diagnostic-reset.	Invalid input: <token>.

FPS Error Log

Event	Message
RMON_BOOT_CRASH_RECORD1	Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset. On detection on local serial
RMON_BOOT_CRASH_RECORD1	SMM: Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset. On detection on SMM serial console and signaled to AMM
RMON_BOOT_CRASH_RECORD1	STKM: Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset. On detection on non-commander serial console and signaled to commander
RMON_BOOT_CRASH_RECORD1	User has initiated diagnostic reset via the serial console. Sw_panic() message
RMON_BOOT_CRASH_RECORD1	SMM: User has initiated diagnostic reset via the serial console. Sw_panic() message when triggered via SMM

Table Continued

Event	Message
RMON_BOOT_CRASH_RECORD1	STKM: User has initiated diagnostic reset via the serial console. Sw_panic() message when triggered via non-commander
Console print	STKM: HA Sync in progress; user initiated diagnostic request via the serial console rejected. Retry after sometime. Printed on the device console. When standby is in sync state, we don't want to crash the commander. So we report to the user to retry later
Console print	STKM: Member is booting; user initiated diagnostic request via the serial console rejected. Retry after sometime. Printed on the device console. When the member is till booting, it doesn't have the commander member number, thus we can't issue UIDC on the commander. So we report to the user to retry later.

User initiated diagnostic crash via the serial console

Remotely triggers a diagnostic reset of the switch via a serial console. This reset reboots the switch and collects diagnostic data for debugging an application hang, a system hang or any other rare occurrence. Diagnostic reset is controlled via FPS options.

The serial sequence to initiate the User Initiated Diagnostic Reset via Serial console is Ctrl+S, Ctrl+T, Ctrl+Q, Ctrl+T, Ctrl+S.

Front-panel-security diagnostic-reset serial-console

In the configure context:

Syntax

```
front-panel-security diagnostic-reset serial-console
```

Enables the diagnostic-reset via serial console. Allows the user to perform diagnostic reset by keying-in diagnostic reset sequence.

Front-panel-security diagnostic-reset serial-console

```
front-panel-security diagnostic-reset serial-console
```

```
Diagnostic Reset      - Enabled
clear-button         - Disabled
serial-console       - Enabled
```

[No] front-panel-security diagnostic-reset serial-console

In the configure context:

Syntax

```
[no] front-panel-security diagnostic-reset serial-console
```

Description

Disables the diagnostic-reset via serial console.

No front-panel-security diagnostic-reset serial-console

```
no front-panel-security diagnostic-reset serial-console
```

```
Diagnostic Reset      - Disabled
```



Disabling the diagnostic reset prevents the switch from capturing diagnostic data on those rare events where the switch becomes unresponsive to user input because of unknown reasons. Ensure that you are familiar with the front panel security options before proceeding.

Serial console error messages

Error	Message
RMON_BOOT_CRASH_RECORD1	Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset.
RMON_BOOT_CRASH_RECORD1	SMM: Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset.
RMON_BOOT_CRASH_RECORD1	STKM: Diagnostic reset sequence detected on serial console; user has initiated diagnostic reset.
RMON_BOOT_CRASH_RECORD1	User has initiated diagnostic reset via the serial console.
RMON_BOOT_CRASH_RECORD1	SMM: User has initiated diagnostic reset via the serial console.
RMON_BOOT_CRASH_RECORD1	STKM: User has initiated diagnostic reset via the serial console.
Console print	STKM: HA Sync in progress; user initiated diagnostic request via the serial console rejected. Retry after sometime.
Console print	STKM: Member is booting; user initiated diagnostic request via the serial console rejected. Retry after sometime.

Overview

Auto device detection

The command `device-profile` enables the user to define profiles and configure the associations of profiles to each device type. By creating a device profile, parameters will be defined for a connection interface by device type. To configure each parameter under a profile name, a context level is provided.

The command `device-profile name <PROFILE NAME>` configures for the default values. The default value is permissible when no user-defined profile is created.

To associate each device type with a device profile, a context level is created which authorizes the user to enable or disable the profile by device-type. Only the device type `aruba-ap` is supported.

Rogue AP isolation

The command `rogue-ap-isolation` configures each device and blocks, logs, or allows a rogue AP when detected. The command enables or disables rogue AP isolation.

The command `clear rogue-ap-isolation` is provided to clear the detected rogue AP device MAC address.

Show commands are provided to display the configuration and status of the profiles. Another show command will display the list of rogue APs detected.

Jumbo frames on a device port

Configure jumbo frame support for the device port. Jumbo frames are not permissible by default.

Enabling jumbo frame support in a profile might affect other ports with different profiles. When a profile has jumbo frame enabled and is applied to any port, all other ports that are members of any VLAN listed in the profile will also have jumbo frame support.

Configuration commands

allow-jumbo-frames

Syntax

```
allow-jumbo-frames
```

Description

Configure jumbo frame support for the device port. Jumbo frames are not enabled by default.

Enabling jumbo frame support in a profile affects other ports with different profiles. When a profile has jumbo frames enabled and is applied to any port, all other ports that are members of any VLAN listed in the profile will also have jumbo frame support.

Validation rules

Validation	Error/Warning/Prompt
Invalid jumbo command.	Invalid input.
If jumbo frame support is configured on a VLAN for which the device profile had overridden the configuration, display the existing warning.	This configuration change will be delayed because a device profile that enables jumbo frame support is applied to a port in this VLAN.

Default AP Profile

Creates a user-defined profile.

The profile name is a valid character string with the maximum permissible length of 32. The default profile is named `default-ap-profile` and cannot be modified.

The default configuration parameters may be modified using the command `device-<PROFILE NAME> default-ap-profile` . Up to four different profiles may be configured.

The `[no]` command removes the user-defined profiles.

device-profile

From within the configure context:

Syntax

```
device-profile <PROFILE-NAME> <DEVICE-TYPE>
```

Description

Create port configuration profiles and associate them with devices. When a configured device type is connected on a port, the system will automatically apply the corresponding port profile. When the device is disconnected, the profile is removed after a 2 minute delay. Connected devices are identified using LLDP.

Options

<PROFILE-NAME> Specify the name of the profile to be configured.

<DEVICE-TYPE> Specify an approved device-type to configure and attach a profile to.

Parameters

allow-jumbo-frames	Configure jumbo frame support for the device port.
untagged-vlan <VLAN-ID>	Configure this port as an untagged member of specified VLAN.
tagged-vlan <VLAN-LIST>	Configure this port as a tagged member of the specified VLANs.
cos <COS-VALUE>	Configure the Class of Service (CoS) priority for traffic from the device.
ingress-bandwidth <PERCENTAGE>	Configure ingress maximum bandwidth for the device port.
egress-bandwidth <PERCENTAGE>	Configure egress maximum bandwidth for the device port.
poe-max-power <WATTS>	Configure the maximum PoE power for the device port (in watts).

poe-priority

Configure the PoE priority for the device port.

Usage

[no] device-profile name <PROFILE-NAME>

[no] device-profile type <DEVICE>

Associating a device with a profile

To associate an Aruba access point (AP) device-type to a user-defined profile, use the context `HPE Switch (device-aruba-ap) #`. All Aruba access points use the identifier **aruba-ap**.

The `[no]` form of the command removes the device type association and disables the feature for the device type.

The feature is disabled by default.

device-profile type

From within the configure context:

Syntax

device-profile type

Description

Configure an approved device-type and attach the profile. The profile configuration is applied to any port where this device type is connected.

Approved device types

aruba-ap	Aruba access point device.
aruba-switch-router	Aruba switch or router device.
cisco-phone	Cisco phone device.
cisco-switch-router	Cisco switch or router device.
hpe-switch-router	HPE switch or router device.

Options

From within the **device-aruba-ap** context

associate <PROFILE-NAME>	Associated the specified device type by profile name.
enable	Enables the automatic profile association.
disable	Disables the automatic profile association.

Usage

[no] device-profile type <DEVICE> [associate <PROFILE-NAME> | enable | disable]

Configuring the rogue-ap-isolation command

Used to configure the `rogue-ap-isolation` command. A block/log option may be configured for when a rogue AP is identified by the switch. The block/log option may be enabled or disabled. The default action is to block a rogue AP.

The `whitelist` command is used to configure any specific MAC addresses excluded from the rogue AP list. The whitelist configuration is saved in the configuration. The whitelist supports 128 MACs.

The `[no]` form the command is used to remove the MAC address individually by specifying the MAC.

rogue-ap-isolation

Within the configure context:

Syntax

```
rogue-ap-isolation
```

Description

Configure rogue AP isolation and rogue AP Whitelist MAC addresses for the switch. When enabled, the system detects the MAC address of rogue access points and takes the specified action for traffic or from that address. The whitelist is used to add MAC addresses of approved access points to the whitelist.

Options

action	Configure the action to take for rogue AP packets. Actions available are enable, disable, block, log, and whitelist.
block	Block and logs traffic to or from any rogue access points.
log	Log traffic to or from any rogue access points.
enable	Enable the rogue AP Isolation.
disable	Disable the rogue AP Isolation.
whitelist <MAC-ADDRESS>	Configures rogue AP Whitelist MAC addresses for the switch. This option is used to add MAC addresses of approved access points to the whitelist.
<MAC-ADDR>	Specify the MAC address of the device to be moved from the Rogue AP list to the whitelist.

Usage

```
rogue-ap-isolation [enable | disable]
```

```
rogue-ap-isolation action [log | block]
```

```
[no] rogue-ap-isolation whitelist <MAC-ADDRESS>
```

Show commands

show device-profile

Syntax

Within the configure context:

```
show device-profile
```

Description

Show device profile configuration and status.

config Show the device profile configuration details for a single, or all, profiles.

status Show currently applied device profiles.

Usage

```
show device-profile config <PROFILE-NAME>
```

```
show device-profile status
```

show device-profile config

```
Switch# Show device-profile config
Device Profile Configuration
Configuration for device profile : default-ap-profile
  untagged-vlan    : 1
  tagged-vlan     : None
  ingress-bandwidth : 100%
  egress-bandwidth : 100%
  cos             : 0
  speed-duplex    : auto
  poe-max-power   : 33W
  poe-priority    : High
  allow-jumbo-frames: Enabled

Configuration for device profile : profile1
  untagged-vlan    : 10
  tagged-vlan     : 40,50,60
  ingress-bandwidth : 10%
  egress-bandwidth : 95%
  cos             : 4
  speed-duplex    : auto-10
  poe-max-power   : 20W
  poe-priority    : Low
```

show device-profile config profile1

```
Switch# Show device-profile config profile1
Device Profile Configuration
Configuration for device profile : profile1
  untagged-vlan    : 10
  tagged-vlan     : 40,50,60
  ingress-bandwidth : 10%
  egress-bandwidth : 95%
  cos             : 4
```

```
speed-duplex      : auto-10
poe-max-power     : 20W
poe-priority      : Low
```

show command device-profile status

Syntax

```
show device-profile [config | status]
```

Description

Displays the device-profile configuration or device-profile status.

Options

config Show device profile configuration details for a single profile or all profiles.

status Show currently applied device profiles status.

show device-profile status

```
Switch# show device-profile status

Device Profile Status
Port           Device Type      Applied Device Profile
-----
5              aruba-ap         profile1
10             aruba-ap         profile1
```

Show rogue-ap-isolation

Syntax

```
show rogue-ap-isolation
```

Description

Show rogue access point information.

Options

whitelist Show rogue access point whitelist information.

Usage

```
show rogue-ap-isolation whitelist
```

show rogue-ap-isolation

```
Switch# show rogue-ap-isolation

Rogue AP Isolation
Rogue AP Status : Enable
Rogue AP Action : Block
Rogue AP MAC           Neighbor Device
-----
```



```
11:22:33:44:55:66  
aa:bb:cc:dd:ee:ff
```

```
00:12:34:56:67:89  
00:98:45:56:67:89
```

show rogue-ap-isolation whitelist

```
Switch# show rogue-ap-isolation whitelist
```

```
Rogue AP Whitelist Configuration
```

```
Rogue AP MAC
```

```
-----
```

```
11:22:33:44:55:66
```

```
aa:bb:cc:dd:ee:ff
```

Overview

Every client is associated with a user role. User roles associate a set of attributes for authenticated clients (clients with authentication configuration) and unauthenticated clients, applied to each user session. User roles must be enabled globally.



Local user roles are only supported when running YA software.

Examples of user roles are:

- Employee = All access
- Contractor = Limited access to resources
- Guest = Browse Internet

Each user role determines the client network privileges, frequency of reauthentication, applicable bandwidth contracts, and other permissions. There are a maximum of 32 administratively configurable user roles available with one predefined and read-only user role called **denyall**.

A user role consists of optional parameters such as:

- Captive portal profile Specifies the URL via:
 - **captive-portal profile**
 - or
 - Vendor Specific Attribute (VSA). RADIUS: HP **HP-Captive-Portal-URL** = <http://...>
- Ingress user policy

L3 (IPv4 and/or IPv6) ordered list of Classes with actions, with an implicit deny all for IPv4 and IPv6.

- Reauthentication period

The time that the session is valid for. The default is 0 unless the user role is overridden. The default means that the reauthentication is disabled.



Reauthentication period is required to override the default of 0.

- Untagged VLAN (either VLAN ID or VLAN-name)

VLAN precedence order behavior:

- If configured, untagged VLAN specified in the user role (VSA Derived Role, UDR, or Initial Role).
- Statically configured untagged and/or tagged VLANs of the port the user is on.

Operational notes

- When user roles are enabled, all users that are connecting on ports where authentication is configured will have a user role applied. User role application happens even if the user fails to authenticate. If the user cannot be authenticated, the “Initial Role” will be applied to that user.
- The user role may be applied in one of two ways:
 - Vendor Specific Attribute (VSA)

Type: RADIUS: Hewlett-Packard-Enterprise

Name: HPE-User-Role

ID: 25

Value: <myUserRole>

The RADIUS server (ClearPass Policy Manager) determines application of the VSA Derived Role. The role is sent to the switch via a RADIUS VSA. The VSA Derived Role will have the same precedence order as the authentication type (802.1x, WMA).

- User Derived Role (UDR

)The User Derived Role is part of Local MAC authentication (LMA) and is applied when user roles are enabled and LMA is configured.

UDR will have the same precedence as LMA. Precedence behavior of the authentication types will be maintained, (802.1x -> LMA -> WMA (highest to lowest)).

Restrictions

- User roles cannot be enabled when BYOD redirect, MAC authentication failure redirect, or enhanced web-based authentication are enabled.
- Web-based authentication is not supported on the same port with other authentication methods when user roles are enabled.
- `show port-access <AUTH-TYPE>` commands are not supported when user-roles are enabled. The command `show port-access clients [detail]` is the only way to see authenticated clients with their associated roles.
- `aaa port-access auth <port> control` commands are not supported when user roles are enabled.
- `unauth-vid` commands are not supported when user roles are enabled.
- `auth-vid` commands are not supported when user roles are enabled.

Limitations for web-based authentication

Cannot be combined with other authentication types on same port.

Limitations for LMA

Reauthentication period and captive portal profile are not supported.

Error messages

Action	Error message
Attempting to enable BYOD Redirect when user roles are enabled.	BYOD redirect cannot be enabled when user roles are enabled.
Attempting to enable MAFR when user roles are enabled.	MAC authentication failure redirect cannot be enabled when user roles are enabled.
Attempting to enable enhanced web-based authentication when user roles are enabled.	Enhanced web-based authentication cannot be enabled when user roles are enabled.
Attempting to enable web-based authentication when other authentication types are enabled for the same port, and user roles are enabled.	Web-based authentication cannot be enabled with other authentication types on this port when user roles are enabled.
<code>switch (config)# show port-access mac-based clients</code>	User roles are enabled. Use <code>show port-access clients</code> to view client information.

Table Continued

Action	Error message
<pre>switch (config)# aaa port-access authenticator e8 control autho</pre> <p>Attempting to enable local user role when MAFR, BYOD, or EWA are enabled.</p>	<p>802.1x control mode, Force Authorized/Unauthorized, cannot be set when user roles are enabled.</p> <p>User roles cannot be enabled when BYOD redirect, MAC authentication failure redirect, or enhanced web-based authentication are enabled.</p>

Captive-portal commands

Overview

The Captive Portal profile defines the web address that a user is redirected to for Captive Portal authentication. If the url is blank, a RADIUS VSA will be used.



There is a predefined profile called **use-radius-vsa** that is already configured to use the RADIUS VSA.

Two captive portal profiles are supported:

- Predefined and read-only
 - Predefined and read-only profile name is `use-radius-vsa`.
- Customized

[no] aaa authentication captive-portal profile

Syntax

```
[no] aaa authentication captive-portal profile <PROFILE-STR> [url <URL-STR>]
```

Description

Create a captive-portal profile. Profiles are used in user roles to direct the user to a designated captive portal server. When the profile includes a web address, that web address is always used to contact the server. When no web address is specified, it is obtained from the RADIUS VSA.



A profile does not have to be pre-existing in the switch for it to be configured to a user role.

Options

profile	Configure a captive portal profile.
<PROFILE-STR>	Configure a captive portal profile string 64 characters long.
url	Configure the captive portal server web address.
<URL-STR>	Configure the captive portal server web address string.

Usage

Switch# aaa authentication captive-portal profile <NAME>

Switch# aaa authentication captive-portal profile <NAME> url <URL>

Validation rules

Validation	Error/Message/Prompt
Attempts made to remove a nonexisting profile will return an error: switch# no aaa authentication captive-portal profile NON_EXISTING_PROFILE	Captive portal profile NON_EXISTING_PROFILE not found.
When including the configured web address after the web address parameter: [no] aaa authentication captive-portal profile myCaptivePortalProfile url http://myCPPM.local/guest/captive_portal_login.php	Invalid input: http://blablabla.com
A profile name with invalid syntax produces an error: Switch# aaa authentication captive-portal-profile "this is an invalid name"	#aaa authentication captive-portal-profile "this is an invalid name" Invalid character ' ' in name.
When trying to modify a profile that is predefined, switch# aaa authentication captive-portal-profile name use-radius-vsa	Captive portal profile use-radius-vsa is read only and cannot be modified
A profile name that is too long produces an error: switch# aaa authentication captive-portal-profile test342...;ldklsdjflkdsjflk	The name must be fewer than 64 characters.
When attempting to configure more than the number of admin configured profiles,switch# aaa authentication captive-portal-profile profileNumber2	No more captive portal profiles may be created.

Policy commands

Overview

These commands create a context that may be used to classify the policy. From the existing `policy` command, a new policy type called **user** was added. The new actions are specific to **policy user**:

- redirect
- permit
- deny



Only L3 classes (IPv4 and IPv6) are currently supported.
The user policy includes “implicit deny all rules” for both IPv4 and IPv6 traffic.

policy user

Syntax

```
policy user <POLICY-NAME>
```

Description

Create and enter newly created user policy context.

Usage

```
Switch (config)# policy user employee
```

[no] policy user

Syntax

```
[no] policy user <POLICYNAME>
```

Description

Delete and remove specified user policy from switch configuration.

Operating notes

- The user policy will include implicit “deny all” rules for both IPv4 and IPv6 traffic.
- `ipv4` or `ipv6` classes must specify source address as *any*. Specifying host addresses or subnets will result in the following error message:

```
Switch (policy-user)# class ipv4 class25 action priority 0
User policies cannot use classes that have a source IP address specified.
```

- *permit* and *deny* are mutually exclusive.
- *ip-precedence* and *dscp* are mutually exclusive.

Usage

```
switch (config)# no policy user employee
```

policy resequence

Syntax

```
policy resequence <POLICYNAME> <START><INCREMENT>
```

Description

Resequence classes and remarks configured within specified user policy. The usage shows resequencing classes and remarks within user policy “employee” starting at 200 and incrementing by 2.

Usage

```
Switch (config)# policy user employee 200 2
```

Commands in the policy-user context

Create classes inside of the **policy** context before you apply actions to them.

(policy-user)# class

Within the **policy-user** context:

Syntax

```
(policy-user)# [no] [<SEQUENCE-NUMBER>] class ipv4 | ipv6 <CLASS-NAME> [action permit | deny |  
redirect captive portal] | [action dscp | ip—precedence <CODEPOINT | PRECEDENCE>] [action priority  
<PRIORITY>] | [action rate-limit kbps <RATE>]
```

Description

Associate a class with ACL or QoS actions for this policy.

Options

Options

deny	Deny all traffic.
DSCP	Specify an IP DSCP.
IP-precedence	Specify the IP precedence.
permit	Permit all traffic.
priority	Specify the priority.
rate-limit	Configure rate limiting for all traffic.
redirect	Specify a redirect destination.

Usage

```
Switch(policy-user)# class ipv6 employeelpv6Http action deny
```

```
Switch(policy-user)# class ipv4 http action redirect captive-portal
```

```
Switch(policy-user)# class ipv4 dnsDhcp action permit
```

User role configuration

aaa authorization user-role

Syntax

```
aaa authorization user-role [enable | disable] [initial-role <ROLE-STR>] [[name <ROLE>]]
```

Description

Configure user roles. A user role determines the client network privileges, the frequency of reauthentication, applicable bandwidth contracts, along with other permissions. Every client is associated with a user role or the client is blocked from access to the network.

Options

enable Enable authorization using user roles.

disable Disable authorization using user roles.

initial-role The default initial role “denyall” is used when no other role applies. If a client connects to the switch and does not have a user role associated, then the initial role is used. Any role can be configured as initial role using this option.

The initial role may be assigned if:

- captive-portal profile is configured with a web address, but the Captive Portal VSA is sent from RADIUS
- captive-portal profile is configured to use the RADIUS VSA but no Captive Portal VSA is sent.
- captive-portal feature is disabled when the captive-portal profile is referenced in the applied user role to the client.
- The user role feature is enabled with RADIUS authentication, but no user role VSA is returned.
- User role does not exist.
- Not enough TCAM resource available.
- Access-Reject from RADIUS.
- User role VSA is sent along with invalid attributes.
- RADIUS not reachable.
- VLAN configured on the user role does not exist.
- Captive Portal profile does not exist.
- User policy configured on the user role does not exist.
- Reauthentication period is enabled (nonzero) in the user role for LMA.
- Captive Portal profile is included in the user role for LMA.

name
<NAME-STR> Create or modify a user-role. Role name identifies a user-role. When adding a user-role, a new context will be created. The context prompt will be named “user-role” (user-role)#.

Usage

```
Switch# aaa authorization user-role enable
```

```
Switch# aaa authorization user-role disable
```

```
Switch# aaa authorization user-role name <ROLE1>
```

```
Switch# [no] aaa authorization user-role enable
```

```
Switch# [no] aaa authorization user-role name <ROLE1>
```

```
Switch# aaa authorization user-role initial-role <ROLE1>
```

```
Switch# aaa authorization user-role name <MYUSERROLE> policy <MYUSERPOLICY>
```

```
Switch# aaa authorization user-role name <MYUSERROLE> captive-portal-profile  
<MYCAPTPORTPROFILE>
```

```
Switch# aaa authorization user-role name <MYUSERROLE> vlan-id <VID>
```

```
Switch# aaa authorization user-role name <MYUSERROLE> reauth-period <0-999999999>
```

Error log

Scenario	Error Message
If the user tries to delete a user-role configured as the initial role	User role <INITIAL_ROLE_NAME> is configured as the initial role and cannot be deleted.
If the user attempts to configure more than the number of administrator configured roles	#aaa authorization user-role name roleNumber33. No more user roles can be created.
If the user enters a role name that is too long	switch# aaa authorization user-role test342....jflkdsjflk. The name must be fewer than 64 characters long.
If the user enters a role name with invalid syntax	switch# aaa authorization user-role name "this is an invalid name". Invalid character '' in name.
If the user tries to delete a nonexistent user-role	User role <NON_EXISTING_ROLE_NAME> not found.
Switch# aaa authorization user-role name <DENYALL>	User role <DENYALL> is read only and cannot be modified.

captive-portal-profile

From within the **user-role** context:

Syntax

```
captive-portal-profile <PROFILE_NAME>
```

Description

Assigns a captive portal profile to the user role. The predefined captive portal profile, `use-radius-vsa`, indicates that the redirect web address must be sent via RADIUS.

To clear a captive portal profile from the user role, use the [no] version of the command.

policy

From within the **user-role** context:

Syntax

```
policy <POLICY_NAME>
```

Description

Assigns a user policy to the user role. To clear a policy from the user role, use the [no] version of the command.



Modification of the user policy, or class contained in a user policy, will force users consuming that user policy via a user role to be deauthenticated.

reauth-period

From within the user-role context:

Syntax

```
reauth-period <VALUE>
```

Description

Set the reauthentication period for the user role. Use [0] to disable reauthentication. For RADIUS-based authentication methods, it will override the RADIUS session timeout. It also overrides any port-based reauth-period configuration with the exception that LMA does not support a reauth-period.

Options

<VALUE> Valid values are 0 – 999,999,999; a required configuration in user roles and it defaults to 0.

(user-role)# reauth-period 100

Set the reauthentication value for the current user role:

```
(user-role)# reauth-period 100
```

(user-role)# reauth-period 0

0 is used to disable reauthentication, and it is the default value.

```
(user-role)# reauth-period 0
```

Validation rules

Validation	Error/Warning/Prompt
(user-role)# reauth-period 10000000	Invalid input: 10000000000000000000

VLAN commands



The VLAN must be configured on the switch at the time the user role is applied. Only one of VLAN-name or VLAN-ID is allowed for any user role.

Modification of the VLAN will force users assigned to that VLAN via a user role to be deauthenticated.

vlan-id

From within the user-role context:

Subcommand syntax

```
vlan-id <VLAN-ID>
```

Description

Create a VLAN with id VLAN-ID.

Use the [no] version of the command when clearing the VLAN-ID from the user role:

Usage

```
(user-role)# no vlan-id
```

vlan-name

From within the **user-role** context:

Subcommand syntax

```
vlan-name <VLAN-NAME>
```

Description

Create a VLAN with the name VLAN-NAME. Only one of VLAN-NAME or VLAN-ID is allowed for any user role.

Use the [no] version of the command when clearing the VLAN from the user role, by name:

Usage

```
(user-role)# no vlan-name
```

vlan-id 100

```
(user-role)# vlan-id 100
```

vlan-name vlan100

```
(user-role)#vlan-name VLAN100
```

VLAN range commands

This command is executed from a global configuration context.

VLANs specified by VLAN-ID-LIST

Syntax

```
[no] vlan <VLAN-ID-LIST>
```

Description

Creates VLANs specified by the VLAN-ID-LIST and returns to the global configuration context. Use the [no] version of the command to delete the VLANs specified by the VLAN-ID-LIST.

Examples

```
config# vlan 2-15
config# vlan 5,10,13-20,25
config# no vlan 2-10
config# no vlan 2,5,15-18,25
```

VLANs specified by VLAN-ID-LIST and tag specified ports specified by PORT-LIST

Syntax

```
[no] vlan <VLAN-ID-LIST> tagged <PORT-LIST>
```

Description

Creates VLANs specified by the VLAN-ID-LIST and tags the ports specified by the PORT-LIST to the VLAN-ID-LIST. If VLANs already exist, the tagging of ports specified by the PORT-LIST is performed.

Use the [no] version of the command to remove the tagged PORT-LIST from a range of VLANs specified by the VLAN-ID-LIST. After command execution, CLI returns to the global configuration context.

Examples

```
config# vlan 2-15 tagged A1-A20
config# vlan 5,10,13-20,25 tagged A1-A5,L2,L5-L10
config# vlan 2-20 tagged all
config# no vlan 2-15 tagged A1-A5
config# no vlan 5,10,13-20 tagged A1-A5,L6
```

Applying a UDR

UDR can be used to assign user roles locally (that is, without RADIUS). LMA has been extended to allow applying a user role to a MAC address, MAC group, MAC mask, or MAC OUI.

aaa port-access local-mac apply user-role

Syntax

```
[no] aaa port-access local-mac apply user-role <Role-Name> [ mac-oui <MAC-OUI> | mac-mask <MAC-MASK> | mac-addr <MAC-ADDR> | mac-group <MAC-GROUP-NAME>]
```

Description

Apply user roles.

Options

mac-addr	To apply user role with MAC address.
mac-group	To apply user role with MAC group.
mac-mask	To apply user role with MAC Mask.
mac-oui	To apply user role with MAC OUI.

Usage

```
[no] aaa port-access local-mac apply user-role <MYUSERROLE> [mac-oui <MAC-OUI>]
```

```
[no] aaa port-access local-mac apply user-role <MYUSERROLE> [mac-mask <MAC-MASK>]
```

```
[no] aaa port-access local-mac apply user-role <MYUSERROLE> [mac-addr <MAC-ADDR>]
```

```
[no] aaa port-access local-mac apply user-role <MYUSERROLE> [mac-group <MAC-GROUP-NAME>]
```

Show commands

show captive-portal profile

Syntax

```
show captive-portal profile
```

Description

Show Captive Portal profile configuration.

show captive-portal profile

```
(config)# show captive-portal profile
```

```
Captive Portal Profile Configuration
Name : use-radius-vsa
Type : predefined
URL  :

Name : myCaptivePortalProfile
Type : custom
URL  : http://mycppm.local/guest/captive_portal_login.php
```

show user-role

Syntax

```
show user-role [<ROLE-NAME>] [detailed]
```

Description

Show users role configuration.

Options

<ROLE-NAME>	Show user roles by role-name.
<ROLE-NAME> detailed	Show user roles in detail by role-name.

show user-role

```
Switch# show user-role

User Roles

Enabled      : <Yes/No>
Initial Role : denyall

Type         Name
-----
local        Employee
local        Guest
predefined   denyall
```

show user-role <ROLE-NAME>

```
Switch# show user-role captivePortalwithVSA

User Role Information

Name          : captivePortalwithVSA
Type          : local
Reauthentication Period (seconds) : 0
Untagged VLAN : 610
Captive Portal Profile : use-radius-vsa
Policy        : cppolicy
```

show user-role detailed

The example shows how to configure user roles to use Clearpass as a Captive Portal. The Captive Portal URL is specified in a RADIUS VSA.

```
Switch# show user-role captivePortalwithVSA detailed

User Role Information
  Name                : captivePortalwithVSA
  Type                : local
  Reauthentication Period (seconds) : 0
  VLAN                : 610
  Captive Portal Profile : use-radius-vsa
  URL                 : (use RADIUS VSA)
  Policy              : cppolicy

Statements for policy "cppolicy"
policy user "cppolicy"
  10 class ipv4 "cppm" action permit
  20 class ipv4 "steal" action redirect captive-portal
  30 class ipv4 "other" action permit
  exit

Statements for class IPv4 "cppm"
class ipv4 "cppm"
  10 match tcp 0.0.0.0 255.255.255.255 1.0.9.15 0.0.0.0 eq 80
  20 match tcp 0.0.0.0 255.255.255.255 1.0.9.15 0.0.0.0 eq 443
  exit

Statements for class IPv4 "steal"
class ipv4 "steal"
  10 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 80
  20 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 443
  exit

Statements for class IPv4 "other"
class ipv4 "other"
  10 match udp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 53
  20 match udp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 67
  30 match icmp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
  exit
```

show port-access clients

Syntax

```
show port-access clients [detailed]
```

Description

Use this command to display the status of active authentication sessions.

show port-access clients

Port Access Client Status

Port	Client Name	MAC Address	IP Address	User Role	Type	VLAN
1/A18	001517581ec4	001517-581ec4	10.108.1.201	ixia1	MAC	108

A7	000c29-5121fc	n/a	denyall	LOCAL
A8	000c29-d12996	n/a	myrole	LOCAL 42

show port-access clients detailed

```
Switch (config)# show port-access clients detailed
```

Port Access Client Status Detail

Client Base Details :

Port	: 1/A18	Authentication Type	: mac-based
Client Status	: authenticated	Session Time	: 11 seconds
Client Name	: 001517581ec4	Session Timeout	: 60 seconds
MAC Address	: 001517-581ec4		
IP	: 10.108.1.201		

User Role Information

Name	: ixial
Type	: local
Reauthentication Period (seconds)	: 60
Untagged VLAN	: 108
Tagged VLANs	:
Captive Portal Profile	:
Policy	: policyIxial

Statements for policy "policyIxial"

```
policy user "policyIxial"  
  10 class ipv4 "classIxial" action rate-limit kbps 11000  
  exit
```

Statements for class IPv4 "classIxial"

```
class ipv4 "classIxial"  
  10 match ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255  
  exit
```


Overview

The Port QoS Trust feature restricts which packet QoS information may be used to determine inbound queue servicing and any priority information to be permitted into the local hop.

Port QoS Trust Mode configuration allows preservation or removal of the inbound QoS priorities carried in Layer 2 (the VLAN cos or Priority CodePoint (PCP) value, known as the 802.1p priority tag) and/or in Layer 3 (the IP-ToS byte, in IP-Precedence or IP-Diffserv mode). The different modes let the customer trust all, some, or no packet priority fields.

The per-port configuration enables the customer to trust some sources or devices and not others. This feature is mutually exclusive with any active port-priority configuration.

Configuration commands

qos trust

Syntax

```
qos trust [default|dot1p|dscp|ip-prec|none|device [none|<DEVICE-TYPE>]]
```

Description

Set the QoS Trust Mode configuration for the port.

Options

default	Trust 802.1p priority and preserve DSCP or IP-ToS.
device <DEVICE-TYPE>	On approved devices, trust IP-ToS Differentiated-Services in IP packets, and use the DSCP-MAP to remark the 802.1p priority. If the DSCP codepoint does not have an associated priority, the priority will be remarked to 0. On unapproved devices, trust 802.1p priority and preserve any IP- ToS values.
dot1p	Trust 802.1p priority and preserve DSCP or IP-ToS.
dscp	Trust IP-ToS Differentiated-Services in IP packets, and use the DSCP-MAP to remark the 802.1p priority. If the DSCP codepoint does not have an associated 802.1p priority, the priority will be remarked to 0.
ip-precedence	Trust IP-ToS IP-Precedence mode in IP packets and remark the 802.1p priority.
none	Do not trust either the 802.1p priority or the IP-ToS values.

QoS trust devices

aruba-ap	Aruba Access point device.
none	Clear all trusted devices from port.



Both SNMP and the CLI will verify that the current QoS Port Priority and desired QoS Trust Mode configuration are not mutually exclusive (and conversely).

qos dscp-map

Syntax

```
qos dscp-map <CODEPOINT> priority <PRIORITY> [name <NAME> | default | legacy]
```

Description

Modifies DSCP mapping.

Options

default Returns switch to the fully mapped factory-default configuration.

legacy Restore the legacy default behavior (partial mapping) used in earlier code releases.

Show commands

show qos trust

Syntax

```
show qos trust [device] <PORT>
```

Description

Shows port-based QoS trust configuration

Options

device Show list of trusted devices per-port.

<port> Show trusted devices on a single port.

Usage

```
show qos trust [device | [ethernet <PORT-LIST> ]
```

show qos trust

```
switch# show qos trust
```

```
Port-based qos Trust Configuration
```

Port	Trust Mode	Device Trust State	-----	---	----
A1	Default				
A2	Default				
A3	Device**		Trusted		
A4	IP-Prec				

```

A5      Dot1p      |
A5      None       |
A5      DSCP       |
A5      Device**  |
A5      Dot1p     |

```

** For a list of trusted devices per-port, use the command `show qos trust device`. To show trusted devices on a single port, use the command `show qos trust device <PORT>`.

show qos trust device

```

switch# show qos trust device

Port-Based QoS Trust Configuration

  Port      Trusted Devices
  -----
  A1        aruba-ap
  A2        aruba-ap
  A4        aruba-ap

```

show qos trust device <PORT>

```

switch# show qos trust device <PORT>

Port A4 QoS Trust Configuration
  Current state: Trusted

  Trusted Devices: aruba-ap

```

Validation rules

Validation	Error/Warning/Prompt
<pre>qos trust <UNSUPPORTEDDEVICETYPE> E></pre>	Invalid input: %s
<pre>no qos trust <ANYVALUE></pre>	Invalid command. To disable trust for a port, use <code>qos trust none</code> . To return to the default configuration and leave priority information unchanged, use <code>qos trust default</code> .
QoS priority when trust mode is anything other than <code><NONE></code> or <code><DEFAULT></code> .	The port QoS trust mode must be <code><DEFAULT></code> or <code><NONE></code> to configure the QoS port priority feature.
QoS DSCP when trust mode is anything other than <code><NONE></code> or <code><DEFAULT></code> .	The port QoS trust mode must be <code><DEFAULT></code> or <code><NONE></code> to configure the QoS port priority feature.

Table Continued

Validation	Error/Warning/Prompt
QoS trust dot1.p when any port QoS priority is enabled.	The port QoS priority feature must be disabled before configuring this port QoS trust mode.
QoS trust ip-prec when any port QoS priority is enabled.	The port QoS priority feature must be disabled before configuring this port QoS trust mode.
QoS trust DSCP when any port QoS priority is enabled.	The port QoS priority feature must be disabled before configuring this port QoS trust mode.
QoS trust device when any port QoS priority is enabled.	The port QoS priority feature must be disabled before configuring this port QoS trust mode.

Net-service Overview

Net-service names are used as alias in defining ACL rules for defined lists. An alias of net-service will configure a list of hosts, networks, or subnets.

Extended ACL can have both source IP, destination IP and port number along with protocol in its ACE. An alias-based ACE for an extended ACL therefore allows the use of an alias of net-service protocol and destination port.

Limitations

- Alias-based ACE will not support access-control based on source port which is a limitation of the net-service command. The use of net-service will also restrict operators specified for port number to `equals` and `range`.
 - Operators `lt`, `gt`, `equal`, `negative`, and `range` for the source port in the ACL rule are not specified using the options available in net-service.
 - Operators `lt`, `gt`, `negative` are not specified for destination port using the options available in net-service.
 - Only the ACL will be affected when changes are made to an existing net-service. Either the rule must be reapplied to the ACL or the switch must be rebooted to affect the service.

net-service [tcp | udp | port]

Syntax

```
[no] net-service <NAME-STR> [tcp | udp | <PROTOCOL>]  
port <PORT-LIST>
```

Description

Configures net-service.

Parameters

protocol IP protocol number.
Range: 0-255

TCP Configure an alias for a TCP protocol.

UDP Configure an alias for a UDP protocol.

port Specify a single port or a list of noncontiguous port numbers, by entering up to six port numbers, separated by commas or range of ports.
Range: 0-65535

Example net-service tcp-service tcp 100

```
net-destination src-ip  
host 10.120.0.1  
host 10.91.1.1  
host 10.0.100.12
```

```

net-destination destn-ip
  host 16.90.51.12
  host 10.93.24.1

net-service tcp-service tcp 100
ip access-list extended "acl1"
  permit alias src-ip alias destn-ip alias tcp-service

```

Net-destination overview

The use of net-destination and net-service helps reduce effort required to configure ACL rules.

Net-destination is a list of hosts, networks, or subnets that are used to configure an ACL rules.

There are two types of ACLs supported and configured on the switch:

- Standard
- Extended

Standard

The standard ACL can have an IP source or network in the ACE. Defining the alias-based ACE for standard ACL, only use an alias of net-destination for the source.

Example - standard

```

net-destination src-ip
  host 10.120.0.1
  host 10.91.1.1
  host 10.0.100.12

net-destination destn-ip
  host 16.90.51.12
  host 10.93.24.1

net-service tcp-service tcp 100
ip access-list extended "acl1"
  permit alias src-ip alias destn-ip
    alias tcp-service

```

Extended

The extended ACL can have both source IP, destination IP and port number along with protocol in its ACE. Defining an alias-based ACE for an extended ACL can use an alias of net-destination for the source and destination and an alias of net-service for the protocol and destination port. Alias-based ACE will not support access-control based on source port which is a limitation of the net-service command. The use of net-service will also restrict the operators that can be specified for port number to `equals` and `range`.

Example - extended

```

HP-Switch-5406Rz12(config)# ip access-list extended aext1
HP-Switch-5406Rz12(config-ext-nacl)#
  permit tcp host 10.100.12.1 gt 23 16.90.0.0 /16 range 200 400
HP-Switch-5406Rz12(config-ext-nacl)# exit

```

Limitations

- Limited to IPv4 addresses per syntax.
- Any changes made to an existing net-destination that is used by an ACL, will be applied on the ACL only when the rule is reapplied to it or when switch is rebooted.

- The number of entries for a single net-destination is limited. The number of net-destinations configurable on a switch is also limited.
- A considerable amount of memory (for global structures) will be allocated when alias-based ACEs are configured which may cause issues on a switch with low memory.
- The Host or Domain name cannot be specified as an entry in a net-destination.
- Application level gateway will not be supported as the existing ACL infra does not support ALG.
- SNMP support to configure and delete net-destination, net-service, and the alias-based rules will not be provided.
- The 'invert' and 'range' option have been deprecated as per ArubaOS-Switch 7.4 CLI Reference Guide and hence will not be supported. However, the functionality of 'invert' option can be achieved through the 'deny' rule.
- RADIUS server-based ACL application to interface/VLAN will not be supported for ACLs with alias-based rules.

net-destination host |position | network

Syntax

```
[no] netdestination <NAME-STR> [host <IP-ADDR>
[position <NUM>] network <IP-ADDR/MASK-LENGTH>
[position <NUM>]]
```

Description

Net-destination is a list of hosts, networks, or subnets that are used to configure an ACL rule.

Parameters

- host** Configures a single IPv4 host.
- network** An IPv4 subnet consisting of an IP address and netmask.
- no** Removes any configured item in list or an entire net-destination.
- position** Specifies the position of a host, network, or range in the net-destination. This optional parameter is specific to a net-destination and may only be used to sort entries in a list.

show net-destination

Syntax

```
show net-destination <NAME-STR>
```

Description

Show a host-specific net-destination.

Networking Websites

Hewlett Packard Enterprise Networking Information Library	www.hpe.com/networking/resourcefinder
Hewlett Packard Enterprise Networking Software	www.hpe.com/networking/software
Hewlett Packard Enterprise Networking website	www.hpe.com/info/networking
Hewlett Packard Enterprise My Networking website	www.hpe.com/networking/support
Hewlett Packard Enterprise My Networking Portal	www.hpe.com/networking/mynetworking
Hewlett Packard Enterprise Networking Warranty	www.hpe.com/networking/warranty

General websites

Hewlett Packard Enterprise Information Library	www.hpe.com/info/EIL
--	--

For additional websites, see [Support and other resources](#).

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:
<http://www.hpe.com/assistance>
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:
<http://www.hpe.com/support/hpesc>

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

Hewlett Packard Enterprise Support Center

www.hpe.com/support/hpesc

Hewlett Packard Enterprise Support Center: Software downloads

www.hpe.com/support/downloads

Software Depot

www.hpe.com/support/softwaredepot

- To subscribe to eNewsletters and alerts:

www.hpe.com/support/e-updates

- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

www.hpe.com/support/AccessToSupportMaterials



Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

<http://www.hpe.com/support/selfrepair>

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information

HPE Get Connected	www.hpe.com/services/getconnected
HPE Proactive Care services	www.hpe.com/services/proactivecare
HPE Proactive Care service: Supported products list	www.hpe.com/services/proactivecaresupportedproducts
HPE Proactive Care advanced service: Supported products list	www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central	www.hpe.com/services/proactivecarecentral
Proactive Care service activation	www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty for your product, see the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* document, available at the Hewlett Packard Enterprise Support Center:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional warranty information

HPE ProLiant and x86 Servers and Options	www.hpe.com/support/ProLiantServers-Warranties
HPE Enterprise Servers	www.hpe.com/support/EnterpriseServers-Warranties
HPE Storage Products	www.hpe.com/support/Storage-Warranties
HPE Networking Products	www.hpe.com/support/Networking-Warranties

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

www.hpe.com/info/reach

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

www.hpe.com/info/ecodata

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

www.hpe.com/info/environment

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**docsfeedback@hpe.com**). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.

Introduction


TR-069 is a technical specification created by the **Broadband Forum**. The TR-069 protocol specifies client and server requirements to manage devices across the Internet by using a client server architecture to provide communication between the CPE (Customer Premises Equipment) and the ACS (Auto Configuration Server). A protocol helps to manage complex networks where many devices such as modems, routers, gateways, VoIP phones and mobile tablets compete for resources. TR-069 defines the CPE WAN Management Protocol (CWMP) protocol necessary to remotely manage end-user devices. ACS provides automatic configuration for these devices.



CWMP is automatically enabled. To conserve resources, reconfigure this setting using the `cwmp disable` command.

TR-069 defines an auto-configuration architecture which provides the following primary capabilities:

- Auto-configuration and dynamic service provisioning
- Software/firmware image management
- Status and performance monitoring
- Diagnostics
- Bidirectional SOAP/HTTP based protocol



Advantages of TR-069

- TR-069 can manage devices with dynamic IP addresses.
 - TR-069 use Organization Unique ID (OUI) and serial number rather than IP to identify a device.
- TR-069 can manage devices in a private network.
 - The HPE ACS BIMS (an iMC module) uses HTTP to communicate with the device, and the session is initiated by the device, so BIMS can pass through NAT to manage the device.
- TR-069 is secure.
 - TR-069 can use HTTPS to communicate with or transfer files to/from the device; it is more secure than TFTP, FTP or Telnet.
- TR-069 is suitable for WAN management across internet.
- TR-069 is suitable for zero-touch configuration.
 - The zero-configuration mechanism is defined in the TR-069 specification.
- TR-069 is suitable for large-scale device management.
 - TR-069 support distributed architecture. The ACS can be distributed to multiple servers, each ACS can manage part of devices.

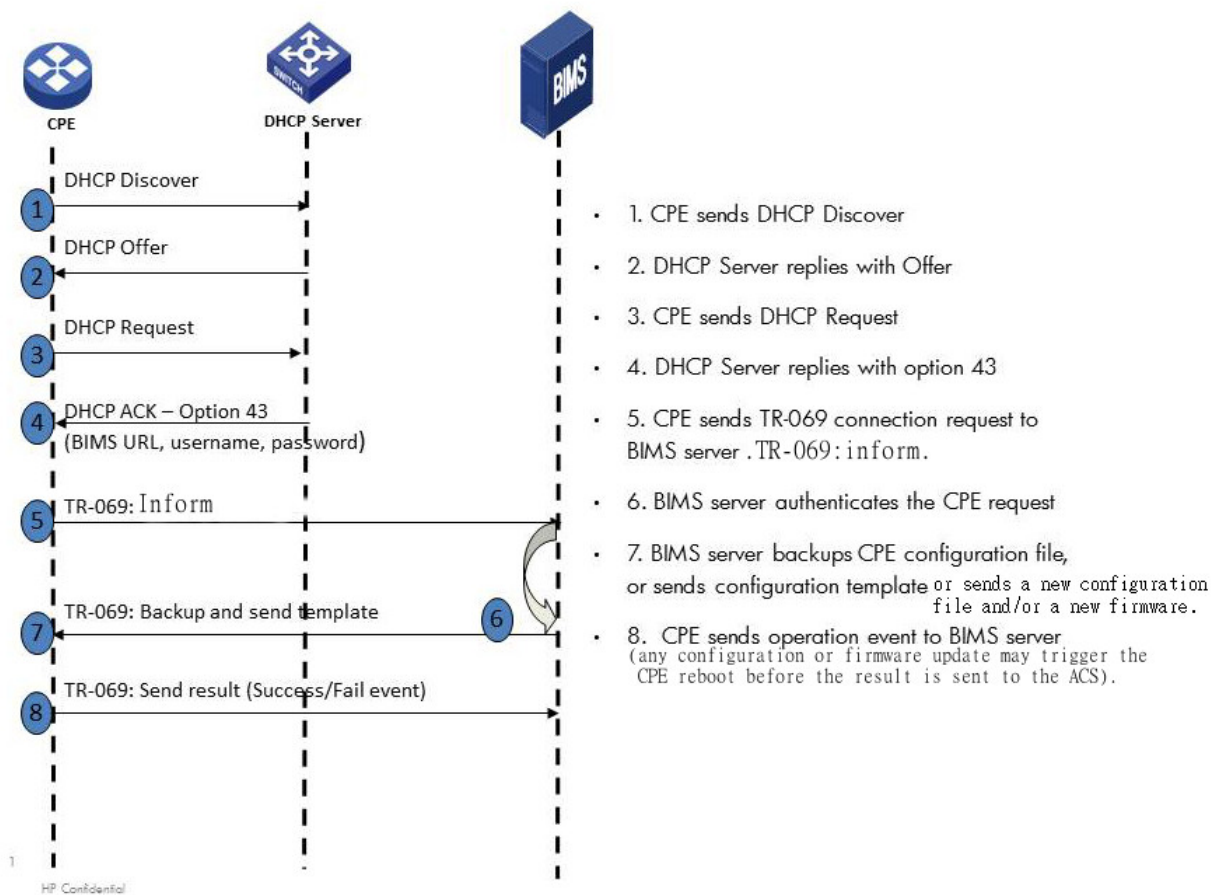
Zero-touch configuration process

Auto configuration or “zero-touch” deployment is a recurring customer requirement, especially for remote-office deployments. New devices introduced inside a private network require management tools be co-located to configure them or update firmware, or require manual intervention to do configuration. TR-069 allows managing devices that reside in a private network via HTTP(S), enabling a new set of deployment and management models today, not possible using SNMP.

The client side, when configured, will contact the server at a predefined URL, using HTTP or HTTPS as protocol. After authentication, the ACS is able to perform the following basic operations:

- Update CPE Configuration.
- Update CPE TR-069 parameters.
- Update CPE firmware.
- Reboot CPE (backup, startup, and running configurations)
- Run CPE ping diagnostics.
- Reset CPE to factory default.
- Get periodic Status (several parameters can be retrieved depending on what is supported).


Since TR-069 uses HTTP, it can be used across a WAN. If the CPE can reach the URL, it can be managed. TR-069 is mostly a push protocol where the client periodically sends information without server requests. This allows for greater scalability over traditional SNMP based tools, which are also bounded to work within the LAN, while TR-069 can offer management to remote offices.



Zero-touch configuration for Campus networks

In this example, the following steps to configure CPEs for a Campus Network environment.


1. Pre-configuration for all CPEs in BIMS.
2. CPEs get BIMS parameters from DHCP server.
3. CPEs initiate a connection to BIMS, then BIMS deploys the pre-configuration to CPEs.



Zero-touch configuration for Branch networks

In this example, the following steps to configure CPEs for a Branch network environment.

1. Create the basic configuration for your spoke device manually, using the username/password from ISP and BIMS URL.
2. The IPsec VPN configuration is generated by IVM and deployed by BIMS.
3. The IPsec VPN tunnel is automatically created.
4. The device in the branch private network can DHCP relay to HQ to continue the zero touch configuration.



Zero-touch configuration setup and execution

1. DHCP configuration
2. BIMS configuration
3. Execution

CLI commands

Configuration setup

Within the configure mode:

Syntax:

```
cwmp
```


- acs** Configure Auto Configuration Server (ACS) access.
- cpe** Configure Customer Premises Equipment (CPE) access.
- disable** Disable the CPE WAN Management Protocol.



CWMP is automatically enabled. To conserve resources, reconfigure this setting using the `cwmp disable` command.

- enable** Enable the CPE WAN Management Protocol.

Syntax:

```
[no] cwmp
```

- acs** Configure Auto Configuration Server (ACS) access.
- cpe** Configure Customer Premises Equipment (CPE) access.
- enable** Enable the CPE WAN Management Protocol.

ACS password configuration

Syntax:

```
cwmp acs
```

- password** Configure the password used for authentication when the switch connects to the ACS.
- url** Configure the URL of the ACS.
- username** Configure the username used for authentication when the switch connects to the ACS.

When encrypt-credentials is off

Syntax:

```
cwmp acs password
```

- plaintext** Configure the password used for authentication when the switch connects to the ACS.

When encrypt-credentials is on

Syntax:

```
cwmp acs password
```

- encrypted-key** An encrypted password generated with the `encrypt-credentials` command.

plaintext Configure the password used for authentication when the switch connects to the ACS.

Encrypt-credential on

```
    cwmp acs password encrypted-key
```

ASCII-STR Enter an ASCII string (maximum length: 384 characters).

Plaintext password

```
    cwmp acs password plaintext
```

PASSWORD-STR A plaintext password used for ACS authentication (maximum length: 256 characters).

ACS URL configuration

Syntax:

```
    cwmp acs url
```

URL-STR The URL of the ACS (maximum length: 256 characters).

ACS username configuration

Syntax:

```
    cwmp acs username
```

USERNAME-STR A username for ACS authentication (maximum length: 256 characters).

CPE configuration

Syntax:

```
    cwmp cpe
```

password Configure the password used for authentication when the ACS connects to the switch.

username Configure the username used for authentication when the ACS connects to the switch.

CPE password configuration

When encrypt-credentials is on

Syntax:

```
cwmp cpe password
```

encrypted-key An encrypted password generated with the 'encrypt-credentials' command.

plaintext Configure the password used for authentication when the ACS connects to the switch.

Syntax:

```
cwmp cpe password encrypted-key
```

ASCII-STR Enter an ASCII string (maximum length: 384 characters).

When encrypt-credentials is off

Syntax:

```
cwmp cpe [password]
```

plaintext Configure the password used for authentication when the ACS connects to the switch

Syntax:

```
cwmp cpe
```

PASSWORD-STR A plaintext password used for ACS authentication (maximum length: 256 characters).

CPE username configuration

Syntax:

```
cwmp cpe [username]
```

USERNAME-STR A username for ACS authentication (maximum length: 256 characters).

Enable/disable CWMP

Syntax:

```
cwmp [enable|disable]
```

Show commands

CWMP configuration and status query

Syntax:

```
show cwmp
```

- configuration** Show current CWMP configuration.
- status** Show current CWMP status.

When CWMP is enabled

Syntax:

```
show cwmp configuration
```

CWMP configuration

```
CWMP Configuration
CWMP Status       : Enabled
ACS URL           : http://16.93.62.32:9090
ACS Username      : bims
Inform Enable Status : Enabled
Inform Interval   : 60
Inform Time       : 2014-04-08T06:00:00
Reconnection Timeout : 30
```

CWMP status

```
CWMP Status
CWMP Status       : Enabled
ACS URL           : http://16.93.62.32:9090
ACS URL Origin    : Config
ACS Username      : bims
Connection Status : Disconnected
Data Transfer Status : None
Last ACS Connection Time : Wed Apr 9 16:56:00 2014
Time to Next Connection : 00:00:36
```

When CWMP is disabled

Syntax:

```
show cwmp status
```

CWMP status

```
CWMP Status
CWMP Status       : Disabled
```

CWMP configuration

```
show cwmp configuration
CWMP Configuration
CWMP Status           : Disabled
```

Event logging

The TR-069 client offers some tools to diagnose problems:

- System logging
- Status/control commands

System logging

The CPE implements the following system log notification codes and sample messages:

- **RMON_TR69_INFORM_COMPLETE**
 - INFORM to http://15.29.20.50:9090/ from (IP address not set yet) completed with error.
 - INFORM to http://15.29.20.50:9090/ from 10.0.10.212 completed with error.
 - INFORM to http://15.29.20.50:9090/ from 10.0.10.212 completed successfully.
- **RMON_TR69_AUTH_FAILED**
 - Authentication on ACS http://15.29.20.50:9090/ failed.
- **RMON_TR69_CONN_FAILED**
 - Connection attempts with ACS http://15.29.20.50:9090/ from 10.0.10.212 failed.

To avoid flooding the system log on frequent attempts to connect with the ACS, the following criteria are used with both successful and failed attempts:

1. The very first event is always logged.
2. Any change from success to failure or vice versa is always logged.
3. Repeat success or failure events are logged only once every five minutes.

The HTTP file transfer component supports these system log notification codes and sample messages:

- **RMON_HTTP_XFER_COMPLETE**
 - I 11/19/13 08:06:13 04185 http: Download of http://10.0.11.240:9876/path to DestinationFile completed successfully.
 - I 11/19/13 08:06:13 04185 http: Upload of SourceFile to http://10.0.11.240:9876/path completed successfully.
- **RMON_HTTP_CONN_FAILED**
 - W 11/19/13 08:06:13 04186 http: Connection to http://10.0.11.240:9876/path failed.
- **RMON_HTTP_TIMED_OUT**
 - W 11/19/13 08:06:13 04192 http: Download of http://10.0.11.240:9876/path to DestinationFile timed out.
 - W 02/20/14 00:32:17 04192 http: Upload of SourceFile to http://10.0.11.240:9876/path timed out.
- **RMON_HTTP_NO_SPACE**
 - W 11/19/13 08:06:13 04189 http: Upload of SourceFile to http://10.0.11.240:9876/path canceled because of insufficient memory.
- **RMON_HTTP_REQ_FAILED**
 - W 11/19/13 08:06:13 04190 http: Upload of SourceFile to http://10.0.11.240:9876/path failed (errno 13).
 - W 11/19/13 08:06:13 04190 http: Upload of SourceFile to http://10.0.11.240:9876/path failed (errno 1).

- W 11/19/13 08:06:13 04190 http: Download of http://10.0.11.240:9876/path to DestinationFile failed (errno 13).
- W 11/19/13 08:06:13 04190 http: Download of http://10.0.11.240:9876/path to DestinationFile failed (errno 1).
- W 11/19/13 08:06:13 04190 http: Download of http://10.0.11.240:9876/path to DestinationFile failed (errno 17).
- **RMON_HTTP_WRONG_FILE**
 - W 11/19/13 08:06:13 04191 http: Download canceled because file http://10.0.11.240:9876/path is malformed or incompatible.
 - W 11/19/13 08:06:13 04191 http: Download canceled because file http://10.0.11.240:9876/path is malformed or incompatible.
- **RMON_HTTP_FILE_NOT_FOUND**
 - W 11/19/13 08:06:13 04200 http: Upload of SourceFile to http://10.0.11.240:9876/path canceled because of inexistent file.

Status/control commands

The following commands help assess the general state of TR-069 and control the source of the ACS configuration record:

Table 29: Status/control commands

Command	Result
show cwmp status	<pre> CWMP is Enabled ACS URL : https://16.93.62.32:9443 ACS URL is set by : Config ACS Username : bims Connection status : Disconnected Data transfer status : None Time of last successful connection : Thu Feb 20 01:16:59 2014 Interval upon to next connection : Null </pre>
show cwmp configuration	<pre> CWMP is Enabled ACS URL : https://16.93.62.32:9443 ACS Username : bims Inform Enable Status : Disabled Inform Interval : 3559 Inform Time : Reconnection times : 30 </pre>
[no] dhcp tr69-acs-url	Prevents using any ACS information from DHCP

Acronym	Definition
ACL	Access Control List
AMP	AirWave Management Platform
AP	Access Point
BYOD	Bring Your Own Device
BPS	Backplane Stacking
CoA	Change of Authorization
CLI	Command Line Interface
CPPM	ClearPass Policy Manager
DHCP	Dynamic Host Configuration Protocol
DoS	Denial-of-Service
EWA	Enhanced Web Authentication
IP	Internet Protocol
HA	High Availability
HMAC-SHA1	Hash-based Message Authentication Code used with the SHA-1 cryptographic hash function.
HTTP	Hypertext Transfer Protocol
HTTPS	Secure Hypertext Transfer Protocol
ID	Identifier
IP	Internet Protocol
L3	The third, or routing, layer of the open systems interconnection (OSI) model. The network layer routes data to different LANs and Wide Area Networks (WANs) based on network addresses.

Table Continued

Acronym	Definition
LAN	Local Area Network
MAC	Media Access Control
MAFR	MAC Authentication Failure Redirect
MAS	Management Interface Specification
NMS	Network Management System
PVOS	ArubaOS-Switch Operating System
RADIUS	Remote Authentication Dial In User Service
SNMP	Simple Network Management Protocol
VLAN	Virtual Local Area Network
VSA	Vendor Specific Attribute
VSF	Virtual Switching Framework
ZTP	Zero Touch Provisioning