

HP Switch Software

Management and Configuration Guide for RA.15.14

Abstract

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. The information in this guide is subject to change without notice.

Applicable Products

HP Switch 2620-series (J9623A, J9624A, J9625A, J9626A, J9627A)



© Copyright 2012 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP 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.

The information contained herein is subject to change without notice. The only warranties for HP 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. HP shall not be liable for technical or editorial errors or omissions contained herein.

Acknowledgments

Microsoft®, Windows®, Windows® XP, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

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

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

Warranty

For the software end user license agreement and the hardware limited warranty information for HP Networking products, visit www.hp.com/networking.

Contents

1 Product Documentation.....	15
About your switch manual set.....	15
Printed publications.....	15
Electronic publications.....	15
2 Time Protocols.....	16
General steps for running a time protocol on the switch.....	16
TimeP time synchronization.....	16
SNTP time synchronization.....	16
Selecting a time synchronization protocol.....	16
Disabling time synchronization.....	17
SNTP: Selecting and configuring.....	17
Viewing and configuring SNTP (Menu).....	18
Viewing and configuring SNTP (CLI).....	19
Configuring (enabling or disabling) the SNTP mode.....	20
Enabling SNTP in Broadcast Mode.....	21
Enabling SNTP in unicast mode (CLI).....	22
Changing the SNTP poll interval (CLI).....	24
Changing the SNTP server priority (CLI).....	24
Disabling time synchronization without changing the SNTP configuration (CLI).....	24
Disabling the SNTP Mode.....	25
SNTP client authentication.....	25
Requirements.....	26
Configuring the key-identifier, authentication mode, and key-value (CLI).....	26
Configuring a trusted key.....	27
Configuring a key-id as trusted (CLI).....	27
Associating a key with an SNTP server (CLI).....	27
Enabling SNTP client authentication.....	28
Configuring unicast and broadcast mode for authentication.....	28
Viewing SNTP authentication configuration information (CLI).....	29
Viewing all SNTP authentication keys that have been configured on the switch (CLI).....	29
Viewing statistical information for each SNTP server (CLI).....	29
Saving configuration files and the include-credentials command.....	30
TimeP: Selecting and configuring.....	31
Viewing, enabling, and modifying the TimeP protocol (Menu).....	32
Viewing the current TimeP configuration (CLI).....	33
Configuring (enabling or disabling) the TimeP mode.....	34
Enabling TimeP in manual mode (CLI).....	34
SNTP unicast time polling with multiple SNTP servers.....	37
Displaying all SNTP server addresses configured on the switch (CLI).....	37
Adding and deleting SNTP server addresses.....	38
Adding addresses.....	38
Deleting addresses.....	38
Operating with multiple SNTP server addresses configured (Menu).....	38
SNTP messages in the Event Log.....	39
3 Port Status and Configuration.....	40
Viewing port status and configuring port parameters.....	40
Connecting transceivers to fixed-configuration devices.....	40
Viewing port configuration (Menu).....	41
Configuring ports (Menu).....	42
Viewing port status and configuration (CLI).....	43

Dynamically updating the show interfaces command (CLI/Menu).....	43
Customizing the show interfaces command (CLI).....	44
Error messages associated with the show interfaces command.....	45
Note on using pattern matching with the show interfaces custom command.....	45
Viewing port utilization statistics (CLI).....	46
Operating notes for viewing port utilization statistics.....	46
Viewing transceiver status (CLI).....	46
Operating notes.....	47
Enabling or disabling ports and configuring port mode (CLI).....	47
Enabling or disabling flow control (CLI).....	48
Configuring auto-MDIX.....	50
Manual override.....	51
Configuring auto-MDIX (CLI).....	51
Using friendly (optional) port names.....	53
Configuring and operating rules for friendly port names.....	53
Configuring friendly port names (CLI).....	54
Configuring a single port name (CLI).....	54
Configuring the same name for multiple ports (CLI).....	54
Displaying friendly port names with other port data (CLI).....	55
Listing all ports or selected ports with their friendly port names (CLI).....	55
Including friendly port names in per-port statistics listings (CLI).....	56
Searching the configuration for ports with friendly port names (CLI).....	57
Uni-directional link detection (UDLD).....	58
Configuring UDLD.....	59
Configuring uni-directional link detection (UDLD) (CLI).....	59
Enabling UDLD (CLI).....	60
Changing the keepalive interval (CLI).....	60
Changing the keepalive retries (CLI).....	60
Configuring UDLD for tagged ports.....	60
Viewing UDLD information (CLI).....	61
Viewing summary information on all UDLD-enabled ports (CLI).....	61
Viewing detailed UDLD information for specific ports (CLI).....	61
Clearing UDLD statistics (CLI).....	62
4 Power Over Ethernet (PoE/PoE+) Operation.....	63
Introduction to PoE.....	63
PoE terminology.....	63
About PoE operation.....	63
Configuration options.....	63
PD support.....	64
Power priority operation.....	64
Configuring PoE operation.....	65
Disabling or re-enabling PoE port operation.....	65
Enabling support for pre-standard devices.....	65
Configuring the PoE port priority.....	65
Controlling PoE allocation.....	66
Manually configuring PoE power levels.....	67
Changing the threshold for generating a power notice.....	69
PoE/PoE+ allocation using LLDP information.....	69
LLDP with PoE.....	69
Enabling or disabling ports for allocating power using LLDP.....	70
Enabling PoE detection via LLDP TLV advertisement.....	70
LLDP with PoE+.....	70
Overview.....	70
PoE allocation.....	70

Initiating advertisement of PoE+ TLVs.....	71
Viewing PoE when using LLDP information.....	72
Operation Note.....	73
Viewing the global PoE power status of the switch.....	74
Viewing PoE status on all ports.....	75
Viewing the PoE status on specific ports.....	77
Planning and implementing a PoE configuration.....	78
Power requirements.....	79
Assigning PoE ports to VLANs.....	79
Applying security features to PoE configurations.....	79
Assigning priority policies to PoE traffic.....	79
PoE Event Log messages.....	80
5 Port Trunking.....	81
Overview of port trunking.....	81
Port connections and configuration.....	81
Port trunk features and operation.....	82
Fault tolerance	82
Trunk configuration methods.....	82
Dynamic LACP trunk.....	82
Static trunk.....	83
Viewing and configuring a static trunk group (Menu).....	85
Viewing and configuring port trunk groups (CLI).....	87
Viewing static trunk type and group for all ports or for selected ports.....	87
Viewing static LACP and dynamic LACP trunk data.....	88
Dynamic LACP Standby Links.....	88
Configuring a static trunk or static LACP trunk group.....	89
Removing ports from a static trunk group.....	89
Enabling a dynamic LACP trunk group.....	90
Removing ports from a dynamic LACP trunk group.....	90
Viewing existing port trunk groups (WebAgent).....	91
Trunk group operation using LACP.....	91
Default port operation.....	93
LACP notes and restrictions.....	93
802.1X (Port-based access control) configured on a port.....	93
Port security configured on a port.....	94
Changing trunking methods.....	94
Static LACP trunks.....	94
Dynamic LACP trunks.....	94
VLANs and dynamic LACP.....	94
Blocked ports with older devices.....	95
Spanning Tree and IGMP.....	95
Half-duplex, different port speeds, or both not allowed in LACP trunks.....	95
Dynamic/static LACP interoperation.....	96
Trunk group operation using the "trunk" option.....	96
How the switch lists trunk data.....	96
Outbound traffic distribution across trunked links.....	96
Trunk load balancing using port layers.....	98
Enabling trunk load balancing.....	98
6 Port Traffic Controls.....	100
Rate-limiting.....	100
All traffic rate-limiting.....	100
Configuring rate-limiting.....	100
Displaying the current rate-limit configuration.....	101
Operating notes for rate-limiting.....	102

ICMP rate-limiting.....	103
Terminology.....	104
Guidelines for configuring ICMP rate-limiting.....	104
Configuring ICMP rate-limiting.....	104
Using both ICMP rate-limiting and all-traffic rate-limiting on the same interface.....	105
Viewing the current ICMP rate-limit configuration.....	106
Operating notes for ICMP rate-limiting.....	106
Notes on testing ICMP rate-limiting.....	107
ICMP rate-limiting trap and Event Log messages.....	108
Determining the switch port number used in ICMP port reset commands.....	108
Configuring inbound rate-limiting for broadcast and multicast traffic.....	109
Operating Notes.....	110
Jumbo frames.....	110
Terminology.....	111
Operating rules.....	111
Configuring jumbo frame operation.....	111
Overview.....	111
Viewing the current jumbo configuration.....	111
Enabling or disabling jumbo traffic on a VLAN.....	113
Configuring a maximum frame size.....	113
Configuring IP MTU.....	113
SNMP implementation.....	114
Jumbo maximum frame size.....	114
Jumbo IP MTU.....	114
Displaying the maximum frame size.....	114
Operating notes for maximum frame size.....	114
Operating notes for jumbo traffic-handling.....	115
Troubleshooting.....	116
A VLAN is configured to allow jumbo frames, but one or more ports drops all inbound jumbo frames.....	116
A non-jumbo port is generating "Excessive undersize/giant frames" messages in the Event Log.....	116
7 Configuring for Network Management Applications.....	117
Using SNMP tools to manage the switch.....	117
SNMP management features.....	117
SNMPv1 and v2c access to the switch.....	118
SNMPv3 access to the switch.....	118
Enabling and disabling switch for access from SNMPv3 agents.....	118
Enabling or disabling restrictions to access from only SNMPv3 agents.....	119
Enabling or disabling restrictions from all non-SNMPv3 agents to read-only access.....	119
Viewing the operating status of SNMPv3.....	119
Viewing status of message reception of non-SNMPv3 messages.....	119
Viewing status of write messages of non-SNMPv3 messages.....	119
Enabling SNMPv3.....	119
SNMPv3 users.....	120
Adding users.....	120
SNMPv3 user commands.....	121
Listing Users.....	121
Assigning users to groups (CLI).....	122
Group access levels.....	122
SNMPv3 communities.....	123
Mapping SNMPv3 communities (CLI).....	123
SNMP community features.....	124
Viewing and configuring non-version-3 SNMP communities (Menu).....	125

Listing community names and values (CLI).....	125
Configuring community names and values (CLI).....	126
SNMP notifications.....	127
Supported Notifications.....	127
General steps for configuring SNMP notifications.....	128
SNMPv1 and SNMPv2c Traps.....	128
SNMP trap receivers.....	128
Configuring an SNMP trap receiver (CLI).....	129
SNMPv2c informs.....	130
Enabling SNMPv2c informs (CLI).....	130
Configuring SNMPv3 notifications (CLI).....	131
Network security notifications.....	133
Enabling or disabling notification/traps for network security failures and other security events (CLI).....	134
Viewing the current configuration for network security notifications (CLI).....	135
Enabling Link-Change Traps (CLI).....	135
Readable interface names in traps.....	135
Source IP address for SNMP notifications.....	135
Configuring the source IP address for SNMP notifications (CLI).....	136
Viewing SNMP notification configuration (CLI).....	137
Configuring the MAC address count option.....	138
Displaying information about the mac-count-notify option.....	139
Advanced management: RMON.....	140
CLI-configured sFlow with multiple instances.....	140
Configuring sFlow (CLI).....	141
Viewing sFlow Configuration and Status (CLI).....	141
LLDP.....	143
General LLDP operation.....	143
LLDP-MED.....	143
Packet boundaries in a network topology.....	144
LLDP operation configuration options.....	144
Enable or disable LLDP on the switch.....	144
Enable or disable LLDP-MED.....	144
Change the frequency of LLDP packet transmission to neighbor devices.....	144
Change the Time-To-Live for LLDP packets sent to neighbors.....	144
Transmit and receive mode.....	144
SNMP notification.....	145
Per-port (outbound) data options.....	145
Remote management address.....	146
Debug logging.....	146
Options for reading LLDP information collected by the switch.....	146
LLDP and LLDP-MED standards compatibility.....	146
LLDP operating rules.....	146
Port trunking.....	147
IP address advertisements.....	147
Spanning-tree blocking.....	147
802.1X blocking.....	147
Configuring LLDP operation.....	147
Displaying the global LLDP, port admin, and SNMP notification status (CLI).....	147
Viewing port configuration details (CLI).....	148
Configuring Global LLDP Packet Controls.....	149
LLDP operation on the switch.....	149
Enabling or disabling LLDP operation on the switch (CLI).....	149
Changing the packet transmission interval (CLI).....	149
Time-to-Live for transmitted advertisements.....	150

Delay interval between advertisements generated by value or status changes to the LLDP MIB.....	150
Reinitialization delay interval.....	151
Configuring SNMP notification support.....	152
Enabling LLDP data change notification for SNMP trap receivers (CLI).....	152
Changing the minimum interval for successive data change notifications for the same neighbor.....	152
Configuring per-port transmit and receive modes (CLI).....	152
Basic LLDP per-port advertisement content.....	153
Mandatory Data.....	153
Configuring a remote management address for outbound LLDP advertisements (CLI).....	153
Optional Data.....	154
Support for port speed and duplex advertisements.....	154
Configuring support for port speed and duplex advertisements (CLI).....	155
Port VLAN ID TLV support on LLDP.....	155
Configuring the VLAN ID TLV.....	155
Viewing the TLVs advertised.....	156
SNMP support.....	158
LLDP-MED (media-endpoint-discovery).....	158
LLDP-MED endpoint support.....	159
LLDP-MED endpoint device classes.....	160
LLDP-MED operational support.....	160
Tracking LLDP-MED connects and disconnects—topology change notification.....	160
LLDP-MED fast start control.....	161
Advertising device capability, network policy, PoE status and location data.....	161
Network policy advertisements.....	162
VLAN operating rules.....	162
Policy elements.....	162
Enabling or Disabling medTlvEnable.....	162
PoE advertisements.....	163
Location data for LLDP-MED devices.....	164
Configuring location data for LLDP-MED devices.....	164
Configuring coordinate-based locations.....	166
Viewing switch information available for outbound advertisements.....	167
Displaying the current port speed and duplex configuration on a switch port.....	169
Viewing the current port speed and duplex configuration on a switch port.....	170
Viewing advertisements currently in the neighbors MIB.....	170
Displaying LLDP statistics.....	171
Viewing LLDP statistics.....	171
LLDP Operating Notes.....	173
Neighbor maximum.....	173
LLDP packet forwarding.....	173
One IP address advertisement per port.....	173
802.1Q VLAN Information.....	174
Effect of 802.1X Operation.....	174
Neighbor data can remain in the neighbor database after the neighbor is disconnected.....	174
Mandatory TLVs.....	174
Determining the switch port number included in topology change notification traps.....	174
LLDP and CDP data management.....	174
LLDP and CDP neighbor data.....	175
CDP operation and commands.....	176
Viewing the current CDP configuration of the switch.....	176
Viewing the current CDP neighbors table of the switch.....	177
Enabling and Disabling CDP Operation.....	177
Enabling or disabling CDP operation on individual ports.....	178

Configuring CDPv2 for voice transmission.....	178
Filtering CDP information.....	180
Configuring the switch to filter untagged traffic.....	180
Displaying the configuration.....	180
Filtering PVID mismatch log messages.....	181
8 Link Aggregation Control Protocol—Multi-Active Detection (LACP-MAD).....	182
LACP-MAD commands.....	182
Configuration command.....	182
show commands.....	182
clear command.....	182
LACP-MAD overview.....	182
A File transfers.....	183
Overview.....	183
Downloading switch software.....	183
General software download rules.....	183
Using TFTP to download software from a server.....	183
Downloading from a server to primary flash using TFTP (Menu).....	184
Troubleshooting TFTP download failures.....	185
Downloading from a server to flash using TFTP (CLI).....	186
Enabling TFTP (CLI).....	187
Configuring the switch to download software automatically from a TFTP server using auto-TFTP (CLI).....	187
Using SCP and SFTP.....	188
Enabling SCP and SFTP.....	189
Disabling TFTP and auto-TFTP for enhanced security.....	189
Enabling SSH V2 (required for SFTP).....	191
Confirming that SSH is enabled.....	191
Disabling secure file transfer.....	191
Authentication.....	191
SCP/SFTP operating notes.....	191
Troubleshooting SSH, SFTP, and SCP operations.....	192
Broken SSH connection.....	193
Attempt to start a session during a flash write.....	193
Failure to exit from a previous session.....	193
Attempt to start a second session.....	193
Using Xmodem to download switch software from a PC or UNIX workstation.....	194
Downloading to primary flash using Xmodem (Menu).....	194
Downloading to primary or secondary flash using Xmodem and a terminal emulator (CLI)....	194
Switch-to-switch download.....	195
Switch-to-switch download to primary flash (Menu).....	195
Downloading the OS from another switch (CLI).....	196
Downloading from primary only (CLI).....	196
Downloading from either flash in the source switch to either flash in the destination switch (CLI).....	197
Copying software images.....	197
TFTP: Copying a software image to a remote host (CLI).....	197
Xmodem: Copying a software image from the switch to a serially connected PC or UNIX workstation (CLI).....	197
Transferring switch configurations.....	198
TFTP: Copying a configuration file to a remote host (CLI).....	198
TFTP: Copying a configuration file from a remote host (CLI).....	198
TFTP: Copying a customized command file to a switch (CLI).....	199
Xmodem: Copying a configuration file to a serially connected PC or UNIX workstation (CLI).....	199
Xmodem: Copying a configuration file from a serially connected PC or UNIX workstation (CLI)....	200

Transferring ACL command files.....	200
TFTP: Uploading an ACL command file from a TFTP server (CLI).....	201
Xmodem: Uploading an ACL command file from a serially connected PC or UNIX workstation (CLI).....	202
Copying diagnostic data to a remote host, PC or UNIX workstation.....	203
Copying command output to a destination device (CLI).....	203
Copying Event Log output to a destination device (CLI).....	203
Copying crash data content to a destination device (CLI).....	204
B Monitoring and Analyzing Switch Operation.....	205
Overview.....	205
Status and counters data.....	205
Accessing status and counters (Menu).....	205
General system information.....	205
Accessing system information (Menu).....	205
Accessing system information (CLI).....	206
Collecting processor data with the task monitor (CLI).....	207
Switch management address information.....	208
Accessing switch management address information (Menu).....	208
Accessing switch management address information (CLI).....	209
Port Status.....	209
Viewing port status (CLI).....	209
Viewing port status (Menu).....	209
Viewing port and trunk group statistics (WebAgent).....	209
Port and trunk group statistics and flow control status.....	209
Accessing port and trunk statistics (Menu).....	210
Accessing port and trunk group statistics (CLI).....	211
Viewing the port counter summary report.....	211
Viewing a detailed traffic summary for specific ports.....	211
Displaying trunk load balancing statistics.....	211
Clearing trunk load balancing statistics.....	211
Resetting the port counters.....	212
Viewing the switch's MAC address tables.....	212
Accessing MAC address views and searches (CLI).....	212
Listing all learned MAC addresses on the switch, with the port number on which each MAC address was learned.....	212
Listing all learned MAC addresses on one or more ports, with their corresponding port numbers.....	212
Listing all learned MAC addresses on a VLAN, with their port numbers.....	212
Finding the port on which the switch learned a specific MAC address.....	212
Accessing MAC address views and searches (Menu).....	212
Viewing and searching per-VLAN MAC-addresses.....	212
Finding the port connection for a specific device on a VLAN.....	213
Viewing and searching port-level MAC addresses.....	214
Determining whether a specific device is connected to the selected port.....	214
Accessing MSTP Data (CLI).....	214
Viewing internet IGMP status (CLI).....	215
Viewing VLAN information (CLI).....	216
WebAgent status information.....	218
C Troubleshooting.....	220
Overview.....	220
Troubleshooting approaches.....	220
Browser or Telnet access problems.....	221
Cannot access the WebAgent.....	221
Cannot Telnet into the switch console from a station on the network.....	221

Unusual network activity.....	222
General problems.....	222
The network runs slow; processes fail; users cannot access servers or other devices.....	222
Duplicate IP addresses.....	222
Duplicate IP addresses in a DHCP network.....	222
The switch has been configured for DHCP/Bootp operation, but has not received a DHCP or Bootp reply.....	222
802.1Q Prioritization problems.....	223
Ports configured for non-default prioritization (level 1 to 7) are not performing the specified action.....	223
Addressing ACL problems.....	223
ACLs are properly configured and assigned to VLANs, but the switch is not using the ACLs to filter IP layer 3 packets.....	223
The switch does not allow management access from a device on the same VLAN.....	224
Error (Invalid input) when entering an IP address.....	224
Apparent failure to log all "deny" matches.....	224
The switch does not allow any routed access from a specific host, group of hosts, or subnet.....	224
The switch is not performing routing functions on a VLAN.....	224
Routing through a gateway on the switch fails.....	224
Remote gateway case.....	225
Local gateway case.....	225
IGMP-related problems.....	225
IP multicast (IGMP) traffic that is directed by IGMP does not reach IGMP hosts or a multicast router connected to a port.....	225
IP multicast traffic floods out all ports; IGMP does not appear to filter traffic.....	226
LACP-related problems.....	226
Unable to enable LACP on a port with the interface <i><port-number></i> lacp command.....	226
Mesh-related problems.....	226
Traffic on a dynamic VLAN does not get through the switch mesh.....	226
Port-based access control (802.1X)-related problems.....	226
The switch does not receive a response to RADIUS authentication requests.....	226
The switch does not authenticate a client even though the RADIUS server is properly configured and providing a response to the authentication request.....	227
During RADIUS-authenticated client sessions, access to a VLAN on the port used for the client sessions is lost.....	227
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.....	227
The supplicant statistics listing shows multiple ports with the same authenticator MAC address.....	227
The show port-access authenticator <i><port-list></i> command shows one or more ports remain open after they have been configured with control unauthorized.....	227
RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch.....	228
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 <i><port-list></i> initialize.....	228
A trunked port configured for 802.1X is blocked.....	229
QoS-related problems.....	229
Loss of communication when using VLAN-tagged traffic.....	229
Radius-related problems.....	229
The switch does not receive a response to RADIUS authentication requests.....	229
RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch.....	229
MSTP and fast-uplink problems.....	230
Broadcast storms appearing in the network.....	230

STP blocks a link in a VLAN even though there are no redundant links in that VLAN.....	230
Fast-uplink troubleshooting.....	230
SSH-related problems.....	231
Switch access refused to a client.....	231
Executing IP SSH does not enable SSH on the switch.....	231
Switch does not detect a client's public key that does appear in the switch's public key file (show ip client-public-key).....	231
An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages.....	231
Client ceases to respond ("hangs") during connection phase.....	231
TACACS-related problems.....	231
Event Log.....	231
All users are locked out of access to the switch.....	232
No communication between the switch and the TACACS+ server application.....	232
Access is denied even though the username/password pair is correct.....	232
Unknown users allowed to login to the switch.....	232
System allows fewer login attempts than specified in the switch configuration.....	232
TimeP, SNTP, or Gateway problems.....	233
The switch cannot find the time server or the configured gateway.....	233
VLAN-related problems.....	233
Monitor port.....	233
None of the devices assigned to one or more VLANs on an 802.1Q-compliant switch are being recognized.....	233
Link configured for multiple VLANs does not support traffic for one or more VLANs.....	233
Duplicate MAC addresses across VLANs.....	233
Disabled overlapping subnet configuration.....	234
Fan failure.....	235
Mitigating flapping transceivers.....	235
Viewing transceiver information.....	237
Viewing information about transceivers (CLI).....	238
MIB support.....	238
Viewing transceiver information.....	238
Information displayed with the detail parameter.....	239
Viewing transceiver information for copper transceivers with VCT support.....	242
Testing the Cable.....	242
Using the Event Log for troubleshooting switch problems.....	244
Event Log entries.....	244
Using the Menu.....	252
Using the CLI.....	253
Clearing Event Log entries.....	254
Turning event numbering on.....	254
Using log throttling to reduce duplicate Event Log and SNMP messages.....	254
Log throttle periods.....	255
Example of event counter operation.....	256
Reporting information about changes to the running configuration.....	256
Debug/syslog operation.....	257
Debug/syslog messaging.....	257
Debug/syslog destination devices.....	257
Debug/syslog configuration commands.....	258
Configuring debug/syslog operation.....	260
Viewing a debug/syslog configuration.....	261
Debug command.....	263
Debug messages.....	263
Filtering debug messages by debug type.....	265
Debug destinations.....	266

Logging command.....	267
Configuring a syslog server.....	268
Deleting syslog addresses in the startup configuration.....	268
Verifying the deletion of a syslog server address.....	268
Blocking the messages sent to configured syslog servers from the currently configured debug message type.....	268
Disabling syslog logging on the switch without deleting configured server addresses.....	269
Sending logging messages using TCP.....	269
Adding a description for a Syslog server.....	270
Adding a priority description.....	271
Configuring the severity level for Event Log messages sent to a syslog server.....	271
Configuring the system module used to select the Event Log messages sent to a syslog server.....	272
Operating notes for debug and Syslog.....	272
Diagnostic tools.....	273
Port auto-negotiation.....	273
Ping and link tests.....	273
Ping test.....	273
Link test.....	273
Executing ping or link tests (WebAgent).....	273
Testing the path between the switch and another device on an IP network.....	274
Halting a ping test.....	275
Issuing single or multiple link tests.....	275
Tracing the route from the switch to a host address.....	276
Halting an ongoing traceroute search.....	277
A low maxttl causes traceroute to halt before reaching the destination address.....	277
If a network condition prevents traceroute from reaching the destination.....	278
Viewing switch configuration and operation.....	278
Viewing the startup or running configuration file.....	278
Viewing the configuration file (WebAgent).....	279
Viewing a summary of switch operational data.....	279
Saving show tech command output to a text file.....	280
Customizing show tech command output.....	281
Viewing more information on switch operation.....	282
Searching for text using pattern matching with show command.....	283
Displaying the information you need to diagnose problems.....	285
Restoring the factory-default configuration.....	286
Resetting to the factory-default configuration.....	286
Using the CLI.....	286
Using Clear/Reset.....	287
Restoring a flash image.....	287
Recovering from an empty or corrupted flash state.....	287
DNS resolver.....	288
Basic operation.....	289
Configuring and using DNS resolution with DNS-compatible commands.....	290
Configuring a DNS entry.....	290
Using DNS names with ping and traceroute: Example.....	291
Viewing the current DNS configuration.....	292
Operating notes.....	293
Event Log messages.....	293
Locating a switch (Locator LED).....	293
D MAC Address Management.....	295
Overview.....	295
Determining MAC addresses.....	295

Viewing the MAC addresses of connected devices.....	295
Viewing the switch's MAC address assignments for VLANs configured on the switch.....	295
Viewing the port and VLAN MAC addresses.....	296
E Monitoring Resources.....	298
Displaying current resource usage.....	298
Viewing information on resource usage.....	299
Policy enforcement engine.....	300
Usage notes for show resources output.....	300
When insufficient resources are available.....	301
F Monitoring Resources.....	302
Viewing information on resource usage.....	302
Policy enforcement engine.....	302
Displaying current resource usage.....	302
Usage notes for show resources output.....	303
When insufficient resources are available.....	303
G Daylight Saving Time on HP Switches.....	305
H Scalability: IP Address, VLAN, and Routing Maximum Values.....	307
I Power-Saving Features.....	309
Configuring the savepower LED option.....	309
J Support and Other Resources.....	310
Intended audience.....	310
Related documentation.....	310
Contacting HP.....	310
HP technical support.....	310
Subscription service.....	310
Related information.....	310
HP websites.....	310
Typographical conventions.....	311
Command syntax statements.....	311
Command prompts.....	312
Screen simulations.....	312
Configuration and operation examples.....	312
Keys.....	312
To set up and install the switch in your network.....	312
Physical installation.....	312
Product warranties.....	313
Online help.....	313
Menu interface.....	313
Command-line interface.....	313
HP customer support services.....	313
Before calling support.....	313
Glossary of Terms and Acronyms.....	314
Index.....	320

1 Product Documentation

About your switch manual set

NOTE: For the latest version of all HP switch documentation, including Release Notes covering recently added features, please visit the HP Networking Web site at www.hp.com/networking.

Printed publications

The *Read Me First* included with your switch provides software update information, product notes, and other information. The latest version is also available in PDF format on the HP website, as described in the Note at the top of this page.

Electronic publications

The latest version of each of the publications listed below is available in PDF format on the HP website, as described in the Note at the top of this page.

- *Installation and Getting Started Guide*—Explains how to prepare for and perform the physical installation and connect the switch to your network.
- *Management and Configuration Guide*—Describes how to configure, manage, and monitor basic switch operation.
- *Advanced Traffic Management Guide*—Explains how to configure traffic management features such as VLANs, MSTP, QoS, and Meshing.
- *Multicast and Routing Guide*—Explains how to configure IGMP, PIM, IP routing, and VRRP features.
- *Access Security Guide*—Explains how to configure access security features and user authentication on the switch.
- *IPv6 Configuration Guide*—Describes the IPv6 protocol operations that are supported on the switch.
- *Command Line Interface Reference Guide*—Provides a comprehensive description of CLI commands, syntax, and operations.
- *Event Log Message Reference Guide*—Provides a comprehensive description of event log messages.
- *Release Notes*—Describe new features, fixes, and enhancements that become available between revisions of the main product guide.

2 Time Protocols

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 and SNTP (Simple Network Time Protocol) and a `timesync` command for changing the time protocol selection (or turning off time protocol operation).

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

In the factory-default configuration, the time synchronization option is set to TimeP, with the TimeP mode itself set to Disabled.

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 two 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.

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

- **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

1. Select the time synchronization protocol: SNTP or TimeP (the default).
2. Enable the protocol; the choices are:
 - SNTP: Broadcast or Unicast
 - TimeP: DHCP or Manual
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 or SNTP configuration:

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

SNTP: Selecting and configuring

Table 1 (page 17) shows the SNTP parameters and their operations.

Table 1 SNTP parameters

SNTP parameter	Operation
Time Sync Method	Used to select either SNTP, TIMEP, or None as the time synchronization method.
SNTP Mode	
Disabled	The Default. SNTP does not operate, even if specified by the Menu interface Time Sync Method parameter or the CLI <code>timesync</code> command.
Unicast	Directs the switch to poll a specific server for SNTP time synchronization. Requires at least one server address.
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)	<p>In Unicast Mode: Specifies how often the switch polls the designated SNTP server for a time update.</p> <p>In Broadcast Mode: Specifies how often the switch polls the network broadcast address for a time update.</p> <p>Value is between 30 to 720 seconds.</p>
Server Address	Used only when the SNTP Mode is set to <code>Unicast</code> . 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)

1. From the Main Menu, select:
 2. **Switch Configuration...**
 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 step 6 (page 19). (For Broadcast mode details, see “SNTP time synchronization” (page 16))

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.

NOTE: 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” (page 37).

- 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, HP recommends leaving this value at the default setting of 3 and testing SNTP operation to determine whether any change is necessary.

NOTE: 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” (page 37).

- 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 [Table 1 \(page 17\)](#), on “SNTP parameters” (page 17).)
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:

Example 1 SNTP configuration when SNTP is the selected time synchronization method

```
HP 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 [Example 2](#), even though TimeP is the current time synchronous method, the switch maintains the SNTP configuration.

Example 2 SNTP configuration when SNTP is not the selected time synchronization method

```
HP 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.

Example 3 Display showing IP addressing for all configured time servers and VLANs

```
HP 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:

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:

Example 4 Enabling SNTP operation in Broadcast Mode

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

```
HP Switch(config)# timesync sntp
```

```
HP Switch(config)# sntp broadcast
```

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

- | | |
|--|--|
| <p>1 show sntp displays the SNTP configuration and also shows that TimeP is the currently active time synchronization mode.</p> | <p>2 show sntp again displays the SNTP configuration and shows that SNTP is now the currently active time synchronization mode and is configured for broadcast operation.</p> |
|--|--|
-

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” \(page 37\)](#)

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 <ip-addr>
Deletes the specified SNTP server.
```

NOTE: 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):

```
HP Switch(config)# timesync sntp
```

Selects SNTP.

```
HP Switch(config)# sntp unicast
```

Activates SNTP in unicast mode.

```
HP Switch(config)# sntp server priority 1 10.28.227.141
```

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

Example 5 Configuring SNTP for unicast operation

```
HP 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.

Example 6 Specifying the SNTP protocol version number

```
HP Switch(config)# no sntp server 10.28.227.141 1
HP Switch(config)# sntp server 10.28.227.141 4 2
HP 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
```

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

```
HP 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.

NOTE: 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:

```
HP Switch(config)# sntp server priority 1 10.28.22.141
```

```
HP 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:

```
HP Switch(config)# no timesync
```

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

Example 7 SNTP with time synchronization disabled

```
HP 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.

Example 8 Disabling time synchronization by disabling the SNTP mode

```
HP Switch(config)# no sntp
HP 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.

SNTP client authentication

Enabling SNTP authentication allows network devices such as HP switches to validate the SNTP messages received from an NTP or SNTP server before updating the network time. NTP or SNTP servers and clients must be configured with the same set of authentication keys so that the servers can authenticate the messages they send and clients (HP switches) can validate the received messages before updating the time.

This feature provides support for SNTP client authentication on HP switches, which addresses security considerations when deploying SNTP in a network.

Requirements

You must configure the following to enable SNTP client authentication on the switch.

SNTP client authentication support

- Timesync mode must be SNTP. Use the `timesync sntp` command. (SNTP is disabled by default).
- SNTP must be in unicast or broadcast mode. See “Configuring unicast and broadcast mode for authentication” (page 28).
- The MD5 authentication mode must be selected.
- An SNTP authentication key-identifier (`key-id`) must be configured on the switch and a value (`key-value`) must be provided for the authentication key. A maximum of 8 sets of `key-id` and `key-value` can be configured on the switch.
- Among the keys that have been configured, one key or a set of keys must be configured as trusted. Only trusted keys are used for SNTP authentication.
- If the SNTP server requires authentication, one of the trusted keys has to be associated with the SNTP server.
- SNTP client authentication must be enabled on the HP Switch. If client authentication is disabled, packets are processed without authentication.

All of the above steps are necessary to enable authentication on the client.

SNTP server authentication support

NOTE: SNTP server is not supported on HP Switch products.

You must perform the following on the SNTP server:

- The same authentication key-identifier, trusted key, authentication mode and key-value that were configured on the SNTP client must also be configured on the SNTP server.
- SNTP server authentication must be enabled on the server.

If any of the parameters on the server are changed, the parameters have to be changed on all the SNTP clients in the network as well. The authentication check fails on the clients otherwise, and the SNTP packets are dropped.

Configuring the key-identifier, authentication mode, and key-value (CLI)

This command configures the `key-id`, `authentication-mode`, and `key-value`, which are required for authentication. It is executed in the global configuration context.

Syntax:

```
sntp authentication key-id <key-id> authentication-mode <md5>  
key-value <key-string> [trusted]  
no sntp authentication key-id <key-id>
```

Configures a `key-id`, `authentication-mode` (MD5 only), and `key-value`, which are required for authentication.

The `no` version of the command deletes the authentication key.

Default: No default keys are configured on the switch.

<code>key-id</code>	A numeric key identifier in the range of 1-4,294,967,295 (2^{32}) that identifies the unique key value. It is sent in the SNTP packet.
---------------------	--

<code>key-value <key-string></code>	The secret key that is used to generate the message digest. Up to 32 characters are allowed for <i>key-string</i> .
<code>encrypted-key <key-string></code>	Set the SNTP authentication key value using a base64–encoded aes-256 encrypted string.

Example 9 Setting parameters for SNTP authentication

```
HP Switch(config)# sntp authentication key-id 55 authentication-mode md5
key-value secretkey1
```

Configuring a trusted key

Trusted keys are used in SNTP authentication. In unicast mode, you must associate a `trusted` key with a specific NTP/SNTP server. That key is used for authenticating the SNTP packet.

In unicast mode, a specific server is configured on the switch so that the SNTP client communicates with the specified server to get the date and time.

In broadcast mode, the SNTP client switch checks the size of the received packet to determine if it is authenticated. If the broadcast packet is authenticated, the `key-id` value is checked to see if the same `key-id` value is configured on the SNTP client switch. If the switch is configured with the same `key-id` value, and the `key-id` value is configured as "trusted," the authentication succeeds. Only trusted `key-id` value information is used for SNTP authentication. For information about configuring these modes, see [“Configuring unicast and broadcast mode for authentication” \(page 28\)](#).

If the packet contains `key-id` value information that is not configured on the SNTP client switch, or if the received packet contains no authentication information, it is discarded. The SNTP client switch expects packets to be authenticated if SNTP authentication is enabled.

When authentication succeeds, the time in the packet is used to update the time on the switch.

Configuring a key-id as trusted (CLI)

Enter the following command to configure a `key-id` as trusted.

Syntax:

```
sntp authentication key-id <key-id> trusted
no sntp authentication key-id <key-id> trusted
```

Trusted keys are used during the authentication process. You can configure the switch with up to eight sets of `key-id`/`key-value` pairs. One specific set must be selected for authentication; this is done by configuring the set as `trusted`.

The `key-id` itself must already be configured on the switch. To enable authentication, at least one `key-id` must be configured as `trusted`.

The `no` version of the command indicates the key is unreliable (not trusted).

Default: No key is trusted by default.

For detailed information about trusted keys, see [“Configuring a trusted key” \(page 27\)](#)

Associating a key with an SNTP server (CLI)

Syntax:

```
[no] sntp server priority <1-3> <ip-address | ipv6-address>
<version-num> [ key-id <1-4,294,967,295> ]
```

Configures a `key-id` to be associated with a specific server. The key itself must already be configured on the switch.

The `no` version of the command disassociates the key from the server. This does not remove the authentication key.

Default: No key is associated with any server by default.

<code>priority</code>	Specifies the order in which the configured servers are polled for getting the time.
<code>version-num</code>	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 - 7.
<code>key-id</code>	Optional command. The key identifier sent in the SNTP packet. This <code>key-id</code> is associated with the SNTP server specified in the command.

Example 10 Associating a `key-id` with a specific server

```
HP Switch(config)# sntp server priority 1 10.10.19.5 2 key-id 55
```

Enabling SNTP client authentication

The `sntp authentication` command enables SNTP client authentication on the switch. If SNTP authentication is not enabled, SNTP packets are not authenticated.

Syntax:

```
[no] sntp authentication
```

Enables the SNTP client authentication.

The `no` version of the command disables authentication.

Default: SNTP client authentication is disabled.

Configuring unicast and broadcast mode for authentication

To enable authentication, you must configure either unicast or broadcast mode. When authentication is enabled, changing the mode from unicast to broadcast or vice versa is not allowed; you must disable authentication and then change the mode.

To set the SNTP mode or change from one mode to the other, enter the appropriate command.

Syntax:

```
sntp unicast
```

```
sntp broadcast
```

Enables SNTP for either broadcast or unicast mode.

Default: SNTP mode is disabled by default. SNTP does not operate even if specified by the CLI `timesync` command or by the menu interface `Time Sync Method` parameter.

Unicast	Directs the switch to poll a specific server periodically for SNTP time synchronization. The default value between each polling request is 720 seconds, but can be configured. At least one manually configured server IP address is required.
---------	--

NOTE: At least one `key-id` must be configured as `trusted`, and it must be associated with one of the SNTP servers. To edit or remove the associated `key-id` information or SNTP server information, SNTP authentication must be disabled.

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 (configurable up to 720 seconds) 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.
-----------	---

Viewing SNTP authentication configuration information (CLI)

The `show sntp` command displays SNTP configuration information, including any SNTP authentication keys that have been configured on the switch.

Example 11 SNTP configuration information

```
HP Switch(config)# show sntp
```

```
SNTP Configuration
```

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

Priority	SNTP Server Address	Protocol	Version	KeyId
1	10.10.10.2	3		55
2	fe80::200:24ff:fec8:4ca8	3		55

Viewing all SNTP authentication keys that have been configured on the switch (CLI)

Enter the `show sntp authentication` command, as shown in [Example 12](#).

Example 12 Show sntp authentication command output

```
HP Switch(config)# show sntp authentication
```

```
SNTP Authentication Information
```

```
SNTP Authentication : Enabled
```

Key-ID	Auth Mode	Trusted
55	MD5	Yes
10	MD5	No

Viewing statistical information for each SNTP server (CLI)

To display the statistical information for each SNTP server, enter the `show sntp statistics` command.

The number of SNTP packets that have failed authentication is displayed for each SNTP server address, as shown in [Example 13](#).

Example 13 SNTP authentication statistical information

```
HP Switch(config)# show sntp statistics
```

```
SNTP Statistics
```

```
Received Packets : 0
```

```
Sent Packets : 3
```

```
Dropped Packets : 0
```

SNTP Server Address	Auth Failed Pkts
-----	-----
10.10.10.1	0
fe80::200:24ff:fec8:4ca8	0

Saving configuration files and the include-credentials command

You can use the `include-credentials` command to store security information in the running-config file. This allows you to upload the file to a TFTP server and then later download the file to the HP switches on which you want to use the same settings. For more information about the `include-credentials` command, see "Configuring Username and Password Security" in the *Access Security Guide* for your switch.

The authentication key values are shown in the output of the `show running-config` and `show config` commands only if the `include-credentials` command was executed.

When SNTP authentication is configured and `include-credentials` has not been executed, the SNTP authentication configuration is not saved.

Example 14 Configuration file with SNTP authentication information

```
HP Switch (config) # show config
```

```
Startup configuration:
```

```
.
```

```
.
```

```
.
```

```
timesync sntp
```

```
sntp broadcast
```

```
sntp 50
```

```
sntp authentication
```

```
sntp server priority 1 10.10.10.2.3 key-id 55
```

```
sntp server priority 2 fe80::200:24ff:fec8:4ca8 4 key-id 55
```

NOTE: SNTP authentication has been enabled and a key-id of 55 has been created.

In this example, the `include-credentials` command has not been executed and is not present in the configuration file. The configuration file is subsequently saved to a TFTP server for later use. The SNTP authentication information is not saved and is not present in the retrieved configuration files, as shown in the following example.

Example 15 Retrieved configuration file when `include credentials` is not configured

```
HP Switch (config) # copy tftp startup-config 10.2.3.44 config1
.
.
Switch reboots ...
.
Startup configuration
.
.
.
timesync sntp
sntp broadcast
sntp 50 sntp server priority 1 10.10.10.2.3
sntp server priority 2 fe80::200:24ff:fec8:4ca8 4
.
.
.
```

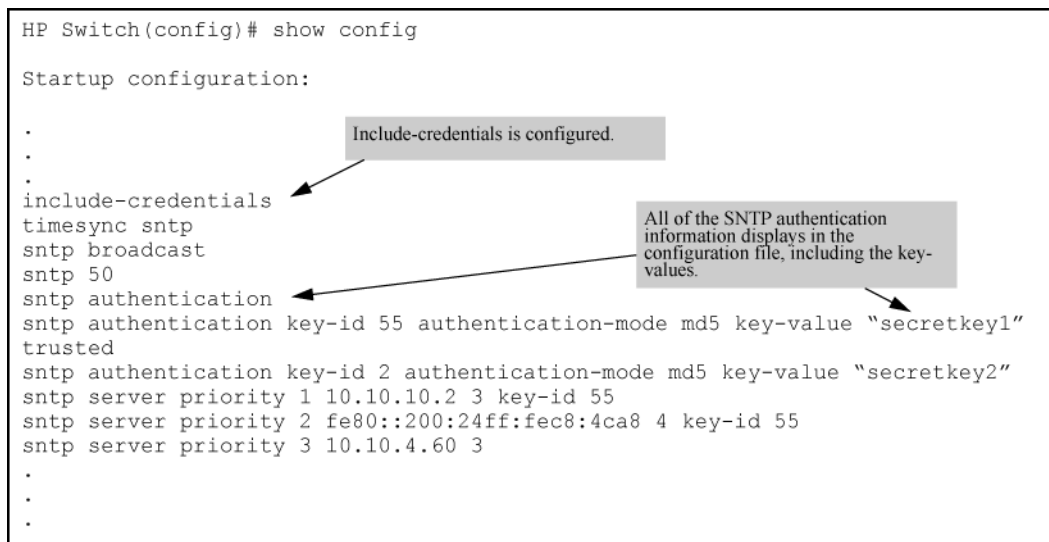
NOTE: The SNTP authentication line and the Key-ids are not displayed. You must reconfigure SNTP authentication.

If `include-credentials` is configured, the SNTP authentication configuration is saved in the configuration file. When the `show config` command is entered, all of the information that has been configured for SNTP authentication displays, including the key-values.

Figure 4 Saved SNTP Authentication information when `include-credentials` is configured

```
HP Switch(config)# show config

Startup configuration:
.
.
.
include-credentials
timesync sntp
sntp broadcast
sntp 50
sntp authentication
sntp authentication key-id 55 authentication-mode md5 key-value "secretkey1"
trusted
sntp authentication key-id 2 authentication-mode md5 key-value "secretkey2"
sntp server priority 1 10.10.10.2 3 key-id 55
sntp server priority 2 fe80::200:24ff:fec8:4ca8 4 key-id 55
sntp server priority 3 10.10.4.60 3
.
.
.
```



TimeP: Selecting and configuring

Table 2 (page 31) shows TimeP parameters and their operations.

Table 2 TimeP parameters

TimeP parameter	Operation
Time Sync Method	Used to select either TIMEP (the default), SNTP, or None as the time synchronization method.
Timep Mode	

Table 2 TimeP parameters (continued)

TimeP parameter	Operation
Disabled	The Default. 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.
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)

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      Server Address :
Tftp-enable [Yes] : Yes                Jumbo Max Frame Size [9216] : 9216
Time Zone [0] : 0                      Jumbo IP MTU [9198] : 9198
Daylight Time Rule [None] : None

Actions->  Cancel  Edit  Save  Help

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

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.

NOTE: 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:

Example 16 TimeP configuration when TimeP is the selected Time synchronization method

```
HP 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:

Example 17 TimeP configuration when TimeP is not the selected time synchronization method

```
HP 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.

Example 18 Display showing IP addressing for all configured time servers and VLANs

```
HP 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.

NOTE: 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):

```
HP Switch(config)# timesync time
```

Selects TimeP.

```
HP Switch(config)# ip timep manual 10.28.227.141
```

Activates TimeP in Manual mode.

Example 19 Configuring TimeP for manual operation

```
HP Switch(config)# timesync timep
```

```
HP Switch(config)# ip timep manual 10.28.227.141
```

```
HP 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)

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:

```
HP 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:

```
HP Switch (config)# no timesync
```

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

Example 20 TimeP with time synchronization disabled

```
HP 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.

Example 21 Disabling time synchronization by disabling the TimeP mode parameter

```
HP Switch(config)# no ip timep

HP 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.

Example 22 How to list all SNTP servers configured on the switch

```
HP 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:

Example 23 Creating additional SNTP server addresses with the CLI

```
HP Switch(config)# sntp server 2001:db8::215:60ff:fe79:8980
HP Switch(config)# sntp server 10.255.5.24
```

NOTE: 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:

```
HP 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.

3 Port Status and Configuration

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)” (page 43)).
- To display information about the transceivers installed on a switch, enter the `show tech receivers` command in the CLI (Example 30 (page 47)).

Table 3 Status and parameters for each port type

Status or parameter	Description
Enabled	Yes (default): The port is ready for a network connection. No: The port will not operate, even if properly connected in a network. Use this setting, for example, if the port needs to be shut down for diagnostic purposes or while you are making topology changes.
Status (read-only)	Up: The port senses a link beat. Down: The port is not enabled, has no cables connected, or is experiencing a network error. For troubleshooting information, see the <i>Installation and Getting Started Guide</i> you received with the switch. See also to Appendix C, "Troubleshooting" (in this manual).
Mode	The port's speed and duplex (data transfer operation) setting. 10/100/1000Base-T Ports: <ul style="list-style-type: none">• Auto-MDIX (default): 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 3. Port Status option under 1. Status and Counters in the menu interface.• 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)• Auto-10: Allows the port to negotiate between half-duplex (HDx) and full-duplex (FDx) while keeping speed at 10 Mbps. Also negotiates flow control (enabled or disabled). HP 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.).• 10HDx: 10 Mbps, half-duplex• 10FDx: 10 Mbps, full-duplex• 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.

Table 3 Status and parameters for each port type (continued)

Status or parameter	Description
	<ul style="list-style-type: none"> 100Hdx: Uses 100 Mbps, half-duplex. 100FDx: Uses 100 Mbps, full-duplex <p>Gigabit Fiber-Optic Ports (Gigabit-SX, Gigabit-LX, and Gigabit-LH):</p> <ul style="list-style-type: none"> 1000FDx: 1000 Mbps (1 Gbps), full-duplex only Auto (default): The port operates at 1000FDx and auto-negotiates flow control with the device connected to the port. <p>Gigabit Copper Ports:</p> <ul style="list-style-type: none"> 1000FDx: 1000 Mbps (1 Gbps), full-duplex only Auto (default): The port operates at 1000FDx and auto-negotiates flow control with the device connected to the port. <p>10-Gigabit CX4 Copper Ports:</p> <ul style="list-style-type: none"> Auto: The port operates at 10 gigabits FDx and negotiates flow control. Lower speed settings or half-duplex are not allowed. <p>10-Gigabit SC Fiber-Optic Ports (10-GbE SR, 10-GbE LR, 10-GbE ER):</p> <ul style="list-style-type: none"> Auto: The port operates at 10 gigabits FDx and negotiates flow control. Lower speed settings or half-duplex are not allowed. <p>NOTE: Conditioning patch cord cables are not supported on 10-GbE.</p>
Auto-MDIX	<p>The switch supports Auto-MDIX on 10Mb, 100Mb, and 1 Gb T/TX (copper) ports. (Fiber ports and 10-gigabit ports do not use this feature.)</p> <ul style="list-style-type: none"> Automdix: Configures the port for automatic detection of the cable type (straight-through or crossover). MDI: Configures the port to connect to a switch, hub, or other MDI-X device with a straight-through cable. MDIX: Configures the port to connect to a PC or other MDI device with a straight-through cable.
Flow control	<ul style="list-style-type: none"> Disabled (default): The port does not generate flow control packets, and drops any flow control packets it receives. Enabled: The port uses 802.3x link layer flow control, generates flow-control packets, and processes received flow-control packets. <p>With the port mode set to Auto (the default) and flow control enabled, the switch negotiates flow control on the indicated port. If the port mode is not set to Auto, or if flow control is disabled on the port, flow control is not used. Note that flow control must be enabled on both ends of a link.</p>
Broadcast limit	<p>Specifies the percentage of the theoretical maximum network bandwidth that can be used for broadcast traffic. Any broadcast traffic exceeding that limit will be dropped. Zero (0) means the feature is disabled.</p> <p>The broadcast-limit command operates at the port context level to set the broadcast limit for a port on the switch.</p> <p>NOTE: This feature is not appropriate for networks that require high levels of IPX or RIP broadcast traffic.</p>

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

Example 24 A switch port status screen

```
===== CONSOLE - MANAGER MODE =====
                        Status and Counters - Port Status

Port      Type      Intrusion
Alert     Enabled Status   Mode     MDI     Flow   Bcast
Mode     Mode     Ctrl     Limit
-----
1         100/1000T No       Yes     Down    100FDx  Auto  off  0
2         100/1000T No       Yes     Down    100FDx  Auto  off  0
3         100/1000T No       Yes     Down    100FDx  Auto  off  0
4         100/1000T No       Yes     Down    100FDx  Auto  off  0
5         100/1000T No       Yes     Down    100FDx  Auto  off  0
6         100/1000T No       Yes     Down    100FDx  Auto  off  0
7         100/1000T No       Yes     Down    100FDx  Auto  off  0
8         100/1000T No       Yes     Down    100FDx  Auto  off  0
9         100/1000T No       Yes     Down    100FDx  Auto  off  0
10        100/1000T No       Yes     Down    100FDx  Auto  off  0
11        100/1000T No       Yes     Down    100FDx  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.
```

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".

1. From the Main Menu, select:

2. Switch Configuration...

2. Port/Trunk Settings

Example 25 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.

- 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:

<code>show interfaces [brief config <port-list>]</code>	
<code>brief</code>	Lists the current operating status for all ports on the switch.
<code>config</code>	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.
<code><port-list></code>	Shows a summary of network traffic handled by the specified ports.

Example 26 The show interfaces brief command listing

```
HP 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

Example 27 The show interfaces config command listing

```
HP 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.

NOTE: Select **Back** to exit the display.

Example

HP 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 **Ctrl-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

Port	Total Bytes	Total Frames	Errors Rx	Drops Tx	Flow Ctrl	Bcast Lim
1	2,164,277	20,366	0	0	off	0
2	0	0	0	0	off	0
3	0	0	0	0	off	0
4	0	0	0	0	off	0
5	0	0	0	0	off	0
6	0	0	0	0	off	0
7	0	0	0	0	off	0
8	0	0	0	0	off	0
9	0	0	0	0	off	0
10	0	0	0	0	off	0
11	0	0	0	0	off	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 [Table 4 \(page 44\)](#).

Table 4 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
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	4 tagged
enabled	port is or is not enabled	yes or no intrusion

Table 4 Supported columns, what they display, and examples *(continued)*

Parameter column	Displays	Examples
intrusion	Intrusion alert status	no
bcast	Broadcast limit	0

Example 28 The custom show interfaces command

```
HP 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 [Example 28](#), the Name column displays only the first four characters of the name. All remaining characters are truncated.

NOTE: 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.

For information on error messages associated with this command and for notes about pattern matching with this command, see [Error messages associated with the show interfaces command](#) (page 45).

Error messages associated with the show interfaces 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. [Example 29](#) shows a sample output from this command.

Example 29 A show interface port-utilization command listing

```
HP 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.

[Example 30](#) shows sample output from the `show tech transceivers` command.

NOTE: Part # column in [Example 30](#) enables you to determine the manufacturer for a specified transceiver and revision number.

Example 30 The show tech transceivers command

```
HP 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-HP switches installed transceiver (see line 23 [Example 30 \(page 47\)](#)), 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)
Check: www.hp.com/rnd/device_help/2_inform for more info.
 - This switch only supports revision B and above transceivers.
Check: www.hp.com/rnd/device_help/2_inform for more info.
 - Self test failure.
 - Transceiver type not supported in this port.
 - Transceiver type not supported in this software version.
 - Not an HP Switch Transceiver.
Go to: www.hp.com/rnd/device_help/2_inform for more info.

Enabling or disabling ports and configuring port mode (CLI)

You can configure one or more of the following port parameters. See [Table 3 \(page 40\)](#) (Broadcast limit (page 41)).

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:

```
HP 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:


```
HP 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:

```
HP Switch(config)# int e c6
HP 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 \(CLI\)”](#) (page 48).

Enabling or disabling flow control (CLI)

NOTE: 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 [Example 26](#) (page 43).) 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. (You can find more information on flow control in [Table 3](#) (page 40).)

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:

Example 31 Configuring flow control for a series of ports

```
HP Switch(config)# int a1-a6 flow-control
```

```
HP Switch(config)# show interfaces brief
```

Status and Counters - Port Status

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

Example 32 Continued from Example 31

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

```
HP Switch(config)# show interfaces brief
```

Status and Counters - Port Status

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

Example 33 Continued from Example 32

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

```
HP Switch(config)# show interfaces brief
```

Status and Counters - Port Status

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

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, see 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

Table 5 (page 51) shows the cabling requirements for the MDI/MDI-X settings.

Table 5 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” \(page 50\)](#).

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`:

Example 34 Displaying the current MDI configuration

```
HP 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

Example 35 Displaying the current MDI operating mode

```
HP Switch(config)# show interfaces brief
```

Status and Counters - Port Status

Port	Type	Intrusion		Status	Mode	MDI Mode	Flow Ctrl	Bcast Limit
		Alert	Enabled					
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 \(CLI\)](#)” (page 55).)
- 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” \(page 53\)](#).

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:

Example 36 Configuring a friendly port name

```
HP Switch(config)# int A3 name
Bill_Smith@10.25.101.73
HP Switch(config)# write mem
HP 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."

Example 37 Configuring one friendly port name on multiple ports

```
HP Switch(config)# int a5-a8 name Draft-Server:Trunk
HP Switch(config)# write mem
HP 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.

Example 38 Friendly port name data for all ports on the switch

```
HP 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
```

Example 39 Friendly port name data for specific ports on the switch

```
HP 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:

Example 40 A friendly port name in a per-port statistics listing

```
HP 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 : not assigned
```

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 [Example 41 “Listing of the startup-config file with a friendly port name configured \(and saved\)”](#) 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`.

Example 41 Listing of the startup-config file with a friendly port name configured (and saved)

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

HP Switch(config)# show config

Startup configuration:
; J9091A Configuration Editor; Created on release xx.15.05.xxxx
hostname "HPSwitch"
interface A0
  name "Print_Server@10.25.101.43"
  exit

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

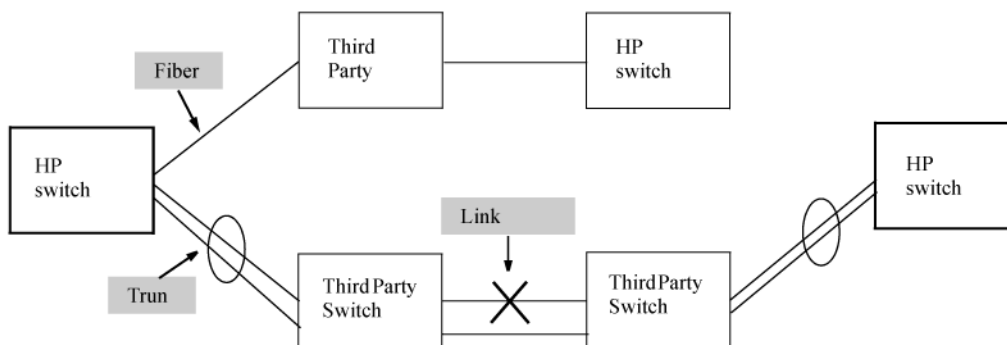
Uni-directional link detection (UDLD)

Uni-directional link detection (UDLD) monitors a link between two HP 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 8 \(page 58\)](#) shows an example.

Figure 8 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 HP 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 HP switches remains undetected. As a result, each switch continues to send traffic on the ports connected to the failed link. When UDLD is enabled on the trunk ports on each HP 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-directional 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 HP 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 \(UDLD\)” \(page 58\)](#).

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:

```
HP Switch(config)#interface a1 link-keepalive
```

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

```
HP Switch(config)#interface a1-a4 link-keepalive
```

NOTE: 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:

```
HP 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:

```
HP 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-HP 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:

```
HP Switch(config)#interface llink-keepalive vlan 22
```

NOTE:

- 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 [Table 3 \(page 40\)](#) 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 9 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 10 Example of show link-keepalive statistics command

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

Port:	1		
Current State:	up	Neighbor MAC Addr:	0000a1-b1c1d1
Ulld Packets Sent:	1000	Neighbor Port:	5
Ulld Packets Received:	1000	State Transitions:	2
Port Blocking:	no	Link-vlan:	1
Port:	2		
Current State:	up	Neighbor MAC Addr:	000102-030405
Ulld Packets Sent:	500	Neighbor Port:	6
Ulld Packets Received:	450	State Transitions:	3
Port Blocking:	no	Link-vlan:	200
Port:	3		
Current State:	off line	Neighbor MAC Addr:	n/a
Ulld Packets Sent:	0	Neighbor Port:	n/a
Ulld Packets Received:	0	State Transitions:	0
Port Blocking:	no	Link-vlan:	1
Port:	4		
Current State:	failure	Neighbor MAC Addr:	n/a
Ulld Packets Sent:	128	Neighbor Port:	n/a
Ulld Packets Received:	50	State Transitions:	8
Port Blocking:	yes	Link-vlan:	1

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

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

Clearing UDLD statistics (CLI)

Enter the following command:

```
HP 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 10 \(page 62\)](#) for an example).

4 Power Over Ethernet (PoE/PoE+) Operation

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 HP Networking website at www.hp.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.

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

NOTE: 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"](#) (page 64).)

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.

NOTE: 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 HP 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.

NOTE: 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 Low 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.

Table 6 shows some examples of PoE priority configuration.

Table 6 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	In this example, the following CLI command sets ports C3 to C17 to Critical : <pre>HP 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	In this example, the following CLI command sets ports C19 to C22 to high : <pre>HP 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	In this example, the CLI command sets ports C23 to C24 to low ² : <pre>HP 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	In this example, the CLI command disables PoE power on ports C1 to C2: <pre>HP Switch(config)# no interface c1-c2 power-over-ethernet</pre> <p>There is no priority setting for the ports in this example.</p>

¹ For a listing of PoE configuration commands with descriptions, see “Configuring PoE operation” (page 65).

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 Table 7 .
<code>value</code>	A user-defined level of PoE power allocated for that port.

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

Table 7 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:

```
HP 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:

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

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

or in interface context:

```
HP Switch(eth-A6) # poe-allocate-by value
```

2. Select a value:

```
HP Switch(config) # int A6 poe-value 15
```

or in interface context:

```
HP Switch(eth-A6) # poe-value 15
```

To view the settings, enter the `show power-over-ethernet` command, shown in [Example 42](#).

Example 42 PoE allocation by value and the maximum power delivered

```
HP Switch(config)# show power-over-ethernet A6
```

```
Status and Counters - Port Power Status for port A7
```

```
Power Enable      : Yes
Priority          : low
AllocateBy       : value
Detection Status  : Delivering
LLDP Detect      : enabled
Configured Type  :
Value            : 15 W 1
Power Class      : 2
Over Current Cnt : 0
Power Denied Cnt : 0
MPS Absent Cnt  : 0
Short Cnt       : 0
Voltage         : 55.1 V
Current        : 154 mA
Power           : 8.4 W
```

1 Maximum power delivered.

If you set the PoE maximum value to less than what the PD requires, a fault occurs, as shown in [Example 43](#).

Example 43 PoE power value set too low for the PD

```
HP Switch(config)# int A7 poe-value 4

HP Switch(config)# show power-over-ethernet A7

Status and Counters - Port Power Status for port A7

Power Enable      : Yes
Priority          : low
AllocateBy       : value
Detection Status  : fault 1
LLDP Detect      : enabled
Configured Type  :
Value           : 4 W
Power Class     : 2

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

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

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

Changing the threshold for generating a power notice

You can configure one of the following thresholds:

- A global power threshold that applies to all ports on the switch. This setting acts as a trigger for sending a notice when the PoE power consumption on any PoE port installed in the switch crosses the configured global threshold level. (Crossing the threshold level in either direction—PoE power usage either increasing or decreasing— triggers the notice.) The default setting is 80%.

Syntax:

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

This command specifies the PoE usage level (as a percentage of the PoE power available) at which the switch generates a power usage notice. This notice appears as an SNMP trap and a corresponding Event Log message, and occurs when the power consumption crosses the configured threshold value. That is, the switch generates a notice whenever the power consumption either exceeds or drops below the specified percentage of the total PoE power available.

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

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.

NOTE: 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:

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

or in interface context:

```
HP Switch(eth-A7) # poe-lldp-detect enabled
```

For more information on PoE/PoE+ and LLDP, see [“PoE/PoE+ allocation using LLDP information” \(page 69\)](#).

Enabling PoE detection via LLDP TLV advertisement

Use this command and insert the desired port or ports:

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

For more information on LLDP, see [“PoE/PoE+ allocation using LLDP information” \(page 69\)](#).

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.

NOTE: 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 [ enabled | disabled ]
```

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 poeplus_config` command).

For the command syntax, see “Initiating advertisement of PoE+ TLVs” (page 71). 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 ]
```

Enables or disables port(s) 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:

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

or in interface context:

```
HP Switch(eth-7) # PoE-lldp-detect enabled
```

NOTE: 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 [Example 44](#).

Example 44 Port with LLDP configuration information obtained from the device

```
HP Switch (config)# show power-over-ethernet brief
```

```
Status and Counters - Port Power Status
```

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

```
Available: 300 W Used: 0 W Remaining: 300 W
```

```
Module A Power
```

```
Available: 300 W Used: 5 W Remaining: 295 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	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 poeplus_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

NOTE: 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.

Example 45 LLDP port configuration information with PoE

```
HP 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:
```

Example 46 shows an example of the local device power information using the `show lldp info local-device <port-list>` command.

Example 46 Local power information

```
HP 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
```

Example 47 shows the remote device power information using the `show lldp info remote-device <port-list>` command.

Example 47 Remote power information

```
HP 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 RA.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 08/04/10 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 08/04/10 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 [ brief | [[ethernet]] <port-list> |  
[ <all> ] ]
```

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

- **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” (page 75) .
<code><port-list></code>	Displays PoE information for the ports in port-list. See “Viewing the PoE status on specific ports” (page 77) .

The `show power-over-ethernet` displays data similar to that shown in [Example 48](#).

Example 48 Output for the `show power-over-ethernet` command

```
HP 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” (page 65) .)
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.

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> <ul style="list-style-type: none"> 0: 0.44 to 12.95 watts can be drawn by the PD. Default class. 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

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

Example 49 Output for the show power-over-ethernet brief command

```
HP 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**:

Example 50 Showing the PoE information by slot

```
HP 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" (page 65) .)
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.
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:

Example 51 Output for the `show power-over-ethernet <port-list>` command

```
HP 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
Power Denied Cnt : 0
MPS Absent Cnt   : 0
Short Cnt        : 0

Voltage           : 0 V
Power             : 0 W
Current           : 0 mA
```

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 *HP PoE/PoE+ Planning and Implementation Guide* which is available on the HP Networking web site at www.hp.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 *HP 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 chapter 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 [Table 8](#).

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 chapter 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 www.hp.com/networking. Auto search the model number for your switch, for example "HP Switch 2920", then select the device from the list and click on **Product manuals**. Click on the "User guide" link under **Manuals**.

5 Port Trunking

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 11 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.

- △ CAUTION:** 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.

NOTE:

Link connections

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

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](#)
- Trunk: Non-Protocol—[Trunk group operation using the "trunk" option](#)

Up to 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.

NOTE: 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, HP 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” \(page 91\)](#).)

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`:

```
HP 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:

```
HP Switch(config)# no int c1-c4 lacp
```

Removes the ports from the trunk.

```
HP 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 9 Trunk types used in static and dynamic trunk groups

Trunking method	LACP	Trunk
Dynamic	Yes	No
Static	Yes	Yes

Table 10 describes the trunking options for LACP and Trunk protocols.

Table 10 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” (page 94).) • 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” (page 91).</p>
Trunk (non-protocol)	<p>Provides manually configured, static-only trunking to:</p> <ul style="list-style-type: none"> • Most HP 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” (page 96).</p>

Table 11 General operating rules for port trunks

Media:	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, HP Switch
---------------	---

Table 11 General operating rules for port trunks *(continued)*

	<p>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. HP Switch 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.</p> <p>Example 52 Recommended port mode setting for LACP</p> <hr/> <pre>HP Switch(config)# show interfaces config</pre> <p>Port Settings</p> <table border="1" data-bbox="580 558 1434 674"> <thead> <tr> <th>Port</th> <th>Type</th> <th>Enabled</th> <th>Mode</th> <th>Flow Ctrl</th> <th>MDI</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>10/100TX</td> <td>Yes</td> <td>Auto</td> <td>Enable</td> <td>Auto</td> </tr> <tr> <td>2</td> <td>10/100TX</td> <td>Yes</td> <td>Auto</td> <td>Enable</td> <td>MDI</td> </tr> </tbody> </table> <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 “Trunk group operation using LACP” (page 91).</p>	Port	Type	Enabled	Mode	Flow Ctrl	MDI	1	10/100TX	Yes	Auto	Enable	Auto	2	10/100TX	Yes	Auto	Enable	MDI
Port	Type	Enabled	Mode	Flow Ctrl	MDI														
1	10/100TX	Yes	Auto	Enable	Auto														
2	10/100TX	Yes	Auto	Enable	MDI														
<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.</p> <p>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 “How the switch lists trunk data” (page 96)</p> <p>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 “Outbound traffic distribution across trunked links” (page 96).</p>																		
<p>Spanning Tree:</p>	<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.</p> <p>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 <code>Trk1</code>—and does not list the individual ports in the trunk.) For example, if ports <code>C1</code> and <code>C2</code> are configured as a static trunk named <code>Trk1</code>, they are listed in the Spanning Tree display as <code>Trk1</code> and do not appear as individual ports in the Spanning Tree displays. See Example 53 (page 85).</p> <p>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>NOTE: A dynamic LACP trunk operates only with the default Spanning Tree settings. Also, this type of trunk appears in the CLI show spanning-tree 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.</p> <p>In the below example, ports <code>C1</code> and <code>C2</code> are members of <code>TRK1</code> and do not appear as individual ports in the port configuration part of the listing.</p>																		

Table 11 General operating rules for port trunks *(continued)*

Example 53 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

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.</p> <p>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>NOTE: For a dynamic LACP trunk to operate in a VLAN other than the default VLAN (DEFAULT_VLAN), GVRP must be enabled. See “Trunk group operation using LACP” (page 91).</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:</p> <pre>< port-list> Command cannot operate over a logical port.</pre>
Monitor port:	<p>NOTE: A trunk cannot be a monitor port. A monitor port can monitor a static trunk but cannot monitor a dynamic LACP trunk.</p>

Viewing and configuring a static trunk group (Menu)

- ① **IMPORTANT:** 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.

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 `E`d*i*t) and then use the arrow keys to access the port trunk parameters.

Figure 12 Example of the menu screen for configuring a 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    -     -
C6    10/100TX | Yes     Auto   Disable    -     -

Actions->  Cancel  Edit  Save  Help

Select Yes to enable the port, No to disable.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.

```

These two columns indicate static trunk status.
(For dynamic LACP trunk status, use the CLI show lacp command—page 12-3.)

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.
 - 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".
 - 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 chapter "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 13 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:
 - LACP
 - 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 figures [Example 54 \(page 87\)](#) and [Example 55](#) 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:

Example 54 Listing specific ports belonging to static trunks

```
HP 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 [Example 55](#), the command does not include a port list, so the switch lists all ports having static trunk membership.

Example 55 A show trunk listing without specifying ports

```
HP 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 13 \(page 93\)](#).)

Example 56 A show LACP listing

```
HP Switch> show lacp
```

Port	LACP Enabled	Trunk Group	LACP		LACP Status	Admin Key	Oper Key
			Port Status	Partner			
A1	Active	Trk1	Up	Yes	Success	0	250
A2	Active	Trk1	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 13 \(page 93\)](#).

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”.

Example 57 A Dynamic LACP trunk with one standby link

```
HP Switch> show lacp
```

Port	LACP Enabled	Trunk Group	LACP		LACP Status	Admin Key	Oper Key
			Port Status	Partner			
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

- ❗ **IMPORTANT:** 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 table on [Table 9](#) 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:

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.

```
HP Switch(config)# trunk c4-c6 trk2 trunk
```

Removing ports from a static trunk group

- ⚠ **CAUTION:** 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, HP 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:

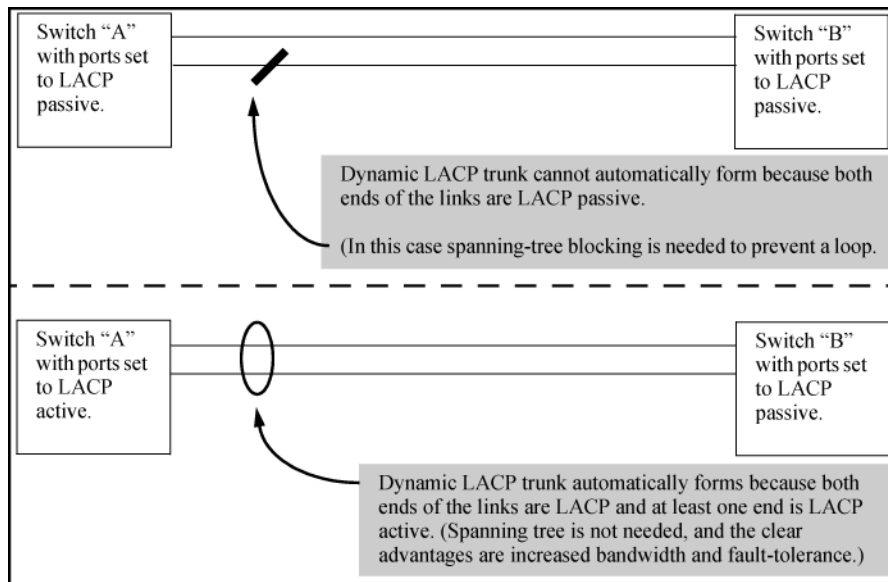
```
HP Switch(config)# no trunk c4-c5
```

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 14 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.

```
HP 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.)

- ⚠ **CAUTION:** 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, HP 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:

```
HP Switch(config)# no interface c6 lacp
HP 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.

NOTE: 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, HP 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 [Table 12 \(page 91\)](#).

Table 12 LACP trunk types

LACP port trunk configuration	Operation
Dynamic LACP	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 60, depending on how many dynamic and static trunks are currently on the switch. (The switch allows a maximum of 60 trunk groups in any combination of static and dynamic trunks.)

Table 12 LACP trunk types *(continued)*

LACP port trunk configuration	Operation
	<p>NOTE: 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” (page 94).</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="582 653 963 768" data-label="Diagram"> <pre> graph LR subgraph Switch1 [Switch 1] direction TB P1[Port X LACP Enable: Active] P2[Port Y LACP Enable: Active] end subgraph Switch2 [Switch 2] direction TB 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>NOTE: 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> <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” (page 94).) • 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. (The table on Table 9 (page 83) lists 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.</p> <p>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. [Table 13 \(page 93\)](#) lists the elements of per-port LACP operation. To display this data for a switch, execute the following command in the CLI:

```
HP Switch> show lacp
```

Table 13 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.</p> <p>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>NOTE: In the default switch configuration, LACP is disabled for all ports.</p>
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 the table on Table 10.)</p> <p>NOTE: 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.

```
HP Switch(config)# aaa port-access authenticator b1
LACP has been disabled on 802.1x port(s).
HP Switch(config)#
```

The switch does not allow you to configure LACP on a port on which port access (802.1X) is enabled. For example:

```
HP Switch(config)# int b1 lacp passive
Error configuring port < port-number > : LACP and 802.1x cannot
be run together.
HP 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:

```
HP Switch(config)# port-security a17 learn-mode static address-
limit 2 LACP has been disabled on secured port(s).
HP Switch(config)#
```

The switch does not allow you to configure LACP on a port on which port security is enabled. For example:

```
HP Switch(config)# int a17 lacp passive
Error configuring port A17: LACP and port security cannot be
run together.
HP 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.

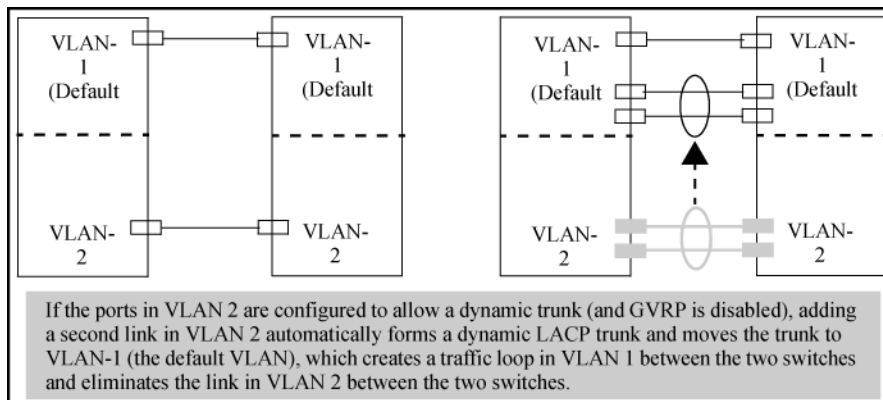
Example 58 Blocked ports with LACP

```
HP 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 15 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 `show trunk` and `show interfaces` commands.

Dynamic LACP trunk group

Appears in the output from the CLI `show lacp` 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

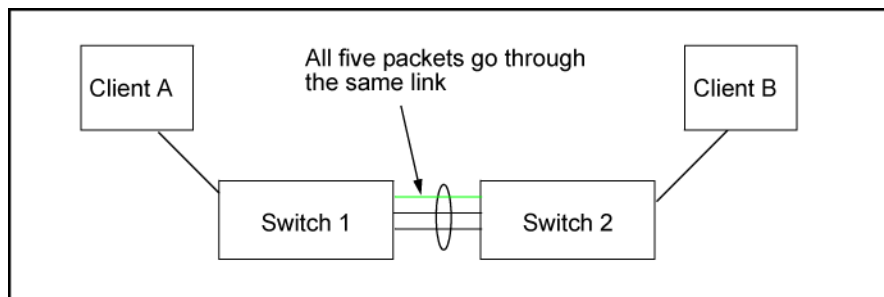
Outbound traffic distribution across trunked links

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 [Figure 16 \(page 97\)](#). 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 16 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 [Figure 17 \(page 97\)](#) showing a three-port trunk, traffic could be assigned as shown in [Table 14 \(page 97\)](#).

Figure 17 Example of port-trunked network




Table 14 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.

Trunk load balancing using port layers

Trunk load balancing using port layers allows the use of TCP/UDP source and destination port number for trunk load balancing. This is in addition to the current use of source and destination IP address and MAC addresses. Configuration of Layer 4 load balancing would apply to all trunks on the switch. Only non-fragmented packets will have their TCP/UDP port number used by load balancing. This ensures that all frames associated with a fragmented IP packet are sent through the same trunk on the same physical link.

The priority for using layer packet information when this feature is enabled is as follows:

1. L4-based: If the packet protocol is an IP packet, use Layer 4, or Layer 3, or Layer 2 information, whichever is present, in that order.
2. L3-based: If the packet protocol is an IP packet, use Layer 3, or Layer 2 information, whichever is present, in that order.
3. L2-based: If the packet protocol is an IP packet use Layer 2 information.
4. For all options, if the packet is not an IP packet, use Layer 2 information.

Enabling trunk load balancing

Enter the following command to enable load balancing.

Syntax:

```
trunk-load-balance <L2-based | L3-based | [L4-based]>
```

This option enables load balancing based on port layer information. The configuration is executed in global configuration context and applies to the entire switch.

Default: L3-based load balancing

L2-based: Load balance based on Layer 2 information.

L3-based: Load balance based on Layer 3 information if present, or Layer 2 information.

L4-based: Load balance on Layer 4 port information if present, or Layer 3 if present, or Layer 2.

Example 59 Enabling L4-based trunk load balancing

```
HP Switch(config)# trunk-load-balance L4 based
```

Example 60 Output when L4-based trunk load balancing is enabled

```
HP Switch(config)# show trunk
```

```
Load Balancing Method: L4-based, L2-based if non-IP traffic
```

Port	Name	Type	Group	Type
41		100/1000T	Trk1	Trunk
42		100/1000T	Trk1	Trunk

Note in [Example 61 “Running config file when L4-based trunk load balancing is enabled”](#) that in if L4 trunk load balancing is enabled, a line appears in the running-config file. If it is not enabled, nothing appears as this is the default and the default values are not displayed.

Example 61 Running config file when L4-based trunk load balancing is enabled

```
HP Switch(config)# show running-config
```

```
Running configuration
```

```
; J9091A Configuration Editor; Created on release #K.15.02.0001x
```

```
hostname "Switch"  
module 1 type J8702A  
module 5 type J9051A  
module 7 type J8705A  
module 10 type J8708A  
module 12 type J8702A  
trunk-load-balance L4-based  
vlan 1  
    name "DEFAULT_VLAN"  
    untagged A1-A24, G1-G24, J1-J4, L1-L24  
    ip address dhcp-bootp  
    tagged EUP  
    no untagged EDP  
    exit  
snmp-server community "public" unrestricted
```

6 Port Traffic Controls

Rate-limiting

- △ **CAUTION:** Rate-limiting is intended for use on edge ports in a network. It is not recommended for use on links to other switches, routers, or servers within a network, or for use in the network core. Doing so can interfere with applications the network requires to function properly.

All traffic rate-limiting

Rate-limiting for all traffic operates on a per-port basis to allow only the specified bandwidth to be used for inbound traffic. When traffic exceeds the configured limit, it is dropped. This effectively sets a usage level on a given port and is a tool for enforcing maximum service level commitments granted to network users. This feature operates on a per-port level and is not configurable on port trunks. Rate-limiting is designed to be applied at the network edge to limit traffic from non-critical users or to enforce service agreements such as those offered by Internet Service Providers (ISPs) to provide only the bandwidth for which a customer has paid.

NOTE: Rate-limiting also can be applied by a RADIUS server during an authentication client session. For further details, see the chapter "RADIUS Authentication and Accounting" in the *Access Security Guide* for your switch.

The switches also support ICMP rate-limiting to mitigate the effects of certain ICMP-based attacks. For more information, see "ICMP rate-limiting" (page 103).

Configuring rate-limiting

Syntax:

```
[no] int <port-list> rate-limit all <in> < percent <0-100>
| kbps <0-10000000>>
```

Configures a traffic rate limit (on non-trunked ports) on the link. The `no` form of the command disables rate-limiting on the specified ports.

The `rate-limit all` command controls the rate of traffic sent or received on a port by setting a limit on the bandwidth available. It includes options for:

- Rate-limiting on inbound traffic.
- Specifying the traffic rate as either a percentage of bandwidth, or in terms of bits per second.

(Default: Disabled.)

<code>in</code>	Specifies a traffic rate limit on inbound traffic passing through that port.
<code>percent or kbps</code>	Specifies the rate limit as a percentage of total available bandwidth, or in kilobits per second.

For more details on configuring rate-limiting, see ???.

Notes:

- The `rate-limit icmp` command specifies a rate limit on inbound ICMP traffic only (see "ICMP Rate-Limiting").
- Rate-limiting does not apply to trunked ports (including meshed ports).
- Kbps rate-limiting is done in segments of 1% of the lowest corresponding media speed. For example, if the media speed is 1 Kbps, the value would be 1 Mbps. A 1-100 Kbps rate-limit

is implemented as a limit of 100 Kbps; a limit of 100-199 Kbps is also implemented as a limit of 100 Kbps, a limit of 200-299 Kbps is implemented as a limit of 200 Kbps, and so on.

- Percentage limits are based on link speed. For example, if a 100 Mbps port negotiates a link at 100 Mbps and the inbound rate-limit is configured at 50%, then the traffic flow through that port is limited to no more than 50 Mbps. Similarly, if the same port negotiates a 10 Mbps link, then it allows no more than 5 Mbps of inbound traffic.

Configuring a rate limit of 0 (zero) on a port blocks all traffic on that port. However, if this is the desired behavior on the port, HP recommends using the `<port-list> disable` command instead of configuring a rate limit of 0.

You can configure a rate limit from either the global configuration level or from the port context level. For example, either of the following commands configures an inbound rate limit of 60% on ports 3 - 5:

```
HP Switch(config)# int 3-5 rate-limit all in percent 60
HP Switch(eth-3-5)# rate-limit all in percent 60
```

Displaying the current rate-limit configuration

The `show rate-limit all` command displays the per-port rate-limit configuration in the running-config file.

Syntax:

```
show rate-limit all [<port-list>]
```

Without [*port-list*], this command lists the rate-limit configuration for all ports on the switch.

With [*port-list*], this command lists the rate-limit configuration for the specified ports. This command operates the same way in any CLI context.

Example 62 Listing the rate-limit configuration

```
HP Switch(config)# show rate-limit all
Inbound Rate Limit Maximum %
```

Port	Limit	Mode	Radius Override
1	Disabled	Disabled	No-override
2	500	kbps	No-override
3	50	%	No-override
4	Disabled	Disabled	No-override

NOTE: To view **RADIUS**-assigned rate-limit information, use one of the following command options:

```
show port-access
  web-based clients <port-list> detailed
  mac-based clients <port-list> detailed
  authenticator clients <port-list> detailed
```

For more on **RADIUS**-assigned rate-limits, see the chapter titled "Configuring RADIUS Server Support for Switch Services" in the latest Management and Configuration Guide for your switch.

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.)

Operating notes for rate-limiting

- **Rate-limiting operates on a per-port basis, regardless of traffic priority.** Rate-limiting is available on all types of ports (other than trunked ports) and at all port speeds configurable for these switches.
- **Rate-limiting is not allowed on trunked ports.** Rate-limiting is not supported on ports configured in a trunk group (including mesh ports). Configuring a port for rate-limiting and then adding it to a trunk suspends rate-limiting on the port while it is in the trunk. Attempting to configure rate-limiting on a port that already belongs to a trunk generates the following message:

```
<port-list>: Operation is not allowed for a trunked port.
```
- **Rate-limiting and hardware.** The hardware will round the actual Kbps rate down to the nearest multiple of 64 Kbps.
- **Operation with other features.** Configuring rate-limiting on a port where other features affect port queue behavior (such as flow control) can result in the port not achieving its configured rate-limiting maximum. For example, in a situation where flow control is configured on a rate-limited port, there can be enough "back pressure" to hold high-priority inbound traffic from the upstream device or application to a rate that is lower than the configured rate limit. In this case, the inbound traffic flow does not reach the configured rate and lower priority traffic is not forwarded into the switch fabric from the rate-limited port. (This behavior is termed "head-of-line blocking" and is a well-known problem with flow-control.)
- **Traffic filters on rate-limited ports.** Configuring a traffic filter on a port does not prevent the switch from including filtered traffic in the bandwidth-use measurement for rate-limiting when it is configured on the same port. For example, ACLs, source-port filters, protocol filters, and multicast filters are all included in bandwidth usage calculations.
- **Monitoring (mirroring) rate-limited interfaces.** If monitoring is configured, packets dropped by 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 bandwidth. This is to ensure the strictest possible rate-limiting of all sizes of packets.

NOTE: Rate-limiting is applied to the available bandwidth on a port and not to any specific applications running through the port. If the total bandwidth requested by all applications is less than the configured maximum rate, then no rate-limit can be applied. This situation occurs with a number of popular throughput-testing applications, as well as most regular network applications. Consider the following example that uses the minimum packet size:

The total available bandwidth on a 100 Mbps port "X" (allowing for Inter-packet Gap—IPG), with no rate-limiting restrictions, is:

$$((100,000,000 \text{ bits}) / 8) / 84) \times 64 = 9,523,809 \text{ bytes per second}$$

where:

- The divisor (84) includes the 12-byte IPG, 8-byte preamble, and 64-bytes of data required to transfer a 64-byte packet on a 100 Mbps link.
- Calculated "bytes-per-second" includes packet headers and data. This value is the maximum "bytes-per-second" that 100 Mbps can support for minimum-sized packets.

Suppose port "X" is configured with a rate limit of 50% (4,761,904 bytes). If a throughput-testing application is the only application using the port and transmits 1 Mbyte of data through the port, it uses only 10.5% of the port's available bandwidth, and the rate-limit of 50% has no effect. This is because the maximum rate permitted (50%) exceeds the test application's bandwidth usage (126,642-164,062 bytes, depending upon packet size, which is only 1.3% to 1.7% of the available total). Before rate-limiting can occur, the test application's bandwidth usage must exceed 50% of the port's total available bandwidth. That is, to test the rate-limit setting, the following must be true:

$$\text{bandwidth usage } (0.50 \times 9,523,809)$$

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 or trunk. 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.

-
- △ CAUTION:** 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.**
-

NOTE: 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.

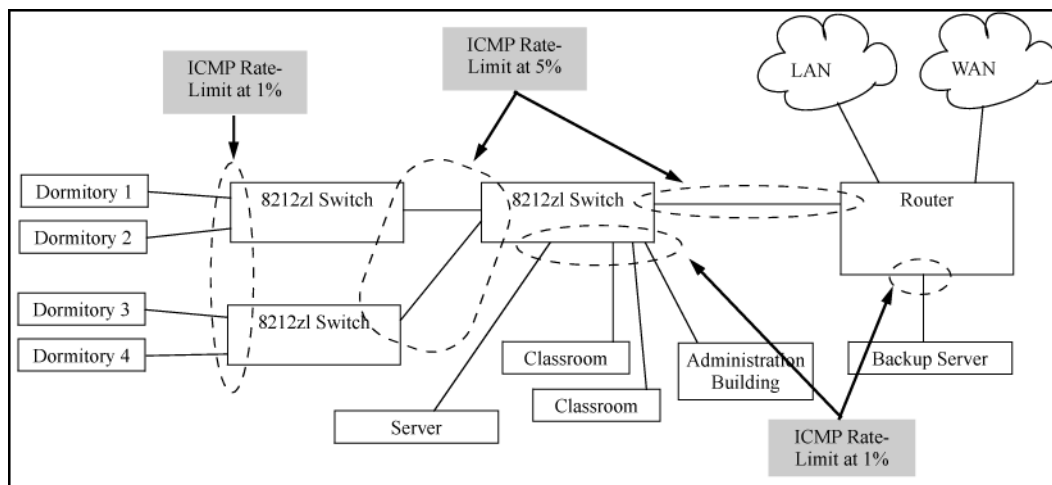
Terminology

All-traffic rate-limiting:	Applies a rate-limit to all traffic (including ICMP traffic) on an interface.
ICMP rate-limiting:	Applies a rate-limit to all inbound ICMP traffic received on an interface, but does not limit other types of inbound traffic.
Spoofed ping:	An ICMP echo request packet intentionally generated with a valid source IP address and an invalid destination IP address. Spoofed pings are often created with the intent to oversubscribe network resources with traffic having invalid destinations.

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 18 \(page 104\)](#) 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 18 Example of ICMP rate-limiting



Configuring ICMP rate-limiting

For detailed information about ICMP rate-limiting, see ["ICMP rate-limiting" \(page 103\)](#).

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 <i>Caution</i> on 103.
<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:

```
HP Switch(config) # int a3-a5 rate-limit icmp 1  
HP Switch(eth-A3-A5) # rate-limit icmp 1
```

NOTE: 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"](#) (page 105).

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.

NOTE: 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 the first six ports in the module in slot "B":

Example 63 Listing the rate-limit configuration

```
HP Switch(config)# show rate-limit icmp b1-b6
```

```
Inbound ICMP Rate Limit Maximum Percentage
```

Port	Mode	Rate Limit
B1	Disabled	Disabled
B2	kpbs	100
B3	%	5
B4	%	1
B5	%	1
B6	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"](#) (page 106).

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 or the `setmib` command.

Syntax:

```
setmib hpIcmpRateLimitPortAlarmflag.<internal-port-#> -i 1
```

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 A1 on a switch would use the following `setmib` 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
or
HP Switch(config)# setmib hpicmpratelimitportalarmflag.
1 -i 1
```

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 [\(page 108\)](#). 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 (slot/number) identity.

To match the port's external slot/number to the internal port number, use the `walkmib ifDescr` command, as shown in the following figure:

Figure 19 Matching internal port numbers to external slot/port numbers

```
HP Switch# walkmib ifDescr
┌ ifDescr.1 = A1 ───┐
│ ifDescr.2 = A2   │
│ ifDescr.3 = A3   │
│ .               │
│ .               │
│ .               │
│ ifDescr.23 = A23 │
│ ifDescr.24 = A24 │
└ ifDescr.27 = B1 ───┘
│ ifDescr.28 = B2  │
│ ifDescr.29 = B3  │
│ .               │
│ .               │
│ .               │
│ ifDescr.48 = B22 │
│ ifDescr.49 = B23 │
│ ifDescr.50 = B24 │
│ .               │
│ .               │
│ .               │
└──────────────────┘
```

← Beginning and Ending of Port Number Listing for Slot

← Beginning and Ending of Port Number Listing for Slot

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:

```
HP Switch(config)# interface 3 rate-limit bcst in percent 10
```

or

```
HP Switch(config)# interface 3
HP Switch(eth-3)# rate-limit bcst in percent 10
```

Syntax:

```
rate-limit <bcst | mcast> in percent <0-100>
```

```
[no] rate-limit <bcst | [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 [Example 64](#). Only 50% of the inbound broadcast traffic will be forwarded.

Example 64 Inbound broadcast rate-limiting of 50% on port 3

```
HP Switch(config)# int 3
HP Switch(eth-3)# rate-limit bcst in percent 50
```

```
HP Switch(eth-3)# show rate-limit bcst
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 [Example 65](#). Only 20% of the multicast traffic will be forwarded.

Example 65 Inbound multicast rate-limiting of 20% on port 3

```
HP Switch(eth-3)# rate-limit mcast in percent 20
HP 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 [Example 66](#).

Example 66 Disabling inbound multicast rate-limiting for port 3

```
HP Switch(eth-3)# no rate-limit mcast in
```

```
HP 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

- This rate-limiting option does not limit unicast traffic.
- This option does not include outbound multicast rate-limiting.

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.

Terminology

Term	Definition
Jumbo Frame	An IP frame exceeding 1522 bytes. The maximum jumbo frame size is 9220 bytes. (This size includes 4 bytes for the VLAN tag.)
Jumbo VLAN	A VLAN configured to allow inbound jumbo traffic. All ports belonging to a jumbo and operating at 1 Gbps or higher can receive jumbo frames from external devices. If the switch is in a meshed domain, all meshed ports (operating at 1 Gbps or higher) on the switch accept jumbo traffic from other devices in the mesh.
MTU (Maximum Transmission Unit)	This is the maximum-size IP frame the switch can receive for Layer 2 frames inbound on a port. The switch allows jumbo frames of up to 9220 bytes.
Standard MTU	An IP frame of 1522 bytes. (This size includes 4 bytes for the VLAN tag.)

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” \(page 113\)](#).

Configuring jumbo frame operation

For detailed information about jumbo frames, see [“Jumbo frames” \(page 110\)](#).

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” (page 113).) See Figure Figure 20.

Figure 20 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 21.)

Figure 21 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 22.)

Figure 22 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)

NOTE: The `jumbo max-frame-size` is set on a GLOBAL level.

Default: 9216 bytes

Configuring IP MTU

NOTE: 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.

```
HP 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” \(page 110\)](#).

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.

Operating notes for jumbo traffic-handling

- HP Switch does not recommend configuring avoice 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 23](#).

Figure 23 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.

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.

7 Configuring for Network Management Applications

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 HP website at: www.hp.com/networking.

From the **Products** menu, select **Network Management**. Then click on **PCM+ Network Management** under the **HP 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)" chapter of the *Advanced Traffic Management Guide* for your switch.

NOTE: 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 HP 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](#)" (page 123))
- 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 HP proprietary MIB (management information base) file. If you are using HP 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 HP Networking website at: www.hp.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:

1. Configure the appropriate SNMP communities. (See “[SNMPv3 communities](#)” (page 123).)
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.)

-
- CAUTION:** 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, HP 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:

1. Enable SNMPv3 for operation on the switch (see “[Enabling SNMPv3](#)” (page 119)).
2. Configure the appropriate SNMP users (see “[SNMPv3 users](#)” (page 120)).
3. Configure the appropriate SNMP communities (see “[SNMPv3 communities](#)” (page 123)).
4. Configure the appropriate trap receivers (see “[SNMP notifications](#)” (page 127)).

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.

-
- CAUTION:** 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).

△ CAUTION: 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 operating in the switch.

Example

Example 67 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
```

The diagram shows three callout boxes with arrows pointing to specific lines in the terminal output:

- Enable SNMPv3**: Points to the `snmpv3 enable` command.
- Create initial user models for SNMPv3 Management Applications**: Points to the `Would you like to create a user that uses SHA? y` prompt.
- Set restriction on non-SNMPv3 messages**: Points to the `Would you like to restrict SHMPv1 and SNMPv2c messages to have read only access` prompt.

SNMPv3 users

NOTE: 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, HP 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:

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” \(page 120\)](#)).
2. Assign users to Security Groups based on their security model with the `snmpv3 group` command (see [“Assigning users to groups \(CLI\)” \(page 122\)](#)).

CAUTION: 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 24 \(page 121\)](#).

Figure 24 Adding SNMPv3 users and displaying SNMPv3 configuration

```
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.
MD5 authentication is enabled and the password is set to "authpass".
Privacy is enabled and the password is set to "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
```

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

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
```

Example 68 “Display of the management stations configured on VLAN 1” displays information about the management stations configured on VLAN 1 to access the switch.

Example 68 Display of the management stations configured on VLAN 1

```
HP Switch# configure terminal
HP Switch(config)# vlan 1
HP 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 25 \(page 122\)](#). For more details on the MIBs access for a given group, see [“Group access levels” \(page 122\)](#).

Figure 25 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” (page 122) .
<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 [Table 6-3 \(page 123\)](#). 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 15 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
commanagerrr	Ver2c or Ver1	ManagerReadView	DiscoveryView
comoperatorrw	Ver2c or Ver1	OperatorReadView	OperatorReadView
comoperatorrr	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.

NOTE: 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 \(CLI\)” \(page 123\)](#)).

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” \(page 123\)](#).

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.

index <index_name>	An index number or title for the mapping. The values of 1 to 5 are reserved and can not be mapped.
name <community_name>	The community name that is being mapped to a group access level.
sec-name <security_name>	The group level to which the community is being mapped.
tag <tag_value>	This is used to specify which target address may have access by way of this index reference.

Example

Figure 26 (page 124) shows the assigning of the Operator community on MgrStation1 to the CommunityOperatorReadWrite group. Any other Operator has an access level of CommunityOperatorReadOnly.

Figure 26 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
```

Two Operator Access Levels

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.


- △ **CAUTION:** 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, HP recommends that you change the write access for the "public" community to "Restricted."

Viewing and configuring non-version-3 SNMP communities (Menu)

1. From the Main Menu, select:
 2. **Switch Configuration...**
 6. **SNMP Community Names**

Figure 27 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.



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” (page 127)).


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 28 Example of the SNMP community listing with two communities



To list the data for only one community, such as the "public" community, use the above command with the community name included. For example:

```
HP 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.

[operator manager]	<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.
[restricted unrestricted]	<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)

```
HP Switch(config)# snmp-server community red-team
manager unrestricted
HP Switch(config)# snmp-server community blue-team
operator restricted
```

To eliminate a previously configured community named "gold-team":

```
HP 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

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” (page 128)
 - “Configuring an SNMP trap receiver (CLI)” (page 129)
 - “Enabling SNMPv2c informs (CLI)” (page 130)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)” (page 131).
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)” (page 135).
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)” (page 136)
 - “Viewing SNMP notification configuration (CLI)” (page 137)

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.

NOTE: 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)

For information about configuring SNMP trap receivers, see ["SNMP trap receivers" \(page 128\)](#).

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 Table 6-2 (page 129)).</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)" (page 130).</p>

Table 16 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:

```
HP 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)” (page 130)) 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” (page 130).

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 inform request if no SNMP response is received. (Default: 3)
<code>timeout</code>	Number of seconds to wait for an acknowledgement before resending the inform request. (Default: 15 seconds)

NOTE: 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 [Example 69 \(page 131\)](#) (note indication of inform Notify Type in bold below):

Example 69 Display of SNMPv2c inform configuration

```
HP 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

Snmp 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.

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](#)" (page 120)). 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 \(CLI\)](#)" (page 122)).
 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</code> <code><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 <parms_name></code>	Name of the SNMPv3 station's parameters file. The parameters filename configured with <code>params parms_name</code> must match the <code>params parms_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 parms_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 snmpv3 targetaddress command in Step 5.</p> <p>If you configure the security model as ver3, you must also configure the message processing value as ver3.</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 ver3 and the security model as ver3, you must also configure a security services level (noauth, auth, or priv).</p>

Example

An example of how to configure SNMPv3 notification is shown here:

Figure 29 Example of an SNMPv3 notification configuration

The diagram shows a configuration session with the following commands:

```
Switch(config)# snmpv3 notify MyNotification tagvalue not_tag
Switch(config)# snmpv3 targetaddress not_addr params not_params 15.255.123.109
                  filter not_info taglist not_tag
Switch(config)# snmpv3 params not_params user NetworkMgr sec-model ver3
                  message-processing ver3 priv
```

Callouts in the diagram:

- Top left: Params _name value in the snmpv3 targetaddress command matches the params _name value in the snmpv3 params command.
- Top right: The tag _name value in snmpv3 notify command matches the tag _name value in the snmpv3 targetaddress command.
- Bottom left: Configuring the security model ver3 requires you to configure message processing ver3 and a security service level.

Network security notifications

By default, a switch is enabled to send the SNMP notifications listed in “Supported Notifications” (page 127) 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)” (page 129)
- “Configuring SNMPv3 notifications (CLI)” (page 131)

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, or 802.1X) 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” (page 133).

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.
dyn-ip-lockdown	If the switch is out of hardware resources needed to program a dynamic IP lockdown rule
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.
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.
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 [Example 70 \(page 135\)](#). Note that command output is a subset of the information displayed with the `show snmp-server` command in [Figure 30 \(page 138\)](#).

Example 70 Display of configured network security notifications

```
HP 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` ((page 136)) and `snmp-server trap-source` ((page 136)) commands.

Configuring the source IP address for SNMP notifications (CLI)

For more information, see “Source IP address for SNMP notifications” (page 135).

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</code> <code>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:

```
HP 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.

NOTE: 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 [Example 71 \(page 137\)](#).

Example 71 Display of source IP address configuration

```
HP 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 30 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
...

```

Configuring the MAC address count option

The MAC Address Count feature provides a way to notify the switch management system when the number of MAC addresses learned on a switch port exceeds the permitted configurable number.

To enable the `mac-count-notify` option, enter this command in global config context.

Syntax

```
[no]snmp-server enable traps mac-count-notify
```

Sends a trap when the number of MAC addresses learned on the specified ports exceeds the configured `<learned-count>` value.

To configure the `mac-count-notify` option on a port or ports, enter this command. When the configured number of MAC addresses is exceeded (the `learned-count`), a trap is sent.

Syntax

```
[no] mac-count-notify traps <port-list> [<learned-count>]
```

Configures `mac-count-notify traps` on the specified ports (or all) for the entire switch.

The `[no]` form of the command disables `mac-count-notify traps`.

`<learned-count>`: The number of MAC addresses learned before sending a trap. Values range between 1-128.

Default: 32

Example 72 Configuring mac-count notify traps on ports 5–7

```
HP Switch (config)# mac-count-notify traps 5-7 50
```

Displaying information about the mac-count-notify option

Use the `show mac-count-notify traps [<port-list>]` command to display information about the configured value for sending a trap, the current count, and if a trap has been sent.

Example 73 Information displayed for the `show mac-count-notify traps` command

```
HP Siwtch (config)# show mac-count-notify traps
```

```
Mac-count-notify Enabled: Yes
```

Port	Count for sending Trap	Count	Trap Sent
-----	-----	-----	-----
1			
2			
3			
4			
5	50	0	No
6	50	2	No
7	50	0	No
8			
9			
...			

The interface context can be used to configure the value for sending a trap.

Example 74 Configuring mac-count-notify traps from the interface context

```
HP Switch (config)# interface 5
```

```
HP Switch (eth-5)# mac-count-notify traps 35
```

The `show snmp-server traps` command displays whether the MAC Address Count feature is enabled or disabled.

Example 75 Information about SNMP traps, including MAC address count being Enabled/Disabled

```
HP Switch(config)# show snmp-server traps
```

```
Trap Receivers
```

```
Link-Change Traps Enabled on Ports [All] : All
```

```
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
```

```
MAC address table changes    : Disabled  
MAC Address Count            : Enabled 1
```

```
Address      Community  Events  Type  Retry  Timeout  
-----  
15.146.194.77 public    None    trap  3      15  
15.255.134.252 public    None    trap  3      15  
16.181.49.167 public    None    trap  3      15  
16.181.51.14 public    None    trap  3      15
```

```
Excluded MIBs
```

1 The notify option is enabled.

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 HP Switch Manager network management software. For more information on PCM+, see the HP Networking web site at www.hp.com/networking.

From the Products menu, select Network Management. Then click on PCM+ Network Management under the HP 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” \(page 140\)](#).

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.

NOTE: 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. [Example 77 \(page 142\)](#) 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.

Example 76 Viewing sflow agent information

```
HP Switch# show sflow agent

Version          1.3;HP;K.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 [Example 77 \(page 142\)](#).

Example 77 Viewing sFlow destination information

```
HP 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 31 \(page 143\)](#). You can specify a list or range of ports for which to view sampling information.

Figure 31 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		Header	Dropped		Polling	
	Enabled	Rate		Samples	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	

NOTE: 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.

LLDP

To standardize device discovery on all HP switches, LLDP will be 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 HP 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.

NOTE: 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:

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 \(media-endpoint-discovery\)” \(page 158\)](#).

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 \(media-endpoint-discovery\)” \(page 158\)](#).

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 \(media-endpoint-discovery\)](#)” (page 158).

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 (see [syntax \(page 149\)](#)).

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 \(media-endpoint-discovery\)](#)” (page 158).

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 (see [syntax \(page 149\)](#)).

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 (see [syntax \(page 150\)](#)).

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 ([Section \(page 152\)](#)). 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 \(page 152\)](#)).

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 ([Section \(page 153\)](#)).

Table 17 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 ^{5, 3}	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 (page 146) ".	
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 17 Data available for basic LLDP advertisements *(continued)*

Data type	Configuration options	Default	Description
System capabilities supported ^{6,3}	Enable/Disable	Enabled	Identifies the switch's primary capabilities (bridge, router).
System capabilities enabled ⁶ ₃	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”](#) (page 147).

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 \(CLI\)”](#) (page 147)).
- 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 \(media-endpoint-discovery\)”](#) (page 158).)

LLDP operating rules

For additional information specific to LLDP-MED operation, see [“LLDP-MED \(media-endpoint-discovery\)”](#) (page 158).

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 \(CLI\)” \(page 153\)](#)). (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 \(CLI\)” \(page 152\)](#).

`show lldp config` produces the following display when the switch is in the default LLDP configuration:

Example 78 Viewing the general LLDP configuration

```
HP 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

NOTE: 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” \(page 152\)](#). For information on the other configurable settings displayed by this command, see [“Configuring per-port transmit and receive modes \(CLI\)” \(page 152\)](#).

Figure 32 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” (page 149).

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)

Example 79 Disabling LLDP

```
HP 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)

NOTE: The `refresh-interval` must be greater than or equal to $(4 \times \text{delay-interval})$. (The default `delay-interval` is 2). For example, with the default `delay-interval`, the lowest `refresh-interval` you can use is 8 seconds ($4 \times 2=8$). Thus, if you want a `refresh-interval` of 5 seconds, you must first change the `delay-interval` to 1 (that is, $4 \times 1 = 4$). 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” (page 150).

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 (4×15).

To reduce the Time-to-Live, you could lower the `holdtime-multiplier` to 2, which would result in a Time-to-Live of 30 seconds.

```
HP 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.

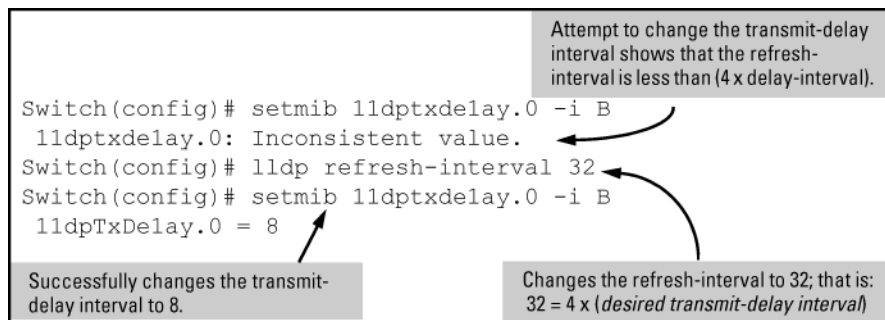
(Default: 2; Range 1–8192)

NOTE: 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 ($32 = 4 \times 8$). (See Figure 33 (page 151).)

Figure 33 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)

For more information, see “Reinitialization delay interval” (page 151).

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:

```
HP 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)

For more information, see Section 1.67.3.2.

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” (page 127).

Example

This command enables SNMP notification on ports 1 - 5:

```
HP 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.

```
HP 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” \(page 153\)](#).

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.)

NOTE: 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:

```
HP 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:

```
HP 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:

```
HP 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” \(page 170\)](#).

Configuring support for port speed and duplex advertisements (CLI)

For more information, see [“Support for port speed and duplex advertisements” \(page 154\)](#).

Syntax:

```
[no] lldp config <port-list> dot3TlvEnable macphy_config
```

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)

NOTE: 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” \(page 155\)](#).

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.

Example

Figure 34 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 [Example 80 \(page 157\)](#) through [Example 82 \(page 158\)](#).

Example 80 Displaying the TLVs for a port

```
HP 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:  
:  
:
```

1 The VLAN ID TLV is being advertised.

Example 81 Local device LLDP information

```
HP 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
```

1 The information that LLDP used in its advertisement.

Example 82 Remote device LLDP information

```
HP 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 K.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 35 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:

- | | | | |
|---------------------|---------------------|---------------------|------------|
| • Hardware revision | • Software revision | • Manufacturer name | • Asset ID |
| • Firmware revision | • Serial number | • Model name | |

- Provide information on network connectivity capabilities (for example, a multi-port VoIP phone with Layer 2 switch capability)
- Support the fast-start capability

NOTE: 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)

NOTE: 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.

Tracking LLDP-MED connects and disconnects—topology change notification

This optional feature provides information an SNMP application can use to track LLDP-MED connects and disconnects. For more information, see “[LLDP-MED \(media-endpoint-discovery\)](#)” (page 158).

Syntax:

```
lldp top-change-notify <port-list>
```

Topology change notification, when enabled on an LLDP port, causes the switch to send an SNMP trap if it detects LLDP-MED endpoint connection or disconnection activity on the port, or an age-out of the LLDP-MED neighbor on the port. The trap includes the following information:

- The port number (internal) on which the activity was detected (For more on internal port numbers, see “[Determining the switch port number included in topology change notification traps](#)” (page 174).)
- The LLDP-MED class of the device detected on the port (“[LLDP-MED endpoint device classes](#)” (page 160).)

The `show running` command shows whether the topology change notification feature is enabled or disabled. For example, if ports A1 to A10 have topology change notification enabled, the following entry appears in the `show running` output:

```
lldp top-change-notify A1-A10
(Default: Disabled)
```

NOTE: To send traps, this feature requires access to at least one SNMP server. For information on configuring traps, see [“SNMP notifications” \(page 127\)](#). Also, if a detected LLDP-MED neighbor begins sending advertisements without LLDP-MED TLVs, the switch sends a top-change-notify trap.

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)

NOTE: 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 \(page 164\)](#))

NOTE: 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 \(CLI\)” \(page 155\)](#)).

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>
```

NOTE: 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 chapter "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"](#) (page 161).

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”](#) (page 164).)
- 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.

capabilities	<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.</p> <p>(Default: enabled)</p> <p>NOTE: This TLV cannot be disabled unless the <code>network_policy</code>, <code>poE</code>, and <code>location_id</code> TLVs are already disabled.</p>
network-policy	<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.</p> <p>(Default: Enabled)</p> <p>NOTE: 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” (page 162).</p>
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” (page 164).</p> <p>(Default: Enabled)</p> <p>NOTE: When disabled, this TLV cannot be enabled unless the <code>capability</code> 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.</p> <p>(Default: Enabled)</p> <p>NOTE: 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” (page 163).</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 Chapter 11, "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

For more information, see "Location data for LLDP-MED devices" (page 164).

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.

NOTE: 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 FR (France), DE (Germany), and IN (India). This field is required in a <code>civic-addr</code>
-------------	--

	<p>command. (For a complete list of country codes, visit www.iso.org.)</p>
WHAT	<p>A single-digit number specifying the type of device to which the location data applies:</p> <ul style="list-style-type: none"> 0: Location of DHCP server 1: Location of switch 2: Location of LLDP-MED endpoint (recommended application) <p>This field is required in a <code>civic-addr</code> command.</p>
Type/Value Pairs (CA-TYPE and CA-VALUE)	<p>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 (<code>CA-TYPE</code>), and the second value in a pair is expected to be the corresponding civic address data (<code>CA-VALUE</code>).</p> <p>For example, if the <code>CA-TYPE</code> for "city name" is "3," the type/value pair to define the city of Paris is "3 Paris."</p> <p>Multiple type/value pairs can be entered in any order, although HP recommends that multiple pairs be entered in ascending order of the <code>CA-TYPE</code>.</p> <p>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.</p> <p>A <code>civic-addr</code> 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.</p> <p><code>CA-TYPE</code>: 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 (<code>CA-VALUE</code>). Some examples of <code>CA-TYPE</code> specifiers include:</p> <ul style="list-style-type: none"> • 3=city • 6=street (name) • 25=building name <p>(Range: 0 - 255)</p> <p>For a sample listing of <code>CA-TYPE</code> specifiers, see Table 6-5 (page 166).</p> <p><code>CA-VALUE</code>: This is the second entry in a type/value pair and is an alphanumeric string containing the location information corresponding to the immediately preceding <code>CA-TYPE</code> entry.</p> <p>Strings are delimited by either blank spaces, single quotes (' ... '), or double quotes ("... ").</p> <p>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 <code>CA-TYPE</code> number identifying the type of data in the string.</p>

	<p>NOTE: A switch port allows one instance of any given CA-TYPE. For example, if a type/value pair of 6 Atlantic (to 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.</p>
elin-addr <emergency-number>	<p>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.</p> <p>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.</p> <p>(Range: 1-15 numeric characters)</p>

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*.

NOTE: Endpoint use of data from a medPortLocation TLV sent by the switch is device-dependent. See the documentation provided with the endpoint device.

Table 18 Some location codes used in CA-TYPE fields¹

Location element	Code	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
city	3	Widgitville
street	6	Main
street number	19	1433
unit	26	Suite 4-N

floor	27	4
room number	28	N4-3

Example 83 shows the commands for configuring and displaying the above data.

Example 83 A civic address configuration

```
HP 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
```

```
HP 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

NOTE: 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 [Example 84 \(page 169\)](#).

Example 84 Displaying the global and per-port information available for outbound advertisements

```
HP 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 3500yl, revision K.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).
-

Example 85 The default per-port information content for ports 1 and 2

```
HP 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]` (Example 47) 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” (page 170).

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 Table 17 (page 145).

Example 86 A global listing of discovered devices

```
HP 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

Example 87 An LLLDP-MED listing of an advertisement received from an LLLDP-MED (VoIP telephone) source

```
HP 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
```

- 1** Indicates 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" (page 171).

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 " Neighbor maximum " (page 173).
Neighbor Entries AgeOut Count	The number of LLDP neighbors dropped on all ports because of Time-to-Live expiring.

Per-Port LLDP Counters:

NumFramesRecvd	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 " Neighbor maximum " (page 173). 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

Example 88 A global LLDP statistics display

```
HP 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	NumFramesRecv	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

Example 89 A per-port LLDP statistics display

```
HP 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` (see [syntax \(page 153\)](#)) 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 \(CLI\)” \(page 150\)](#).

Mandatory TLVs

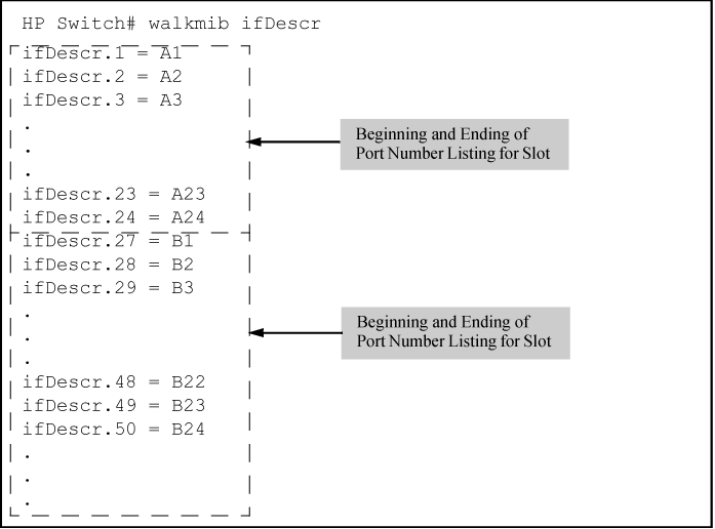
All mandatory TLVs required for LLDP operation are also mandatory for LLDP-MED operation.

Determining the switch port number included in topology change notification traps

Enabling topology change notification on a switch port and then connecting or disconnecting an LLDP-MED endpoint on that port causes the switch to send an SNMP trap to notify the designated management stations. The port number included in the trap corresponds to the internal number the switch maintains for the designated port, and not the port's external (slot/number) identity. To match the port's external slot/number to the internal port number appearing in an SNMP trap, use the `walkmib ifDescr` command, as shown in [Figure 36 \(page 174\)](#).

Figure 36 Matching internal port numbers to external slot/port numbers

```
HP Switch# walkmib ifDescr
| ifDescr.1 = A1 |
| ifDescr.2 = A2 |
| ifDescr.3 = A3 |
| . |
| . |
| ifDescr.23 = A23 |
| ifDescr.24 = A24 |
| ifDescr.27 = B1 |
| ifDescr.28 = B2 |
| ifDescr.29 = B3 |
| . |
| . |
| ifDescr.48 = B22 |
| ifDescr.49 = B23 |
| ifDescr.50 = B24 |
| . |
| . |
| . |
```



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. (HP 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.

NOTE: 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.

NOTE: Because HP 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.

Protocol state	Packet generation	Inbound data management	Inbound packet forwarding
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.

NOTE: 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.

Example 90 “Default CDP configuration” shows the default CDP configuration.

Example 90 Default CDP configuration

```
HP 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.

[[e] port-numb [detail]]	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.
[detail [[e] port-numb]]	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.

Example 91 “CDP neighbors table listing” displays the CDP devices that the switch has detected by receiving their CDP packets.

Example 91 CDP neighbors table listing

```
HP Switch(config)# show cdp neighbors

CDP neighbors information

Port Device ID | Platform | Capability
-----+-----+-----
1 Accounting (0030c1-7fcc40) | J4812A HP Switch. . . | S
2 Researç1-1 (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:

```
HP 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:

```
HP Switch(config)# no cdp enable a1
```

Configuring CDPv2 for voice transmission

Legacy Cisco VOIP phones only support manual configuration or using CDPv2 for voice VLAN auto-configuration. LLDP-MED is not supported. CDPv2 exchanges information such as software version, device capabilities, and voice VLAN information between directly connected devices such as a VOIP phone and a switch.

When the Cisco VOIP phone boots up (or sometimes periodically), it queries the switch and advertises information about itself using CDPv2. The switch receives the VOIP VLAN Query TLV (type 0x0f) from the phone and then immediately sends the voice VLAN ID in a reply packet to the phone using the VLAN Reply TLV (type 0x0e). The phone then begins tagging all packets with the advertised voice VLAN ID.

NOTE: A voice VLAN must be configured before the voice VLAN can be advertised. For example, to configure VLAN 10 as a voice VLAN tagged for ports 1 through 10, enter these commands:

```
HP Switch(config)# vlan 10
HP Switch(vlan-10)# tagged 1-10
HP Switch(vlan-10)# voice
HP Switch(vlan-10)# exit
```

The switch CDP packet includes these TLVs:

- CDP Version: 2
- CDP TTL: 180 seconds
- Checksum

- Capabilities (type 0x04): 0x0008 (is a switch)
- Native VLAN: The PVID of the port
- VOIP VLAN Reply (type 0xe): voice VLAN ID (same as advertised by LLDPMED)
- Trust Bitmap (type 0x12): 0x00
- Untrusted port COS (type 0x13): 0x00

CDP should be enabled and running on the interfaces to which the phones are connected. Use the `cdp enable` and `cdp run` commands.

The `pre-standard-voice` option for the `cdp mode` command allows the configuration of CDP mode so that it responds to received CDP queries from a VoIP phone.

Syntax

```
[no] cdp mode pre-standard-voice [admin-status <port-list>
[ {tx_rx} | {rxonly} ]]
```

Enable CDP-compatible voice VLAN discovery with pre-standard VoIP phones. In this mode, when a CDP VoIP VLAN query is received on a port from pre-standard phones, the switch replies back with a CDP packet that contains the VID of the voice VLAN associated with that port.

NOTE: Not recommended for phones that support LLDP-MED.

pre-standard-voice	Enables CDP-compatible voice VLAN discovery with pre-standard VoIP phones.
admin-status	Sets the port in either transmit and receive mode, or receive mode only. Default: tx-rx. <port-list> Sets this port in transmit and receive mode, or receive mode only. rxonly Enable receive-only mode of CDP processing. tx_rx Enable transmit and receive mode.

```
HP Switch(config)# cdp mode pre-standard-voice admin-status A5 rxonly
```

Example 92 The show cdp output when CDP Run is disabled

```
HP Switch (config)# show cdp
Global CDP information
Enable CDP [yes] : no
```

Example 93 The show cdp output when cdp run and sdp mode are enabled

```
HP Switch(config)# show cdp

Global CDP Information

Enable CDP [Yes] : Yes
CDP mode [rxonly] : pre-standard-voice
CDP Hold Time [180] : 180
CDP Transmit Interval [60] : 60

Port CDP            admin-status
-----
A1    enabled        rxonly
A2    enabled        tx_rx
A3    enabled        tx_rx
```

When CDP mode is not `pre-standard voice`, the `admin-status` column is not displayed.

Example 94 The show cdp output when cdp run and cdp mode rxonly are enabled

```
HP Switch(config)# show cdp

Global CDP Information

Enable CDP [Yes] : Yes
CDP mode [rxonly] : rxonly

Port CDP
-----
A1    enabled
A2    enabled
A3    enabled
```

Example 95 The show running-config when admin-status is configured

```
HP Switch(config)# show running-config

Running configuration:

; J9477A Configuration Editor; Created on release #K.16.09.0000x
; Ver #03:01:1f:ef:f2
hostname "HPSwitch"
module 1 type J9307A
cdp mode pre-standard-voice admin-status A5 RxOnly
```

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)

Example 96 Configuring the switch to ignore packet MAC address learns for an untagged VLAN

```
HP Switch(config) ignore-untagged-mac 1-2
```

Displaying the configuration

Enter the show running-config command to display information about the configuration.

Example 97 Configuration showing interfaces to ignore packet MAC address learns

```
HP Switch(config) show running-config
```

```
Running configuration:
```

```
; J9627 Configuration Editor; Created on release K.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

```
HP Switch(config)# logging filter [filter-name] [sub filter id] <regular-expression> deny
```

Regular-expression The regular expression should match the message which is to be filtered.

Syntax

```
HP Switch(config)# logging filter [filter-name] enable
```

8 Link Aggregation Control Protocol—Multi-Active Detection (LACP-MAD)

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.

```
HP-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
HP-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
HP-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.

```
HP-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 Intelligent Resilient Framework (IRF) 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 IRF virtual device. The active ID is identical to the member ID of the master and is thus unique to the IRF virtual device. When LACP MAD detection is enabled, the members exchange their active IDs by sending extended LACPDUs.

- When the IRF 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 IRF stack, the active IDs in the extended LACPDUs sent by the members in different IRF virtual devices are different, indicating that there are multi-active collisions.

LACP-MAD passthrough helps IRF-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.

A File transfers

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.

For general information about downloading software, see the section starting with “[Downloading switch software](#)” (page 183).

Downloading switch software

HP Switch periodically provides switch software updates through the HP Switch Networking website. For more information, see the support and warranty booklet shipped with the switch, or visit <http://www.hp.com/networking> and click on **software updates**.

NOTE: 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.

NOTE: 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](#)” (page 198).

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 HP Switch Networking website at <http://www.hp.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).

NOTE: 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.

1. In the console Main Menu, select **Download OS** to display the screen in [Figure 37 \(page 184\)](#). (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->  _Cancel      _Edit      eXecute      _Help

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 screen shown in [Figure 38 \(page 184\)](#) 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.

NOTE: 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* for your switch.

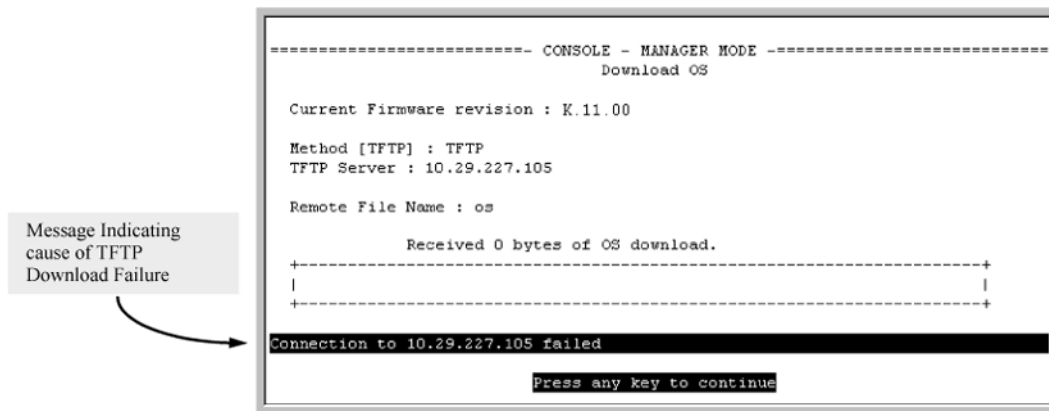
8. After you reboot the switch, confirm that the software downloaded correctly:
 - a. From the Main Menu, select
 2. **Switch Configuration...**
 2. **Port/Trunk Settings**
 - b. Check the **Firmware revision** line.

For troubleshooting information on download failures, see "[Troubleshooting TFTP download failures](#)" (page 185).

Troubleshooting TFTP download failures

When using the menu interface, if a TFTP download fails, the Download OS (Operating System, or software) screen indicates the failure (see [Figure 39](#) (page 185)).

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 HP Switch website. (See "Getting Documentation From the Web".)

NOTE: 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:

1. Execute `copy` as shown below:

Example 98 The command to download an OS (switch software)

```
HP Switch# copy tftp flash 10.28.227.103 k0800.swi
The primary OS Image will be deleted, continue [y/n]? y 1
01431K 2
```

1 This message means that the image you want to upload will replace the image currently in primary flash.

2 Dynamic 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.

NOTE: 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.

Enabling TFTP (CLI)

TFTP is enabled by default on the switch. If TFTP operation has been disabled, you can re-enable it by specifying TFTP client or server functionality with the `tftp [client|server]` command at the global configuration level.

Syntax:

```
[no] tftp [ client | server ]
```

Disables/re-enables TFTP for client or server functionality so that the switch can:

- Use TFTP client functionality to access TFTP servers in the network to receive downloaded files.
- Use TFTP server functionality to upload files to other devices on the network.

Usage notes:

To disable all TFTP client or server operation on the switch except for the auto-TFTP feature, enter the `no tftp [client|server]` command.

When IP SSH file transfer is used to enable SCP and SFTP functionality on the switch, this disables TFTP client and server functionality. Once `ip ssh file transfer` is enabled, TFTP and auto-TFTP cannot be re-enabled from the CLI.

When TFTP is disabled, instances of TFTP in the CLI `copy` command and the Menu interface "Download OS" screen become unavailable.

The `no tftp [client|server]` command does not disable auto-TFTP operation. To disable an auto-TFTP command configured on the switch, use the `no auto-tftp` command described on page ["Configuring the switch to download software automatically from a TFTP server using auto-TFTP \(CLI\)" \(page 187\)](#) to remove the command entry from the switch's configuration.

For information on how to configure TFTP file transfers on an IPv6 network, see the "IPv6 Management Features" chapter in the *IPv6 Configuration Guide* for your switch.

Configuring the switch to download software automatically from a TFTP server using auto-TFTP (CLI)

The `auto-tftp` command lets you configure the switch to download software automatically from a TFTP server.

At switch startup, the auto-TFTP feature automatically downloads a specified software image to the switch from a specified TFTP server and then reboots the switch. To implement the process, you must first reboot the switch using one of the following methods:

- Enter the `boot system flash primary` command in the CLI.
- With the default flash boot image set to primary flash (the default), enter the `boot` or the `reload` command, or cycle the power to the switch. (To reset the boot image to primary flash, use `boot set-default flash primary`.)

Syntax:

```
auto-tftp <ip-addr> <filename>
```

By default, auto-TFTP is disabled. This command configures the switch to automatically download the specified software file from the TFTP server at the specified IP address. The file is downloaded into primary flash memory at switch startup; the switch then automatically reboots from primary flash.

NOTE: To enable auto-TFTP to copy a software image to primary flash memory, the version number of the downloaded software file (for example, K_14_01.swi) must be different from the version number currently in the primary flash image.

The current TFTP client status (enabled or disabled) does not affect auto-TFTP operation. (See “Enabling TFTP (CLI)” (page 187).)

Completion of the auto-TFTP process may require several minutes while the switch executes the TFTP transfer to primary flash and then reboots again.

The `no` form of the command disables auto-TFTP operation by deleting the `auto-tftp` entry from the startup configuration.

The `no auto-tftp` command does not affect the current TFTP-enabled configuration on the switch. However, entering the `ip ssh filetransfer` command automatically disables both `auto-tftp` and `tftp` operation.

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 HP 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).

NOTE: 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 `r`cp (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:

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” \(page 188\)](#).

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 chapter "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:

```
HP Switch(config)# ip ssh filetransfer
```

For information on disabling TFTP and auto-TFTP, see [“Disabling TFTP and auto-TFTP for enhanced security” \(page 189\)](#).

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 [Example 99 \(page 190\)](#).

Example 99 Switch configuration with SFTP enabled

```
HP Switch(config)# ip ssh filetransfer
Tftp and auto-tftp have been disabled. 1
HP 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
```

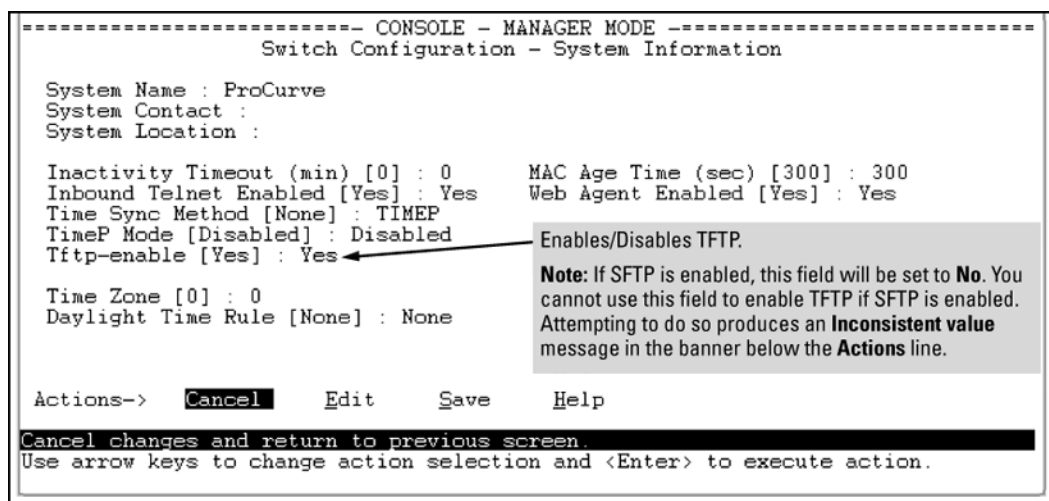
- 1 Enabling SFTP automatically disables TFTP and auto-tftp and displays this message. 2 Viewing the configuration shows that SFTP is enabled and TFTP is disabled.

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 \(page 190\)](#)), 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



- 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.

Enabling SSH V2 (required for SFTP)

```
HP Switch(config)# ip ssh version 2
```

NOTE: 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 HP Switch Series 2500 switches).

Confirming that SSH is enabled

```
HP 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.

NOTE: 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

```
HP 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.

NOTE: 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

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.

NOTE: 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
```

NOTE: 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)

NOTE: The menu interface accesses only the primary flash.

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:

1. Execute the following command in the CLI:

```
HP 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:

```
HP 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.

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:
 - To download the software in the primary flash of the source switch, enter `flash` in lowercase characters.
 - 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:

Example 100 Switch-to-switch, from primary in source to either flash in destination

```
HP Switch# copy tftp flash 10.29.227.13 flash
Device will be rebooted, do you want to continue [y/n]? y
00107K █
```

1 Running 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:

Example 101 Switch-to-switch, from either flash in source to either flash in destination

```
HP 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
```

Copying software images

NOTE: 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:

```
HP 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:

1. Execute the following command:
HP 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 \(CLI\)](#)” (page 198), 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.

NOTE: For greater security, you can perform all TFTP operations using SFTP, as described in the section “[Using SCP and SFTP](#)” (page 188).

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:

```
ProCurve# 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:

```
HP 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 [Example 102](#)).

Example 102 Using the `copy tftp show-tech` command to upload a customized command file

```
HP 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.

NOTE: 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.)

Example 103 The `show tech custom` command

```
HP 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:

```
HP 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:

1. Execute the following command:

```
HP 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.)

Transferring ACL command files

This section describes how to upload and execute a command file to the switch for configuring or replacing an ACL in the switch configuration. Such files should contain only access control entry (ACE) commands. For more on this general topic, including an example of an ACL command file created offline, see the section "Editing ACLs and Creating an ACL Offline" in the "Access Control Lists (ACLs)" chapter of the latest *Access Security Guide* for your switch.

TFTP: Uploading an ACL command file from a TFTP server (CLI)

Syntax:

```
copy tftp command-file <ip-addr> <filename.txt> <unix | pc>
```

Copies and executes the named text file from the specified TFTP server address and executes the ACL commands in the file.

<ip-addr>	The IP address of a TFTP server available to the switch
<filename.txt>	A text file containing ACL commands and stored in the TFTP directory of the server identified by <i>ip-addr</i>
<unix pc>	The type of workstation used for serial, Telnet, or SSH access to the switch CLI

Depending on the ACL commands used, this action does one of the following in the `running-config` file:

- Creates a new ACL.
- Replaces an existing ACL. (See "Creating an ACL Offline" in the "Access Control Lists (ACLs)" chapter in the latest Access Security Guide for your switch.)
- Adds to an existing ACL.

Example

Suppose you:

1. Created an ACL command file named `vlan10_in.txt` to update an existing ACL.
2. Copied the file to a TFTP server at 18.38.124.16.

Using a PC workstation, you then execute the following from the CLI to upload the file to the switch and implement the ACL commands it contains:

```
HP Switch(config)# copy tftp command-file 18.38.124.16  
vlan10_in.txt pc
```

The switch displays this message:

```
Running configuration may change, do you want to continue  
[y/n]?
```

To continue with the upload, press the **[Y]** key. To abort the upload, press the **[N]** key. Note that if the switch detects an illegal (non-ACL) command in the file, it bypasses the illegal command, displays a notice (as shown in [Example 104 \(page 202\)](#)), and continues to implement the remaining ACL commands in the file.

Example 104 Using the `copy` command to download and configure an ACL

```
HP Switch(config)# copy tftp command-file 10.38.124.18 vlan10_in.txt
pc
Running configuration may change, do you want to continue [y/n]?
y
  1. ip access-list extended "155"
  2. deny tcp 0.0.0.0 255.255.255.255 10.10.10.2 0.0.0.0 eq 23 log

  3. permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
  4. show running
Command files are limited to access-list commands. 1
  5. exit
Switch(config)# show running 2
Running configuration:

; J9091A Configuration Editor; Created on release #W.15.05.0000x
; Ver #01:01:00

hostname "HP Switch"
cdp run
ip default-gateway 10.38.248.1
logging 10.38.227.2
snmp-server community "public" unrestricted
ip access-list extended "155"
deny tcp 0.0.0.0 255.255.255.255 10.10.10.2 0.0.0.0 eq 23 log
permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
exit
```

1 This message indicates that the `show running` command just above it is not an ACL command and will be ignored by the switch.

2 Manually executing the `show running` from the CLI indicates that the file was ~~implemented~~ creating ACL 155 in the switch's running ~~configuration~~.

Xmodem: Uploading an ACL command file from a serially connected PC or UNIX workstation (CLI)

Syntax:

```
copy xmodem command-file <unix | pc>
```

Uses Xmodem to copy and execute an ACL command from a PC or UNIX workstation. Depending on the ACL commands used, this action does one of the following in the running-config file:

- Creates a new ACL.
- Replaces an existing ACL. (See "Creating an ACL Offline" in the "Access Control Lists (ACLs)" chapter in the latest *Access Security Guide* for your switch.)
- Adds to an existing ACL.

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:

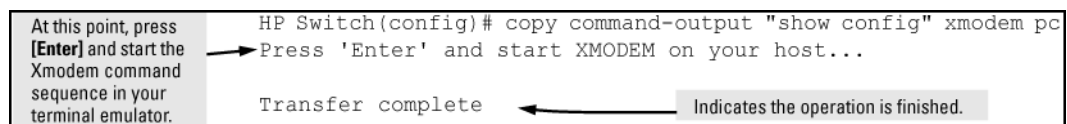
```
copy command-output <"cli-command"> tftp <ip-address> <filepath-filename>
copy command-output <"cli-command"> xmodem
```

These commands direct the displayed output of a CLI command to a remote host, or to a serially connected PC or UNIX workstation.

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



NOTE: The command you specify must be enclosed in double quotation marks.

Copying Event Log output to a destination device (CLI)

Syntax:

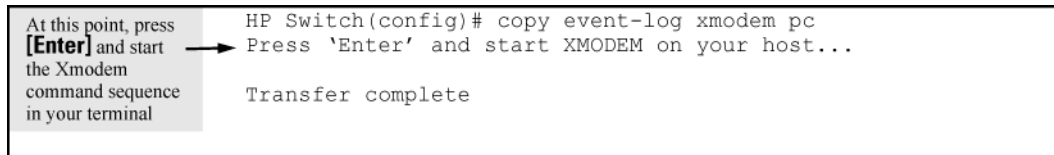
```
copy event-log <tftp | sftp | xmodem>
copy event-log tftp <ip-address> <filepath_filename>
copy event-log xmodem <filename>
```

These commands copy the Event Log content to a remote host, or to a serially connected PC or UNIX workstation.

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 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 or to a serially connected PC or UNIX workstation.

To copy the switch's crash data to a file in a PC:

Example 105 Copying switch crash data content to a PC

```
Switch(config)# copy crash-data xmodem pc
Press 'Enter' and start XMODEM on your host... 1
```

Transfer complete

- 1 At this point press [Enter] and start the Xmodem command sequence in your terminal emulator.

B Monitoring and Analyzing Switch Operation

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 (“[Status and counters data](#)” (page 205)).
- **Counters:** Display details of traffic volume on individual ports (“[Accessing port and trunk statistics \(Menu\)](#)” (page 210)).
- **Event Log:** Lists switch operating events (“[Using the Event Log for troubleshooting switch problems](#)” (page 244)).
- **Configurable trap receivers:** Uses SNMP to enable management stations on your network to receive SNMP traps from the switch.

NOTE: Link test and ping test—analysis tools in troubleshooting situations—are described in Appendix C, “[Troubleshooting](#)” (page 220). See “[Diagnostic tools](#)” (page 273).

Status and counters data

This section describes the status and counters screens available through the switch console interface and/or the WebAgent.

NOTE: You can access all console screens from the WebAgent via Telnet to the console. Telnet access to the switch is available in the **Device View** window under the **Configuration** tab.

Accessing status and counters (Menu)

Beginning at the Main Menu, display the Status and Counters menu by selecting:

1. **Status and Counters**

Figure 43 The Status and Counters menu

```
----- 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...

Displays switch management information including software versions.
To select menu item, press item number, or highlight item and press <Enter>.
```

Each of the above menu items accesses the read-only screens described on the following pages. See the online help for a description of the entries displayed in these screens.

General system information

Accessing system information (Menu)

From the console Main Menu, select:

1. **Status and Counters**

1. General System Information

Figure 44 Example of general switch information

```
=====-- CONSOLE - MANAGER MODE -=====
                        Status and Counters - General System Information

System Contact      :
System Location     :

Firmware revision   : K.11.00           Base MAC Addr      : 0001e7-a09900
ROM Version         : K.11.Z4           Serial Number      : S2600017409

Up Time            : 2 hours            Memory - Total     : 24,588,136
CPU Util (%)       : 1                  Free               : 19,613,568

IP Mgmt - Pkts Rx  : 0                  Packet - Total     : 832
              Pkts Tx : 0                Buffers - Free    : 793
                                           Lowest            : 769
                                           Missed           : 0
                                           24,588,1 6

Actions->  Back  Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

This screen dynamically indicates how individual switch resources are being used. See the online Help for details.

Accessing system information (CLI)

Syntax:

```
show system [ chassislocate | information | power-supply |
temperature | fans ]
```

Displays global system information and operational parameters for the switch.

chassislocate	Shows the chassisLocator LED status. Possible values are On, Off, or Blink. When the status is On or Blink, the number of minutes that the Locator LED will continue to be on or to blink is displayed. (See Example 106 (page 207))
information	Shows global system information and operational parameters for the switch. (See Example 108 (page 207) .)
power-supply	Shows chassis power supply and settings.
temperature	Shows system temperature and settings.
fans	Shows system fan status. (See Example 107 (page 207) .)

Example 106 Command results for show system chassislocate command

```
HP Switch(config)# show system chassislocate
Chassis Locator LED: ON 5 minutes 5 seconds
HP Switch(config)# show system chassislocate
Chassis Locator LED: BLINK 10 minutes 6 seconds
HP Switch(config)# show system chassislocate
Chassis Locator LED: OFF
```

Example 107 System fan status

```
HP Switch(config)# show system fans

Fan Information
  Num | State | Failures
-----+-----+-----
  Sys-1 | Fan OK | 0

0 / 1 Fans in Failure State
0 / 1 Fans have been in Failure State
```

Example 108 Switch system information

```
HP Switch(config)# show system

Status and Counters - General System Information

System Name       : HP Switch
System Contact    :
System Location   :

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None

Software revision : T.13.XX      Base MAC Addr  : 001635-b57cc0
ROM Version       : K.12.12      Serial Number   : LP621KI005

Up Time          : 51 secs      Memory - Total  : 152,455,616
CPU Util (%)     : 3           Free           : 100,527,264

IP Mgmt - Pkts Rx : 0          Packet - Total  : 6750
                Pkts Tx : 0          Buffers Free   : 5086
                                   Lowest  : 5086
                                   Missed  : 0
```

Collecting processor data with the task monitor (CLI)

The task monitor feature allows you to enable or disable the collection of processor utilization data. The `task-monitor cpu` command is equivalent to the existing debug mode command `taskusage -d`. (The `taskUsageShow` command is also available.)

When the `task-monitor` command is enabled, the `show cpu` command summarizes the processor usage by protocol and system functions.

Syntax:

```
[no] task-monitor cpu
```

Allows the collection of processor utilization data.
Only manager logins can execute this command.
The settings are not persistent, that is, there are no changes to the configuration.
(Default: Disabled)

Example 109 The `task-monitor cpu` command and `show cpu` output

```
HP Switch(config)# task-monitor cpu
HP Switch(config)# show cpu
```

```
2 percent busy, from 2865 sec ago
1 sec ave: 9 percent busy
5 sec ave: 9 percent busy
1 min ave: 1 percent busy
```

```
% CPU | Description
-----+-----
 99 | Idle
```

Switch management address information

Accessing switch management address information (Menu)

From the Main Menu, select:

1. Status and Counters ...
2. Switch Management Address Information

Figure 45 Example of management address information with VLANs configured

```
----- CONSOLE - MANAGER MODE -----
                Status and Counters - Management Address Information

Time Server Address : Disabled

VLAN Name      MAC Address      IP Address
-----
DEFAULT VLAN   0001e7-a09900    10.28.227.101
VLAN-22        0001e7-a09900    Disabled
VLAN-33        0001e7-a09900    Disabled

Actions->     Back      Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

This screen displays addresses that are important for management of the switch. If multiple VLANs are *not configured*, this screen displays a single IP address for the entire switch. See the online Help for details.

NOTE: As shown in [Figure 45 \(page 208\)](#), all VLANs on the switches use the same **MAC address**. (This includes both the statically configured VLANs and any dynamic VLANs existing on the switch as a result of GVRP operation.)

Also, the switches use a multiple forwarding database. When using multiple VLANs and connecting a switch to a device that uses a single forwarding database, such as a Switch 4000M, there are cabling and tagged port VLAN requirements. For more information on this topic, see "Multiple VLAN Considerations" in the "Static Virtual LANs (VLANs)" chapter of the *Advanced Traffic Management Guide* for your switch.

Accessing switch management address information (CLI)

Syntax:

```
show management
```

Port Status

The WebAgent and the console interface show the same port status data.

Viewing port status (CLI)

Syntax:

```
show interfaces brief
```

Viewing port status (Menu)

From the Main Menu, select:

1. Status and Counters ...
4. Port Status

Figure 46 Example of port status on the menu interface

Port	Type	Intrusion Alert	Enabled	Status	Mode	Flow Ctrl
A1		No	Yes	Down		off
A2		No	Yes	Down		off
A3		No	Yes	Down		off
A4		No	Yes	Down		off
B1	10/100TX	No	Yes	Up	100FDx	off
B2	10/100TX	No	Yes	Down	10FDx	off
B3	10/100TX	No	Yes	Down	10FDx	off
B4	10/100TX	No	Yes	Down	10FDx	off
B5	10/100TX	No	Yes	Down	10FDx	off
B6	10/100TX	No	Yes	Down	10FDx	off
B7	10/100TX	No	Yes	Down	10FDx	off

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.

Viewing port and trunk group statistics (WebAgent)

1. In the navigation pane of the WebAgent, click Interface.
2. Click Port Info/Config.

For information about this screen, click ? in the upper right corner of the WebAgent screen.

NOTE: To reset the port counters to zero, you must reboot the switch.

Port and trunk group statistics and flow control status

The features described in this section enable you to determine the traffic patterns for each port since the last reboot or reset of the switch. You can display:

- A general report of traffic on all LAN ports and trunk groups in the switch, along with the per-port flow control status (On or Off).
- A detailed summary of traffic on a selected port or trunk group.

You can also reset the counters for a specific port.

The menu interface provides a dynamic display of counters summarizing the traffic on each port. The CLI lets you see a static "snapshot" of port or trunk group statistics at a particular moment.

As mentioned above, rebooting or resetting the switch resets the counters to zero. You can also reset the counters to zero for the current session. This is useful for troubleshooting. See the [\(page 210\)](#), below.

NOTE: The **Reset** action resets the counter display to zero for the current session, but does not affect the cumulative values in the actual hardware counters. (In compliance with the SNMP standard, the values in the hardware counters are not reset to zero unless you reboot the switch.) Exiting from the console session and starting a new session restores the counter displays to the accumulated values in the hardware counters.

Accessing port and trunk statistics (Menu)

From the Main Menu, select:

1. Status and Counters ...
4. Port Counters

Figure 47 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
-----  -
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.
```

To view details about the traffic on a particular port, use the ↓ key to highlight that port number, then select **Show Details**. For example, selecting port A2 displays a screen similar to [Figure 48](#) [\(page 210\)](#), below.

Figure 48 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. (See the “NOTE” [\(page 210\)](#).)

NOTE: Once cleared, statistics cannot be reintroduced.

Accessing port and trunk group statistics (CLI)

Viewing the port counter summary report

Syntax:

```
show interfaces
```

Provides an overview of port activity for all ports on the switch.

Viewing a detailed traffic summary for specific ports

Syntax:

```
show interfaces <port-list>
```

Provides traffic details for the ports you specify.

Displaying trunk load balancing statistics

To display trunk counters information since the trunk was formed with the given ports. If ports are added or removed from the trunk-groups, statistical data is reset.

Syntax:

```
show trunk-statistics <trunk-group>
```

Displays the trunk counter information since the trunk was formed.

Example 110 Output for the show trunk-statistics command

```
HP Switch(config)# show trunk-statistics trk1

Group : Trk1 Ports : 3,4
Monitoring time : 23 hours 15 minutes

Totals

Packets Rx : 3,452,664 Bytes Rx : 14,004,243
Packets Tx : 2,121,122 Bytes Tx : 2,077,566
Packets Tx Drop :

Rates (5 minute weighted average):
Trunk utilization Rx : 30.2 %
Trunk utilization Tx : 78.2 %

Traffic Spread past 5 minutes
```

Port	%Tx	%Rx	Bytes Rx	Bytes Tx	Dropped Frame-Tx
3	27	42	1,223,445	2,112,122	123,122
4	73	58	356,233	993,222	0

Clearing trunk load balancing statistics

To display trunk counters information since the trunk was formed with the given ports. If ports are added or removed from the trunk-groups, statistical data is reset. The data is for a specific trunk.

Syntax:

```
clear trunk-statistics <trunk-group>
```

Clears statistics for all trunks if no trunks identified.

trunk-group: Clears specific trunk counter information since the trunk was formed.

Resetting the port counters

It is useful to be able to clear all counters and statistics without rebooting the switch when troubleshooting network issues. The `clear statistics global` command clears all counters and statistics for all interfaces except SNMP. You can also clear the counters and statistics for an individual port using the `clear statistics <port-list>` command.

Syntax:

```
clear statistics <<port-list> | global>
```

When executed with the `port-list` option, clears the counters and statistics for an individual port.

When executed with the `global` option, clears all counters and statistics for all interfaces except SNMP.

The `show interfaces [<port-list>]` command displays the totals accumulated since the last boot or the last `clear statistics` command was executed. The menu page also displays these totals.

SNMP displays the counter and statistics totals accumulated since the last reboot; it is not affected by the `clear statistics global` command or the `clear statistics <port-list>` command. An SNMP trap is sent whenever the statistics are cleared.

Viewing the switch's MAC address tables

Accessing MAC address views and searches (CLI)

Syntax:

```
show mac-address  
[vlan <vlan-id>]  
[<port-list>]  
[<mac-addr>]
```

Listing all learned MAC addresses on the switch, with the port number on which each MAC address was learned

```
HP Switch# show mac-address
```

Listing all learned MAC addresses on one or more ports, with their corresponding port numbers

For example, to list the learned MAC address on ports A1 through A4 and port A6:

```
HP Switch# show mac-address a1-a4,a6
```

Listing all learned MAC addresses on a VLAN, with their port numbers

This command lists the MAC addresses associated with the ports for a given VLAN. For example:

```
HP Switch# show mac-address vlan 100
```

NOTE: The switches operate with a multiple forwarding database architecture.

Finding the port on which the switch learned a specific MAC address

For example, to find the port on which the switch learns a MAC address of 080009-21ae84:

```
|| Select VLAN : DEFAULT_VLAN ||
```

Accessing MAC address views and searches (Menu)

Viewing and searching per-VLAN MAC-addresses

This feature lets you determine which switch port on a selected VLAN is being used to communicate with a specific device on the network.

From the Main Menu, select:

1. **Status and Counters ...**
5. **VLAN Address Table**

1. The switch then prompts you to select a VLAN.

```
===== CONSOLE - MANAGER MODE =====
                Status and Counters - Address Table

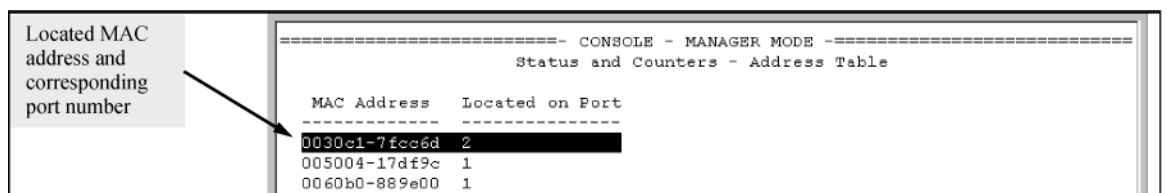
  MAC Address  Located on Port
-----
0030c1-7f49c0  A3
0030c1-7fec40  A1
0030c1-b29ac0  A3
0060b0-17de5b  A3
0060b0-880a80  A2
0060b0-df1a00  A3
0060b0-df2a00  A3
0060b0-e9a200  A3
009027-e74f90  A3
080009-21ae84  A3
080009-62c411  A3
080009-6563e2  A3

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.
```

2. Use the Space bar to select the VLAN you want, and then press **[Enter]**.
The switch then displays the MAC address table for that VLAN (Figure 49 (page 213)).

Figure 49 Example of the address table



```
Located MAC address and corresponding port number
-----
===== CONSOLE - MANAGER MODE =====
                Status and Counters - Address Table

  MAC Address  Located on Port
-----
0030c1-7fcc6d  2
005004-17df9c  1
0060b0-889e00  1
```

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.

1. Proceeding from Figure 49 (page 213), 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 and port number are highlighted if found (Figure 50 (page 214)). If the switch does not find the MAC address on the currently selected VLAN, it leaves the MAC address listing empty.

Figure 50 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
9. 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>.
```

Prompt for selecting the port to search

3. 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.

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 \(page 214\)](#), above:

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.

Accessing MSTP Data (CLI)

Syntax:

```
show spanning-tree
```

Displays the switch's global and regional spanning-tree status, plus 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.

Example

Figure 51 Output from show spanning-tree command

```

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 | Hello | 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.

Viewing internet IGMP status (CLI)

The switch uses the CLI to display the following IGMP status on a per-VLAN basis:

Show command	Output
show ip igmp	Global command listing IGMP status for all VLANs configured in the switch: <ul style="list-style-type: none"> 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
show ip igmp config	Displays the IGMP configuration information, including VLAN ID, VLAN name, status, forwarding, and Querier information.
show ip igmp <vlan-id>	Per-VLAN command listing above, IGMP status for specified VLAN (VID)

Show command	Output
show ip igmp group <ip-addr>	Lists the ports currently participating in the specified group, with port type, Access type, Age Timer data and Leave Timer data.
show ip igmp groups	Displays VLAN-ID, group address, uptime, expiration time, multicast filter type, and the last reporter for IGMP groups.
show ip igmp statistics	Displays IGMP operational information, such as VLAN IDs and names, and filtered and flooding statistics.

Example 111 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 Internal
1	DEFAULT_VLAN	No	No	Yes	125
2	VLAN2	Yes	No	Yes	125
12	New_VLAN	No	No	Yes	125

Example 112 IGMP statistical information

```
HP 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

Viewing VLAN information (CLI)

Show command	Output
show vlan	Lists: <ul style="list-style-type: none"> • Maximum number of VLANs to support • Existing VLANs • Status (static or dynamic) • Primary VLAN
show vlan <vlan-id>	For the specified VLAN, lists: <ul style="list-style-type: none"> • Name, VID, and status (static/dynamic) • Per-port mode (tagged, untagged, forbid, no/auto) • "Unknown VLAN" setting (Learn, Block, Disable) • Port status (up/down)

Example

Suppose that your switch has the following VLANs:

Ports	VLAN	VID
A1-A12	DEFAULT_VLAN	1
A1, A2	VLAN-33	33
A3, A4	VLAN-44	44

The next three examples show how you could list data on the above VLANs.

Example 113 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
```

Note: Because ports A1 and A2 are not members of VLAN-44, it does not appear in this listing.

Example 114 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
```

Example 115 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          Untagged Learn          Up
A3          Untagged Learn          Up
A4          Untagged Learn          Down
A5          Untagged Learn          Up
A6          Untagged Learn          Up
A7          Untagged Learn          Up
```

WebAgent status information

The WebAgent Status screen provides an overview of the status of the switch. Scroll down to view more details. For information about this screen, click on ? in the upper right corner of the WebAgent screen. For an example of a status screen, see [Figure 52 \(page 219\)](#).

Figure 52 Example of a WebAgent status screen

Home > Status [Reboot] ?

Switch Status Change ?

System Name: ProCurve Switch 8212zl

System Location:

System Contact:

System Uptime: 2 days, 2 hours, 32 minutes, 44 seconds

System CPU Util: 0%

System Memory: 117288960 Bytes

Unit Information Change ?

Product Name: ProCurve Switch 8212zl(J9091A)

IP Address: 15.255.133.38

Base MAC Address: 00 18 71 b9 85 00

Serial Number: LP713BX00E

Mgmt Server: http://www.hp.com/rnd/device_help

Firmware Version: K.15.01.0000c.ROMK.15.04

VLAN ?

Name	Status	IP Address
DEFAULT_VLAN	Port-based	15.255.133.38

[Change]

Alert Log ?

Search: [Refresh] [Delete]

Date & Time	Status	Alert	Description
More >>			

ProCurve Switch 8212zl(J9091A)

Power: <input checked="" type="checkbox"/>	FAN	MM 1 Status: ACTIVE
Fault: <input type="checkbox"/>	TMP	MM 2 Status: DOWN/BOO
	POE	

EMPTY CHARACTERS: 1 2 3 4 5 6 7 8 9 10 11 12 13

Details

Port Name:	Totals:	Receive
Enabled:	Bytes:	
Type:	Unicast:	

C Troubleshooting

Overview

This appendix 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: HP periodically places switch software updates on the HP Switch Networking website. HP Switch 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

Use these approaches to diagnose switch problems:

- Check the HP website for software updates that may have solved your problem: www.hp.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 HP 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 chapter "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 HP 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

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:

Example 116 Indication that routing is enabled

```
HP Switch(config)# show running
Running configuration:
; J9091A Configuration Editor; Created on release #k.15.06
hostname " HPswitch "
ip default-gateway 10.33.248.1
ip routing I
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
```

- I** Indicates that routing is enabled, a requirement for ACL operation. (There is an exception. Refer to the **Note**, below.)
-

NOTE: 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.

Example 117 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 10.28.100.100 255.255.255.255 2  
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
```

1 Correct.

2 Incorrect. No mask needed to specify a single host.

3 Incorrect. 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)" chapter 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 118 (page 225)) and applying it outbound on VLAN 1 in Figure 53 (page 225) 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.

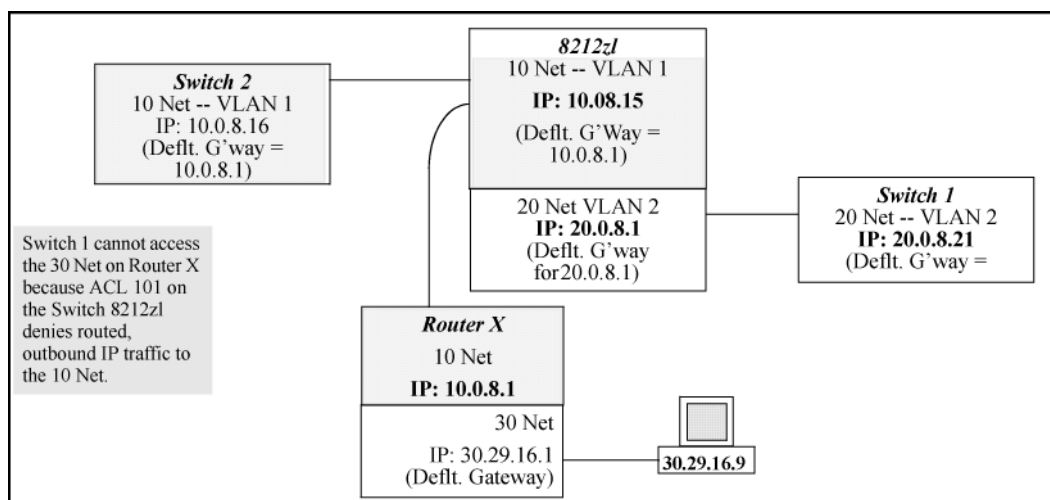
In Figure 53 (page 225), this ACE (see data in bold below) denies access to the 10 Net's 10.0.8.1 router gateway needed by the 20 Net (Subnet mask is 255.255.255.0).

Example 118 ACE blocking an entire subnet

```
HP 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
```

Figure 53 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):

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.

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:

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 (page 225).

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:

1. Use the `no trunk <port-number>` command to disable the static trunk assignment.
2. Execute `interface <port-number> lacp`.

△ CAUTION: 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, HP recommends that you either disable the port or disconnect it from the LAN.

Mesh-related problems

Traffic on a dynamic VLAN does not get through the switch mesh

GVRP enables dynamic VLANs. Ensure that all switches in the mesh have GVRP enabled.

Port-based access control (802.1X)-related problems

NOTE: To list the 802.1X port-access Event Log messages stored on the switch, use `show log 802`.

See also “Radius-related problems” (page 229).

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.

Example 119 Authenticator ports remain "open" until activated

```
HP 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
```

- 1 Port 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.

Example 120 Displaying encryption keys

```
HP 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

Auth Acct DM/ Time
Server IP Addr 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.

NOTE: 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.

Example 121 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

      Auth Acct DM/ Time
Server IP Addr  Port Port CoA 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

- ⚠ **CAUTION:** If you enable MSTP, HP 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 chapter "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 or cycling the power 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-server host` 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.

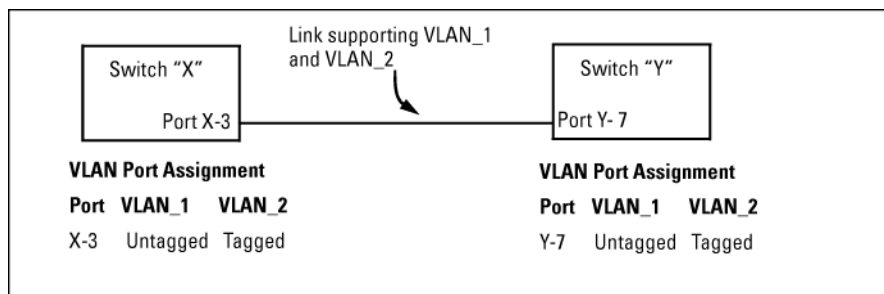
None of the devices assigned to one or more VLANs on an 802.1Q-compliant switch are being recognized

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 54 \(page 233\)](#).

Figure 54 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.

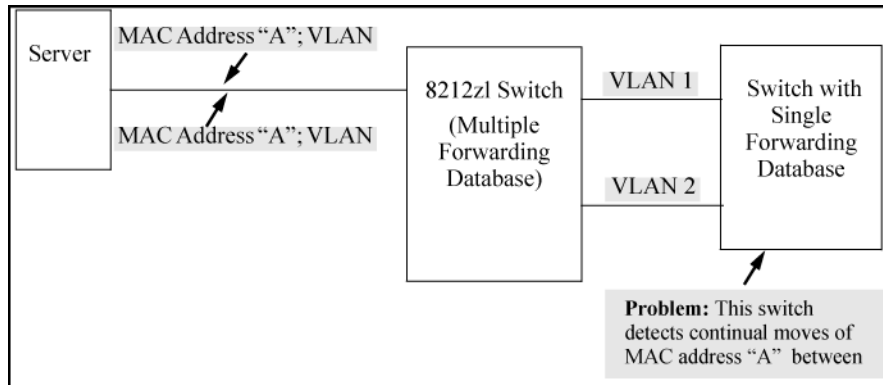
Duplicate MAC addresses across VLANs

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 55 Example of duplicate MAC address



Disabled overlapping subnet configuration

Previous software versions allowed configuration of VLAN IP addresses in overlapping subnets which can cause incorrect routing of packets and result in IP communication failure. As of software version WB.15.09, overlapping subnet configurations are no longer allowed. An overlapping subnet is determined by the configuration order. The subnet that is configured first is valid, but any subsequent IP addresses that overlap are not allowed.

When the switch is booted into software version WB.15.09 or later, and the configuration file includes overlapping subnets, the following occurs:

- The event log provides an error message in the format:
`ip: VLANx : IP initialization failed for vlan x.`
For a multinetted VLAN (multiple IP addresses assigned to the VLAN), only the IP addresses that are overlapping subnets are removed. The other IP addresses on the VLAN are retained and function correctly. The error message can be somewhat misleading; the IP addresses on the VLAN that are not overlapping are initialized correctly.
- The output of the `show ip` command correctly indicates that the overlapping IP address does not exist on the VLANs that have error messages in the event log.
- The output of the `show running-config` command incorrectly indicates that the overlapping IP address is configured. In [Example 122 "An IP address that is not actually configured on the VLAN"](#), the IP address shown in VLAN6 is not actually configured on the VLAN; it has been removed.

Example 122 An IP address that is not actually configured on the VLAN

```
HP Switch(config)# show running-config
.
.
.
vlan 5
  name "VLAN5"
  ip address 11.22.33.1 255.0.0.0
  exit
vlan 6
  name "VLAN6"
  ip address 11.23.34.1 255.255.255.0
  exit
```

The information is retained in the config file to allow you to boot up the switch and have it function as it did when it was configured with earlier software that allows overlapping subnets. If you attempt to remove the overlapping subnet from the VLAN, the switch displays an error message similar to:

The IP address *<ip-address>* is not configured on this VLAN

This occurs because the overlapping IP address has been removed and is not visible to the switch. To resolve this:

- Enter the `show ip` command to determine which addresses are visible to the switch.
- Remove the erroneous IP addresses from the config file by entering the `no ip address` command to remove all the IP addresses from the specific VLAN. Be sure to document the other valid IP addresses on that VLAN so they can be restored after removing the erroneous IP addresses from the config file.

If you go back to a software version prior to WB.15.09 before removing the overlapping IP address, the prior software version enables the overlapping IP subnet.

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.

HP recommends that you replace a failed fan tray assembly within one minute of removing it.

Mitigating flapping transceivers

In traditional HP switches, the state of a link is driven directly by the reported state of the port, which is required for rapid detection of link faults. However, the consequence of this is that a marginal transceiver, optical, or wire cabling, one that "flaps" up and down several times per second, can cause STP and other protocols to react poorly, resulting in a network outage. The link-flap option expands the functionality of the existing fault finder function to include a "link-flap" event and a new action of "warn-and-disable." Together, these additions allow the errant condition to be detected, and the port in question can be optionally disabled.

Syntax:

```
fault-finder <link-flap> sensitivity <low | medium | high>
  > action <warn | warn-and-disable>
```

Default settings: Sensitivity = Medium; Action = Warn

Sensitivity thresholds are static. In a 10-second window, if more than the threshold number of link state transitions (up or down) are detected, the event is triggered. The 10-second window is statically determined, that is, the counters are reset every 10 seconds, as opposed to being a sliding window. The counters are polled twice per second (every 500 milliseconds), and the event is triggered if the sensitivity threshold is crossed at that time.

The sensitivity thresholds are:

High	3 transitions in 10 seconds
Medium	6 transitions in 10 seconds
Low	10 transitions in 10 seconds

Configuring the link-flap event and corresponding action applies to all ports and port types (it is a global setting per FFI event type). Note that normal link transition protocols may prevent link state changes from occurring fast enough to trigger the event for some port types, configurations, and sensitivity settings.

When the link-flap threshold is met for a port configured for warn (for example, `fault-finder link-flap sensitivity medium action warn`), the following message is seen in the switch event log.

```
02672 FFI: port <number>-Excessive link state transitions
```

When the link-flap threshold is met for a port configured for warn-and-disable (for example, `fault-finder linkflap sensitivity medium action warn-and-disable`), the following messages are seen in the switch event log.

```
02672 FFI: port <number>-Excessive link state transitions
```

```
02673 FFI: port <number>-Port disabled by Fault-finder.
```

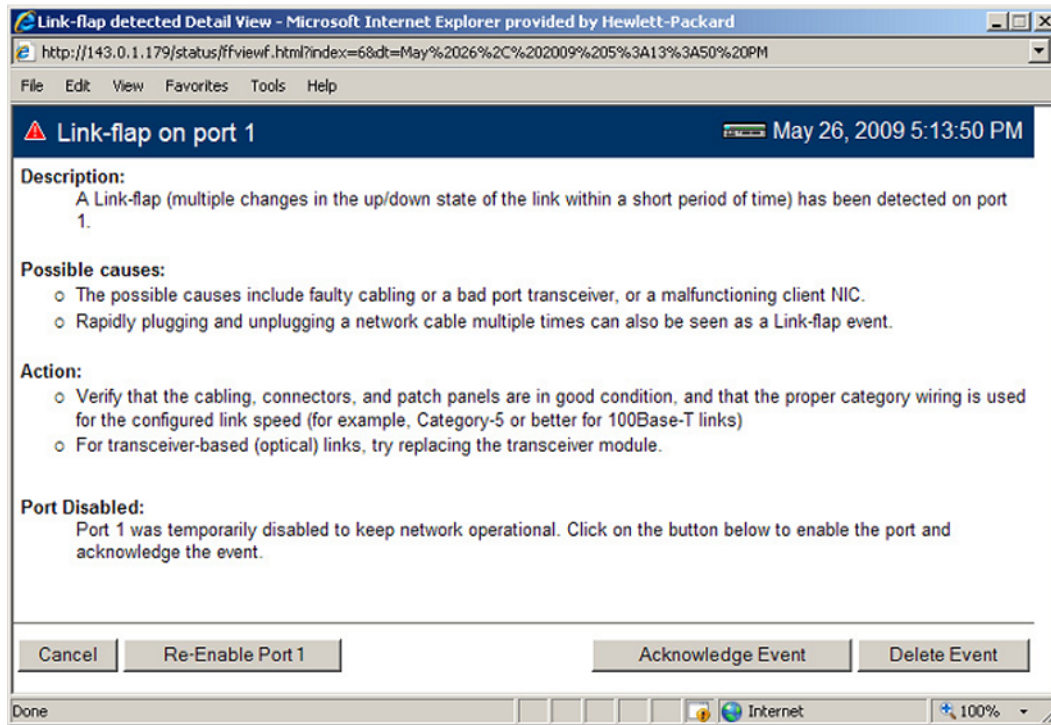
```
02674 FFI: port <number>-Administrator action required to re-enable.
```

The warn-and-disable action is available for all fault-finder events on an individual basis. It may be used, for example, to disable a port when excessive broadcasts are received. Because the fault-generated disabling of a port requires operator intervention to re-enable the port, such configuration should be used with care. For example, link-flap-initiated disablement is not desired on ports that are at the client edge of the network, because link state changes there are frequent and expected.

HP does not recommend automatic disabling of a port at the core or distribution layers when excessive broadcasts are detected, because of the potential to disable large parts of the network that may be uninvolved and for the opportunity to create a denial-of-service attack.

Within the Web Management interface, double-clicking an event on a port that was configured with warn-and-disable and that has met the threshold to trigger the disable action brings up a dialog box with the event details, as shown in [Figure 56 \(page 237\)](#). The event dialog box now contains a button at the bottom of the page, which can be used to re-enable the disabled port. The button remains, even if the port has already been brought up through a prior exercise of it, or if the port was re-enabled via some other interface (for example, the command line). Re-enabling an already enabled port has no effect. The button to acknowledge the event remains unchanged.

Figure 56 Link-flap on port 1 event detail dialog box



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
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

Product #	Description	Support ¹
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)

NOTE: 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.

Example 123 Output for a specified transceiver

```
HP 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.

Example 124 Output when no transceiver is present in specified interface

```
HP Switch(config)# show interfaces transceiver 22
```

No Transceiver found on interface 22

When no ports are specified, information for all transceivers found is displayed.

Example 125 Output when no ports are specified

```
HP 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.

Example 126 Output when "all" is specified

```
HP 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 19 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 [Table 20 \(page 240\)](#), [Table 21 \(page 240\)](#), and [Table 22 \(page 240\)](#) is only displayed when the transceiver supports DOM.

Table 20 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 [Table 21 \(page 240\)](#).

Table 21 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 [Table 22 \(page 240\)](#).

Table 22 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
Laser temperature fault	Laser temperature fault
Laser output power fault	Laser output power fault

Table 22 Alarm and error information (XENPAK transceivers) (continued)

Alarm	Description
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.

Example 127 Detailed information for a 1000SX Mini-GBIC transceiver

```
HP 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.

Example 128 Detailed information for a 10GbE-LR transceiver

```
HP 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.

Example 129 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			

Example 130 Copper cable diagnostic test results

```
HP 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
```

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. An electrical transceiver value is displayed as N/A.

Table 23 General transceiver information *(continued)*

Parameter	Description
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.

NOTE: 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 57 \(page 245\)](#), each Event Log entry is composed of six or seven fields, depending on whether numbering is turned on or not:

Figure 57 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 HP internal diagnostic information.
Date	The date in the format <i>mm/dd/yy</i> when an entry is recorded in the log.
Time	The time in the format <i>hh:mm:ss</i> 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 [no] log-number 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. Table 24 (page 245) lists the different system modules with a description of each one.
Event message	A brief description of the operating event.

Table 24 Event Log system modules

System module	Description	Documented in HP Switch hardware/software guide
802.1x	802.1X authentication: Provides access control on a per-client or per-port basis: <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>
acl	ACLs: Filter layer-3 IP traffic to or from a host to block unwanted IP traffic and block or limit other protocol traffic such as TCP, UDP, IGMP, and ICMP. ACEs specify the filter criteria and an action (permit or deny) to take on a packet if it meets the criteria.	<i>Advanced Traffic Management 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>
arp-protect	Dynamic ARP Protection: Protects the network from ARP cache poisoning. Only valid ARP requests and	<i>Access Security Guide</i>

Table 24 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
	responses are relayed or used to update the local ARP cache. ARP packets with invalid IP-to-MAC address bindings advertised in the source protocol address and source physical address fields are discarded.	
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. HP does not support the transmission of CDP packets to neighbor devices.	<i>Management and Configuration Guide</i>
connfilt	<p>Connection-rate filtering: Used on the network edge to protect the network from attack by worm-like malicious code by detecting hosts that are generating IP traffic that exhibits this behavior and (optionally) either throttling or dropping all IP traffic from the offending hosts.</p> <p>Connection-rate filtering messages include events on virus throttling. Virus throttling uses connection-rate filtering to stop the propagation of malicious agents.</p>	<i>Access Security 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>
cos	<p>Class of Service (CoS): Provides priority handling of packets traversing the switch, based on the IEEE 802.1p priority carried by each packet.</p> <p>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.</p>	<i>Advanced Traffic Management Guide</i>
dca	Dynamic Configuration Arbiter (DCA) determines the client-specific parameters that are assigned in an authentication session.	<i>Access Security 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	<i>Management and Configuration Guide</i>

Table 24 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
	gateway, Timep Server address, and TFTP server address.	
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>
dhcp-snoop	DHCP snooping: Protects your network from common DHCP attacks, such as address spoofing and repeated address requests.	<i>Access Security Guide</i>
dma	Direct Access Memory (DMA): Transmits and receives packets between the CPU and the switch. Not used for logging messages in software release K.13.xx.	—
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>
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>
idm	Identity-driven Management: Optional management application used to monitor and control access to switch.	<i>Advanced Traffic Management 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>
inst-mon	Instrumentation Monitor: Identifies attacks on the switch by generating alerts for detected anomalies.	<i>Access Security Guide</i>

Table 24 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
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 Multicast and Routing Guide</i>
ipaddrmgr	IP Address Manager: Programs IP routing information in switch hardware.	<i>Multicast and Routing 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>
licensing	HP Switch premium licensing: Provides access to expanded features on certain HP switches.	<i>Premium License Installation Guide</i>
kms	Key Management System: Configures and maintains security information (keys) for all routing protocols, including a timing mechanism for activating and deactivating an individual protocol.	<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>
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 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>

Table 24 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
loop_protect	Loop protection: Detects the formation of loops when an unmanaged device on the network drops spanning tree packets and provides protection by transmitting loop protocol packets out ports on which loop protection has been enabled.	<i>Advanced Traffic Management 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>
mgr	HP PCM and PCM+: Windows-based network management solutions for managing and monitoring performance of HP switches. PCM messages also include events for configuration operations.	<i>Management and Configuration Guide</i>
mlد	Multicast Listener Discovery (MLD): IPv6 protocol used by a router to discover the presence of multicast listeners. MLD can also optimize IPv6 multicast traffic flow with the snooping feature.	<i>Multicast and Routing Guide</i>
mtm	Multicast Traffic Manager (MTM): Controls and coordinates L3 multicast traffic for upper layer protocols.	<i>Multicast and Routing 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).	—
pim	Protocol-independent multicast (PIM) routing: Enables IP multicast traffic to be transmitted for multimedia	<i>Multicast and Routing Guide</i>

Table 24 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
	applications throughout a network without being blocked at routed interface (VLAN) boundaries.	
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 Management and Configuration Guide 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>
ratelim	Rate-limiting: Enables a port to limit the amount of bandwidth a user or device may utilize for inbound traffic on the switch.	<i>Management and Configuration Guide</i>
sflow	Flow sampling: sFlow is an industry standard sampling technology, defined by RFC 3176, used to continuously monitor traffic flows on all ports providing network-wide visibility into the use of the network.	<i>Management and Configuration 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>
ssh	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. 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.	<i>Access Security Guide</i>
ssl	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	<i>Access Security Guide</i>

Table 24 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
	station clients capable of SSL/TLS operation.	
stack	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.	<i>Advanced Traffic Management Guide</i>
stp	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.	<i>Advanced Traffic Management Guide</i>
system	Switch management, including system configuration, switch bootup, activation of boot ROM image, memory buffers, traffic and security filters. System messages also include events from management interfaces (menu, CLI, and HP PCM+) used to reconfigure the switch and monitor switch status and performance.	<i>Management and Configuration Guide</i> <i>Access Security Guide</i>
tacacs	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).	<i>Access Security Guide</i>
tcp	Transmission Control Protocol: A transport protocol that runs on IP and is used to set up connections.	<i>Advanced Traffic Management Guide</i>
telnet	Session established on the switch from a remote device through the Telnet virtual terminal protocol.	<i>Management and Configuration Guide</i>
tftp	Trivial File Transfer Protocol: Supports the download of files to the switch from a TFTP network server.	<i>Management and Configuration Guide</i>
timep	Time Protocol: Synchronizes and ensures a uniform time among interoperating devices.	<i>Management and Configuration Guide</i>
udld	Uni-directional Link Detection: 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.	<i>Access Security Guide</i>

Table 24 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
udpf	UDP broadcast forwarding: Supports the forwarding of client requests sent as limited IP broadcasts addressed to a UDP application port on a network server.	<i>Multicast and Routing Guide</i>
update	Updates (TFTP or serial) to HP switch software and updates to running-config and start-up config files	<i>Management and Configuration Guide</i>
usb	Auxiliary port that allows you to connect external devices to the switch.	<i>Installation and Getting Started Guide</i>
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. • A protocol-based VLAN creates a layer-3 broadcast domain for traffic of a particular routing protocol, and comprises member ports that bridge traffic of the specified protocol type among themselves. <p>VLAN messages include events from management interfaces (menu, CLI, and HP 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>
xrrp	Extended Router Redundancy Protocol	—

Using the Menu

To display the Event Log from the Main Menu, select `Event Log`. [Example 131 \(page 253\)](#) shows a sample event log display.

Example 131 An event log display

```
HP 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 [Table 3-3 \(page 253\)](#).

Table 25 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:

```
show logging [-a, -b, -r, -s, -t, -m, -p, -w, -i, -d]
[<option-str>]
```

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.
-p	Displays only performance log events.
-w	Displays only warning log events.
-i	Displays only informational log events.
-d	Displays only debug log events.
<option-str>	Displays all Event Log entries that contain the specified text. Use an <option-str> value with -a or -r to further filter show logging command output.

Example

To display all Event Log messages that have "system" in the message text or module name, enter the following command:

```
HP 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:

```
HP Switch# show logging system
```

Clearing Event Log entries

Syntax:

```
clear logging
```

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.

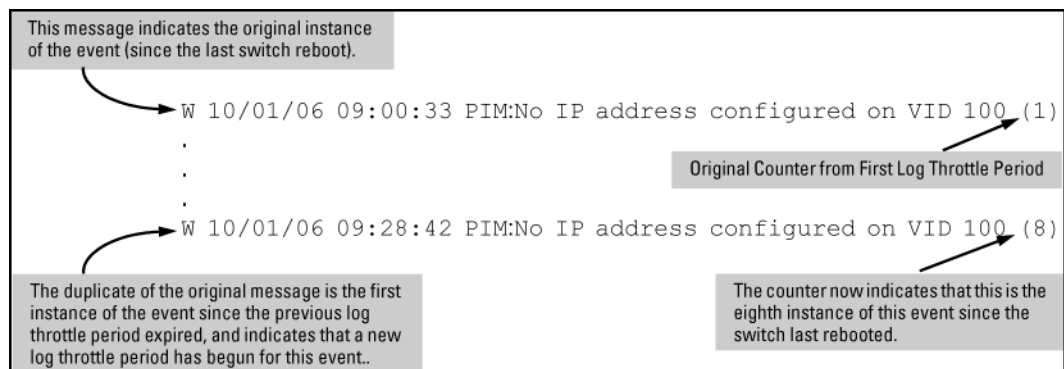
NOTE: In Example 132 “The first instance of an event message and counter” the counter (1) indicates that this is the first instance of this event since the switch last rebooted.

Example 132 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)
```

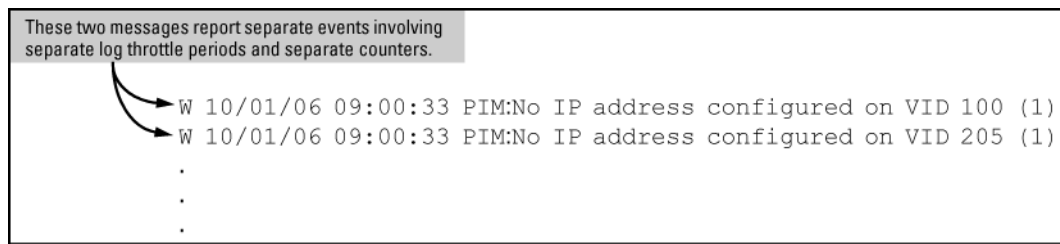
If PIM operation causes the same event to occur six more times during the initial log throttle period, there are no further entries in the Event Log. However, if the event occurs again after the log throttle period has expired, the switch repeats the message (with an updated counter) and starts a new log throttle period.

Figure 58 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 59 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 [Table 3-4 \(page 256\)](#). (The same operation applies for messages sent to any configured SNMP trap receivers.)

Table 26 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</code> <code><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

Example 133 Sending running config changes to the syslog server

```
HP 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:

- | | |
|--|---|
| <ul style="list-style-type: none">• ACL "deny" matches• Dynamic ARP protection events• DHCP snooping events• DIPLD events• Events recorded in the switch's Event Log• IP routing events (IPv4 and IPv6) | <ul style="list-style-type: none">• LLDP events• SNMP events• SSH events• Wireless services events |
|--|---|

- Use the `logging` command to select a subset of Event Log messages to send to an external device for debugging purposes according to:

- | |
|--|
| <ul style="list-style-type: none">• Severity level• System module |
|--|

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” \(page 266\)](#) and [“Configuring a syslog server” \(page 268\)](#).

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.
	facility	(Optional) The logging facility command specifies the destination (facility) subsystem used on a syslog server for debug reports.
	priority-desc	A text string associated with the values of facility, severity, and system-module.
	severity	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.)
debug Command	acl	Sends ACL syslog logging to configured debug destinations. When there is a match with a "deny" statement, directs the resulting message to the configured debug destinations.
	all	Sends debug logging to configured debug destinations for all ACL, Event Log, IP-OSPF, and IP-RIP options.
	cdp	Displays CDP information.
	destination	logging: Disables or re-enables syslog logging on one or more syslog servers configured with the logging <code>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 debug <code><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	fib: Displays IP Forwarding Information Base messages and events.

		<p>forwarding: Sends IPv4 forwarding messages to the debug destinations.</p> <p>packet: Sends IPv4 packet messages to the debug destinations.</p> <p>pim [packet [filter source <ip-addr> vlan <vid>>]]</p> <p>: Enables or disables tracing of PIM messages.</p> <p>Note: When PIM debugging is enabled, the following message displays:</p> <p>PIM Debugging can be extremely CPU intensive when run on a device with an existing high CPU load or on a switch with more than 10 PIM-enabled VLANs. In high load situations, the switch may suffer from protocol starvation, high latency, or even reload. When debugging a switch with more than 10 PIM-enabled VLANs, the "vlan" option in "debug ip pim packet" should be utilized. Debugging should only be used temporarily while troubleshooting problems. Customers are advised to exercise caution when running this command in a highstress production network.</p> <p>rip: Sends RIP event logging to the debug destinations.</p>
	ipv6	<p>dhcpv6-client: Sends DHCPv6 client debug messages to the configured debug destination.</p> <p>dhcpv6-relay: Sends DHCPv6 relay debug messages to the configured debug destination.</p> <p>forwarding: Sends IPv6 forwarding messages to the debug destination(s)</p> <p>nd: Sends IPv6 debug messages for IPv6 neighbor discovery to the configured debug destinations.</p>
	lldp	Sends LLDP debug messages to the debug destinations.
	security	Sends security messages to the debug destination.
	services	Displays debug messages on the services module.
	snmp	Sends snmp messages to the debug destination.

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

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 "a (page 260)" 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 (page 260)" 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

```
HP Switch(config)# logging severity <debug | major | error | warning | info>  
HP 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 [Table 24 \(page 245\)](#) .

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:

```
HP Switch(config)# no logging severity <debug | major | error | warning | info>
HP Switch(config)# no logging system-module <system-module>
```

CAUTION: 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.)

Example 134 Output of the show debug command

```
HP 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 60 Syslog configuration to receive event log messages from specified system module and severity levels

```
HP Switch(config)# show debug
Debug Logging
 Destination: None
 Enabled debug types:
  None are enabled
HP Switch(config)# logging 10.28.38.164
HP Switch(config)# write memory
HP Switch(config)# show debug
Debug Logging
 Destination:
 Logging --
  10.28.38.164
 Facility=user
 Severity=debug
 System module=all-pass
 Enabled debug types:
  event
HP Switch(config)# logging severity error
HP Switch(config)# logging system-module iplock
```

Displays the default debug configuration. (No Syslog server IP addresses or debug types are configured.)

When you configure a Syslog IP address with the **logging** command, by default, the switch enables debug messaging to the Syslog address and the **user** facility on the Syslog server, and sends Event Log messages of all severity levels from all system modules.

You can enter the **logging severity** and **logging system-module** commands to specify a subset of Event Log messages to send to the Syslog server.

As shown at the top of [Figure 60 \(page 262\)](#), 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 61 \(page 263\)](#).

Figure 61 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.

NOTE: To configure a syslog server, use the `logging <syslog-ip-addr>` command. For more information, see [“Configuring a syslog server” \(page 268\)](#).

Debug messages

Syntax:

[no] debug <debug-type>

acl

When a match occurs on an ACL "deny" ACE (with log configured), the switch sends an ACL message to configured debug destinations. For

	<p>information on ACLs, see the "Access Control Lists (ACLs)" chapter in the latest version of the following guides:</p> <ul style="list-style-type: none"> • IPv4 ACLs: <i>Access Security Guide</i> • IPv6 ACLs: <i>IPv6 Configuration Guide</i> <p>NOTE: ACE matches (hits) for permit and deny entries can be tracked using the <code>show statistics <aclv4 aclv6></code> command.</p> <p>(Default: Disabled—ACL messages for traffic that matches "deny" entries are not sent.)</p>
all	<p>Configures the switch to send all debug message types to configured debug destinations.</p> <p>(Default: Disabled—No debug messages are sent.)</p>
cdp	<p>Sends CDP information to configured debug destinations.</p>
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> <p>NOTE: This value does not affect the reception of event notification messages in the Event Log on the switch.</p> <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 forwarding [packet] [rip]]	<p>Sends IP messages to configured destinations.</p>
ip [fib [events]]	<p>For the configured debug destinations:</p> <p>events—Sends IP forwarding information base events.</p>
ip [rip [database event trigger]]	<p>rip <database event trigger>—Enables the specified RIP message type for the configured destination(s).</p> <p>database—Displays database changes.</p>

	<p>event—Displays RIP events.</p> <p>trigger—Displays trigger messages.</p>
<pre>ipv6 [dhcpv6-client dhcpv6-relay nd packet]</pre>	<p>NOTE: See the "IPv6 Diagnostic and Troubleshooting" chapter in the <i>IPv6 Configuration Guide</i> for your switch for more detailed IPv6 debug options.</p> <p>When no debug options are included, displays debug messages for all IPv6 debug options.</p> <p>dhcpv6-client [events packet] —Displays DHCPv6 client event and packet data.</p> <p>dhcpv6-relay [events packet] —Displays DHCPv6 relay event and relay packet data.</p> <p>nd—Displays debug messages for IPv6 neighbor discovery.</p> <p>packet—Displays IPv6 packet messages.</p>
<pre>lldp</pre>	<p>Enables all LLDP message types for the configured destinations.</p>
<pre>security [arp-protect dhcp-snooping dynamic-ip-lockdown port-access port-security radius-server ssh tacacs-server user-profile-mib]</pre>	<p>arp-protect—Sends dynamic ARP protection debug messages to configured debug destinations.</p> <p>dhcp-snooping—Sends DHCP snooping debug messages to configured debug destinations.</p> <p>agent—Displays DHCP snooping agent messages.</p> <p>event—Displays DHCP snooping event messages.</p> <p>packet—Displays DHCP snooping packet messages.</p> <p>dynamic-ip-lockdown—Sends dynamic IP lockdown debug messages to the debug destination.</p> <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>
<pre>snmp <pdu></pre>	<p>Displays the SNMP debug messages.</p> <p>pdu—Displays SNMP pdu debug messages.</p>

Filtering debug messages by debug type

Debug message filtering provides the ability to filter debug messages by debug type. Multiple debug filters can be applied to a debug type.

Syntax:

```
[no] debug <debug type> include [ ip ip-addr list | port  
<port-list> | vlan <vid-list> ]
```

Enables or disables debug message filtering for a debug type. The filter types are:
IPv4 or IPv6 address list Port list VLAN list

Default: Disabled

Example 135 Setting an SNMP event filter for an IP address

```
HP Switch(config)# debug snmp event include ip 10.10.10.1
```

```
HP Switch(config)# show debug
```

```
Debug Logging
```

```
Destination: Session
```

```
Enabled debug types:  
snmp trap include ip 10.10.10.1
```

Example 136 Setting an IP RIP filter for port A4

```
HP Switch(config)# debug ip rip database include port A4
```

```
HP Switch(config)# show debug
```

```
Debug Logging
```

```
Destination: Session
```

```
Enabled debug types:  
ip rip database include port A4  
snmp trap include ip 10.10.10.1
```

Example 137 Setting a filter for fatal SSH messages on a VLAN

```
HP Switch(config)# debug ssh fatal include vlan 2
```

```
HP Switch(config)# show debug
```

```
Debug Logging
```

```
Destination: Session
```

```
Enabled debug types:  
ip rip database include port A4  
snmp trap include ip 10.10.10.1  
ssh (fatal) include vlan 2
```

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” (page 263)) are sent.</p> <p>(Default: Logging disabled)</p> <p>To configure a syslog server IP address, see “Configuring a syslog server” (page 268).</p> <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” (page 272).</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 (HP Switch#_).</p> <p>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.</p> <p>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.</p> <p>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.

CAUTION: 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:

```
HP Switch(config)# no logging severity <debug | major | error | warning | info>
HP 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”](#) (page 263)).
- `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”](#) (page 271) and [“Configuring the system module used to select the Event Log messages sent to a syslog server”](#) (page 272).)

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”](#) (page 258)).

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”](#) (page 263).)

<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”](#) (page 263).)

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.

Example 138 Configuring TCP for logging message transmission using the default port

```
HP Switch(config)# logging 192.123.4.5 tcp  
(Default TCP port 1470 is used.)
```

Example 139 Configuring TCP for logging message transmission using a specified port

```
HP Switch(config)# logging 192.123.4.5 9514  
(TCP port 9514 is used.)
```

Example 140 Configuring UDP for logging message transmission using the default port

```
HP Switch(config)# logging 192.123.4.5 udp  
(Default UDP port 514 is used.)
```

Example 141 Configuring UDP for logging message transmission using a specified port

```
HP 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.) HP 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
cron	cron/at subsystem
sys9	cron/at subsystem
sys10 - sys14	Reserved for system use
local10 - local17	Reserved for system use

Use the `no` form of the command to remove the configured facility and reconfigure the default (user) value.

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.

CAUTION: Entering the `no logging` command removes ALL the syslog server addresses without a verification prompt.

NOTE: The HP enterprise MIB `hpicfSyslog.mib` allows the configuration and monitoring of syslog for SNMP (RFC 3164 supported).

The CLI command is:

Syntax:

```
logging <ip-addr> [control-descr <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

NOTE: To remove the description using SNMP, set the description to an empty string.

Example

Example 142 The logging command with a control description

```
HP 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

Example 143 The logging command with a priority description

```
HP Switch(config)# logging priority-descr severe-pri
```

NOTE: 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 HP 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” \(page 268\)](#).

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.

NOTE: 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” (page 268).)

See Table 24 (page 245) 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.

NOTE: 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.
event (debug type)	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. If no syslog server is configured, the sending of Event Log messages is <code>disabled</code> .
IP (debug type)	Disabled.

- Debugcommands 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:

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" \(page 40\)](#).

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.

NOTE: 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" \(page 288\)](#).

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.

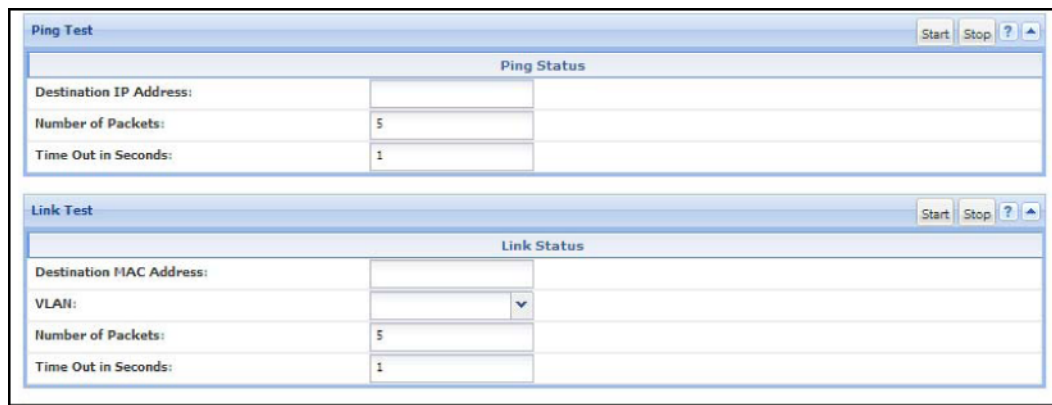
Executing ping or link tests (WebAgent)

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 62 \(page 274\)](#).

Figure 62 Ping test and link test screen on the WebAgent



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> [repetitions <1-10000>]
[timeout <1-60>][source < ip-address | <vlan-id> | loopback
<0-7>> ] [data-size <0-65471>] [data-fill <0-1024>]
```

Syntax:

```
ping6 <ipv6-address | hostname> [repetitions <1-10000>]
[timeout <1-60>][source < ip-address | vlan-id | loopback
<0-7>> ] [data-size <0-65471>] [data-fill <0-1024>]
```

Sends ICMP echo requests to determine if another device is alive.

<ip-address hostname>	Target IP address or hostname of the destination node being pinged
repetitions <1-10000>	Number of ping packets sent to the destination address. Default: 1
timeout <1-60>	Timeout interval in seconds; the ECHO REPLY must be received before this time interval expires for the ping to be successful. Default: 5
source <ip-addr vid loopback <0-7>>	Source IP address, VLAN ID, or loopback address used for the ping. 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

Example 144 Ping tests

```
HP Switch# ping 10.10.10.10
10.10.10.10 is alive, time = 15 ms
```

```
HP 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
```

```
HP Switch# ping 10.10.10.10 timeout 2
10.10.10.10 is alive, time = 10 ms
```

```
HP 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]**.

NOTE: To use the ping (or traceroute) command with host names or fully qualified domain names, see “DNS resolver” (page 288).

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 63 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> [maxttl <1-255>] [minttl <1-255>] [probes <1-5>] [source <ip-address | source-vlan <vid> | loopback <0-7> ] [<timeout 1-120>]
```

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>]	For the current instance of <code>traceroute</code> , changes the minimum number of hops allowed for each probe packet sent along the route. <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. 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)
[maxttl <1-255>]	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)
[<code>probes <1-5></code>]	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)
[[<code>source <ip-addr> <vlan-id></code>]]	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.
[<code>source <ip-addr vid loopback <0-7>>]</code>]	The source IPv4 address, VLAN ID, or Loopback address.
[<code>timeout <1-120></code>]	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

NOTE: 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 64 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

```

Continuing from the previous example (Figure 64 (page 277)), executing `traceroute` with an insufficient `maxttl` for the actual hop count produces an output similar to this:

Figure 65 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 66 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.

show config	Displays the startup configuration.
show running-config	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)

[Example 145](#) shows sample output from the `show tech` command.

Example 145 The show tech command

```
HP 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 : K.14.xx      Base MAC Addr  : 001871-c42f00

ROM Version       : K.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 [“Customizing show tech command output” \(page 281\)](#).

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.


1. In HyperTerminal, click on Transfer | Capture Text...(see Figure 67 (page 281)).

Figure 67 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, as shown in Figure 68 (page 281).

Figure 68 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:

```
HP 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.

Customizing show tech command output

Use the copy show tech command to customize the detailed switch information displayed with the show tech command to suit your troubleshooting needs.

To customize the information displayed with the show tech command:

1. Determine the information that you want to gather to troubleshoot a problem in switch operation.
2. Enter the copy show tech command to specify the data files that contain the information you want to view.

Syntax:

```
copy <source> show-tech
```

Specifies the operational and configuration data from one or more source files to be displayed by the show tech command. Enter the command once for each data file that you want to include in the display.

Default: Displays data from all source files, where *<source>* can be any one of the following values:

<pre>command-output "<command>"</pre>	<p>Includes the output of a specified command in <code>show tech</code> command output.</p> <p>Enter the command name between double-quotation marks, for example, <code>copy "show system" show-tech</code>.</p>
<pre>crash-log</pre>	<p>Copies the contents of the Crash Log to <code>show tech</code> command output.</p>
<pre>event-log</pre>	<p>Copies the contents of the Event Log to <code>show tech</code> command output.</p>
<pre>tftp config <startup-config running-config <ip-addr> <remote-file> <pc unix></pre>	<p>Downloads the contents of a configuration file from a remote host to <code>show tech</code> command output, where:</p> <p><i><ip-addr></i>: Specifies the IP address of the remote host device.</p> <p><i><remote-file></i>: Specifies the pathname on the remote host for the configuration file whose contents you want to include in the command output.</p> <p><code>pc unix</code>: Specifies whether the remote host is a DOS-based PC or UNIX workstation.</p> <p>For more information on using <code>copy tftp</code> commands, see the appendix "File transfers" (page 183).</p>
<pre>xmodem config <startup-config config <filename> command-file <acl-filename.txt> <pc unix></pre>	<p>Copies the contents of a configuration file or ACL command file from a serially connected PC or UNIX workstation to <code>show tech</code> command output, where:</p> <p><code>startup-config</code>: Specifies the name of the startup configuration file on the connected device.</p> <p><code>config <filename></code>: Specifies the pathname of a configuration file on the connected device.</p> <p><code>command-file <acl-filename.txt></code>: Specifies the pathname of an ACL command file on the connected device.</p> <p><code>pc unix</code>: Specifies whether the connected device is a DOS-based PC or UNIX workstation.</p> <p>For more information on using <code>copy xmodem</code> commands, see the appendix "File transfers" (page 183).</p>

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" \(page 285\)](#)).

```
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)" chapter.

```
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" chapter).

```
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.

include	Only the lines that contain the matching pattern are displayed in the output.
exclude	Only the lines that contain the matching pattern are <i>not</i> displayed in the output.
begin	The display of the output begins with the line that contains the matching pattern.

NOTE: Pattern matching is case-sensitive.

Below are examples of what portions of the running config file display depending on the option chosen.

Example 146 Pattern matching with include option

```
HP Switch(config)# show run | include ipv6 1
  ipv6 enable
  ipv6 enable
ipv6 access-list "EH-01"
HP Switch(config)#
```

- 1 Displays only lines that contain "ipv6".
-

Example 147 Pattern matching with exclude option

```
HP Switch(config)# show run | exclude ipv6 1
```

Running configuration:

```
; J9299A Configuration Editor; Created on release #RA.15.XX

hostname "HP Switch"
snmp-server community "notpublic" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged 1-24
  ip address dhcp-bootp
  no untagged 21-24
  exit
vlan 20
  name "VLAN20"
  untagged 21-24
  no ip address
  exit
policy qos "michael"
  exit
  sequence 10 deny tcp 2001:db8:255::/48 2001:db8:125::/48
  exit
no autorun
password manager
```

- 1 Displays all lines that do not contain "ipv6".
-

Example 148 Pattern matching with begin option

```
HP 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
```

- 1 Displays the running config beginning at the first line that contains "ipv6".
-

Example 149 (page 285) 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.

Example 149 The show arp command and pattern matching with the include option

```
HP 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

HP 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

NOTE: HP 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.

NOTE: 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

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.

NOTE: 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.
```

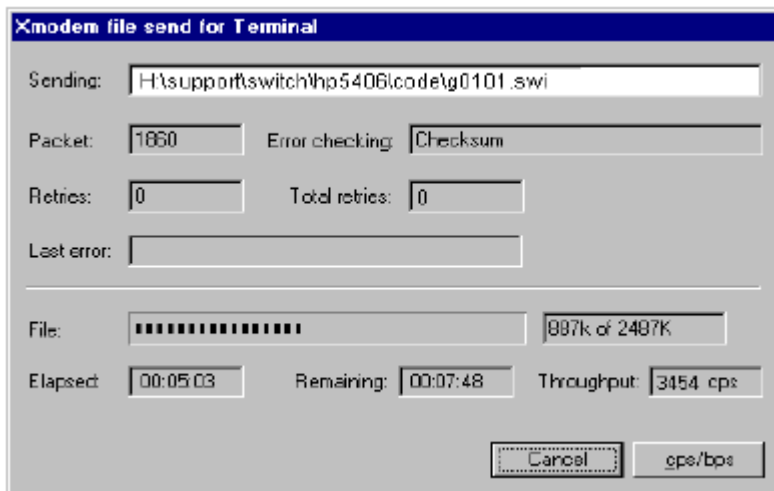
```
=>
```
3. Because the OS file is large, you can increase the speed of the download by changing the switch console and terminal emulator baud rates to a high speed. For example:
 - a. Change the switch baud rate to 115,200 Bps.

```
=> 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 69 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 serverthen:
 - 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 70 Example of using either a host name or a fully qualified domain name



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 `tracert` command using the target's fully qualified DNS name should succeed.

Figure 71 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
  
```

Fully Qualified Host Name for the Target Host

IP Address for Target Host "remote-01"

Configuring and using DNS resolution with DNS-compatible commands

The DNS-compatible commands include ping and traceroute.)

1. Determine the following:
 - The IP address for a DNS server operating in a domain in your network.
 - The priority (1 to 3) of the selected server, relative to other DNS servers in the domain.
 - 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" \(page 289\)](#).) Note that if a domain suffix is not configured, fully qualified domain names can be used to resolve DNS-compatible commands.
 - 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 step 1 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 [Figure 70 \(page 289\)](#).) 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.

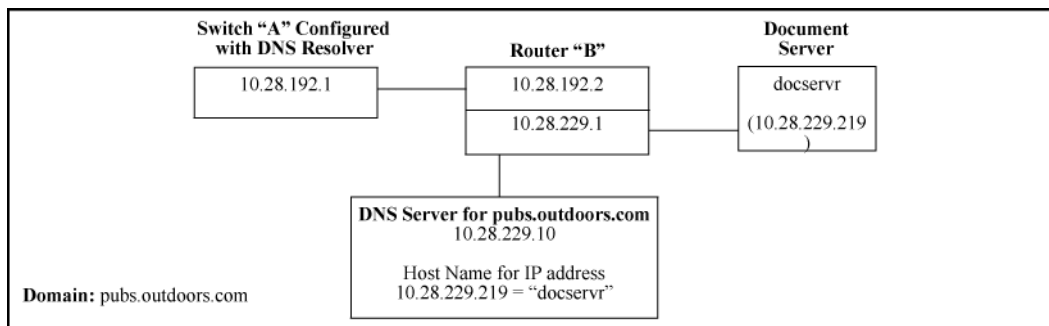
The switch supports one domain suffix entry and three DNS server IP address entries. (See the preceding command description.)

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 72 \(page 291\)](#), 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 72 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.

Example 150 Configuring switch "A" in Figure 72 (page 291) to support DNS resolution

```
HP Switch(config)# ip dns server-address 10.28.229.10
HP Switch(config)# ip dns domain-name pbs.outdoors.com
```

Example 151 Ping and traceroute execution for the network in Figure 72 (page 291)

```
HP Switch(config)# ping docservr
10.28.229.219 is alive, time = 1 ms
```

```
HP 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
```


1 First-Hop Router ("B") **2** Traceroute Target

As mentioned under "Basic operation" (page 289), 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 72 (page 291) as a target:

Figure 73 Example of ping and traceroute execution when only the DNS server IP address is configured

```
HP Switch# ping [docservr.pbs.outdoors.com]
10.28.229.219 is alive, time = 1 ms

HP Switch# traceroute [docservr.pbs.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 74 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   : pbs.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.

Locating a switch (Locator LED)

To locate where a particular switch is physically installed, use the `chassislocate` command to activate the blue Locator LED on the switch's front panel.

Syntax:

```
chassislocate [ blink | on | off ]
```

Locates a switch by using the blue Locate LED on the front panel.

<code>blink <1-1440></code>	Blinks the chassis Locate LED for a specified number of minutes (Default: 30 minutes).
<code>on <1-1440></code>	Turns the chassis Locate LED on for a specified number of minutes (Default: 30 minutes).
<code>off</code>	Turns the chassis Locate LED off.

Example 152 Locating a switch with the `chassislocate` command

```
HP Switch(config)# chassislocate
  blink <1-1440>      Blink the chassis locate led (default 30 minutes).
  off                Turn the chassis locate led off.
  on <1-1440>        Turn the chassis locate led on (default 30 minutes).
HP Switch(config)# chassislocate
```

For redundant management systems, if the active management module failover, the Locator LED does not remain lit.

D MAC Address Management

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” (page 296).

MAC addresses are assigned at the factory. The switch automatically implements these addresses for VLANs and ports as they are added to the switch.

NOTE: 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 VLAN 1 and all other VLANs configured on the switch.)

NOTE: 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.

[<i>port-list</i>]	Lists the MAC addresses of the devices the switch has detected, on the specified ports.
[<i>mac-addr</i>]	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: <code>MAC address <mac-addr> not found.</code>
[<i>vlan <vid></i>]	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.

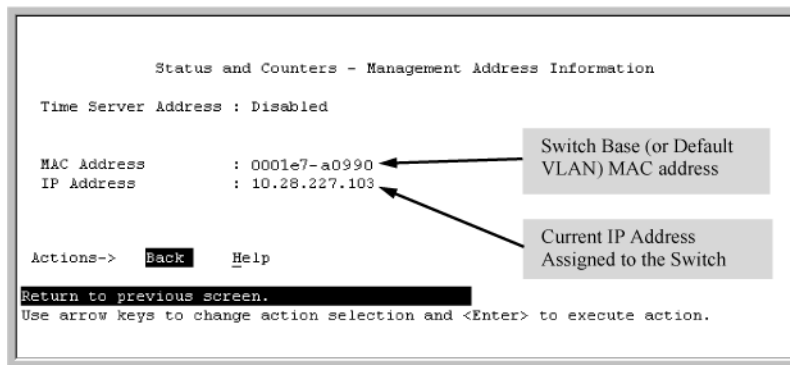
Also, the Base MAC address appears on a label on the back of the switch.

NOTE: 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 75 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.

NOTE: This procedure displays the MAC addresses for all ports and existing VLANs in the switch, regardless of which VLAN you select.

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:

```
HP Switch# walkmib ifPhysAddress
```

(The above command is not case-sensitive.)

Example:

An HP 8212zl switch with the following module configuration shows MAC address assignments similar to those shown in [Figure 76 \(page 297\)](#):

- 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 76 Example of Port MAC address assignments on a switch

```

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 =
  
```

ifPhysAddress.1 - 4: Ports A1 - A4 in Slot A
(Addresses 5 - 24 in slot A are unused.)

ifPhysAddress.49 - 72: Ports C1 - C24 in Slot C
(In this example, there is no module in slot B.)

ifPhysAddress.362 Base MAC Address (MAC Address for default VLAN; VID = 1)

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 *Advanced Traffic Management Guide* for your switch.

Virtual

E Monitoring Resources

Displaying current resource usage

To display current resource usage in the switch, enter the following command:

Syntax:

```
show <qos | access-list | policy> resources
```

Displays the resource usage of the policy enforcement engine on the switch by software feature. For each type of resource, the amount still available and the amount used by each software feature is shown.

<code>show resources</code>	This output allows you to view current resource usage and, if necessary, prioritize and reconfigure software features to free resources reserved for less important features.
<code>qos</code> <code>access-list</code> <code>openflow</code> <code>policy</code>	Display the same command output and provide different ways to access task-specific information. NOTE: See "Viewing OpenFlow Resources" in the <i>OpenFlow Administrators Guide</i> for your switch.

“Displaying current resource usage” (page 299) shows the resource usage on a switch configured for ACLs, QoS, RADIUS-based authentication, and other features:

- The "Rules Used" columns show that ACLs, VT, mirroring, and other features (for example, Management VLAN) have been configured globally or per-VLAN, because identical resource consumption is displayed for each port range in the switch. If ACLs were configured per-port, the number of rules used in each port range would be different.

Example 153 Displaying current resource usage

```
HP Switch(config)# show access-list resources
```

```
Resource usage in Policy Enforcement Engine
```

Ports	Rules	Rules Used			
	Available	ACL	QoS	IDM	Other
1-48	2006	10	5	0	6

Ports	Meters	Meters Used			
	Available	ACL	QoS	IDM	Other
1-48	255		5		0

Ports	Application	Application Port Ranges Used			
	Port Ranges Available	ACL	QoS	IDM	Other
1-48	31	1	0	0	0

```
2 of 16 Policy Engine management resources used.
```

Key:

ACL = Access Control Lists

QoS = Device & Application Port Priority

IDM = Identity Driven Management

Other = Management VLAN, DHCP Snooping, ARP Protection, RA Guard.

Resource usage includes resources actually in use, or reserved for future

use by the listed feature. Internal dedicated-purpose resources, such as

port bandwidth limits or VLAN QoS priority, are not included.

Viewing information on resource usage

The switch allows you to view information about the current usage and availability of resources in the Policy Enforcement engine, including the following software features:

- Access control lists (ACL)
- Quality-of-service (QoS), including device and application port priority, ICMP rate-limiting, and QoS policies
- Dynamic assignment of per-port or per-user ACLs and QoS through RADIUS authentication designated as “IDM”, with or without the optional identity-driven management (IDM) application
- Virus throttling (VT) using connection-rate filtering
- Mirroring policies, including switch configuration as an endpoint for remote intelligent mirroring
- Other features, including:
 - Management VLAN
 - DHCP snooping

- Dynamic ARP protection
- Jumbo IP-MTU

Policy enforcement engine

The policy enforcement engine is the hardware element in the switch that manages QoS, mirroring, and ACL policies, as well as other software features, using the rules that you configure. Resource usage in the policy enforcement engine is based on how these features are configured on the switch:

- Resource usage by dynamic port ACLs is determined as follows:
 - Dynamic port ACLs configured by a RADIUS server (with or without the optional IDM application) for an authenticated client determine the current resource consumption for this feature on a specified slot. When a client session ends, the resources in use for that client become available for other uses.
- When the following features are configured globally or per-VLAN, resource usage is applied across all port groups or all slots with installed modules:
 - ACLs
 - QoS configurations that use the following commands:
 - QoS device priority (IP address) through the CLI using the `qos device-priority` command
 - QoS application port through the CLI using `qos tcp-port` or `qos udp-port`
 - VLAN QoS policies through the CLI using `service-policy`
 - Management VLAN configuration
 - DHCP snooping
 - Dynamic ARP protection
 - Remote mirroring endpoint configuration
 - Mirror policies per VLAN through the CLI using `monitor service`
 - Jumbo IP-MTU
- When the following features are configured per-port, resource usage is applied only to the slot or port group on which the feature is configured:
 - ACLs or QoS applied per-port or per-user through RADIUS authentication
 - ACLs applied per-port through the CLI using the `ip access-group` or `ipv6 traffic-filter` commands
 - QoS policies applied per port through the CLI using the `service-policy` command
 - Mirror policies applied per-port through the CLI using the `monitor all service` and `service-policy` commands
 - ICMP rate-limiting through the CLI using the `rate-limit icmp` command

Usage notes for show resources output

- A 1:1 mapping of internal rules to configured policies in the switch does not necessarily exist. As a result, displaying current resource usage is the most reliable method for keeping track of available resources. Also, because some internal resources are used by multiple features, deleting a feature configuration may not increase the amount of available resources.
- Resource usage includes resources actually in use or reserved for future use by the listed features.

- "Internal dedicated-purpose resources" include the following features:
 - Per-port ingress and egress rate limiting through the CLI using `rate-limit in/out`
 - Per-port or per-VLAN priority or DSCP through the CLI using `qos priority` or `qos dscp`
 - Per protocol priority through the CLI using `qos protocol`
- The "Available" columns display the resources available for additional feature use.
- The "IDM" column shows the resources used for RADIUS-based authentication with or without the IDM option.
- "Meters" are used when applying either ICMP rate-limiting or a QoS policy with a rate-limit class action.

When insufficient resources are available

The switch has ample resources for configuring features and supporting RADIUS-authenticated clients (with or without the optional IDMap application).

If the resources supporting these features become fully subscribed:

- The current feature configuration, RADIUS-authenticated client sessions, and VT instances continue to operate normally.
- The switch generates an event log notice to say that current resources are fully subscribed.
- Currently engaged resources must be released before any of the following actions are supported:
 - Modifying currently configured ACLs, IDM, VT, and other software features, such as Management VLAN, DHCP snooping, and dynamic ARP protection.
You can modify currently configured classifier-based QoS and mirroring policies if a policy has not been applied to an interface. However, sufficient resources must be available when you apply a configured policy to an interface.
 - Acceptance of new RADIUS-based client authentication requests (displayed as a new resource entry for IDM).
Failure to authenticate a client that presents valid credentials may indicate that insufficient resources are available for the features configured for the client in the RADIUS server. To troubleshoot, check the event log.
 - Throttling or blocking of newly detected clients with high rate-of-connection requests (as defined by the current VT configuration).
The switch continues to generate Event Log notifications (and SNMP trap notification, if configured) for new instances of high-connection-rate behavior detected by the VT feature.

F Monitoring Resources

Viewing information on resource usage

The switch allows you to view information about the current usage and availability of resources in the Policy Enforcement engine, including the following software features:

- Access control lists (ACL)
- Quality-of-service (QoS), including device and application port priority, ICMP rate-limiting, and QoS policies
- Dynamic assignment of per-port or per-user ACLs and QoS through RADIUS authentication designated as “IDM”, with or without the optional identity-driven management (IDM) application
- Other features, including:
 - Management VLAN
 - DHCP snooping
 - Dynamic ARP protection
 - Jumbo IP-MTU

Policy enforcement engine

The policy enforcement engine is the hardware element in the switch that manages QoS, mirroring, and ACL policies, as well as other software features, using the rules that you configure. Resource usage in the policy enforcement engine is based on how these features are configured on the switch.

ACLs configured by a RADIUS Dynamic port ACLs configured by a RADIUS server (with or without the optional IDM application) for an authenticated client determine the current resource consumption for this feature on a specified port. When a client session ends, the resources in use for that client become available for other uses.

Resource usage by the following features (when configured globally or per VLAN), applies across all port groups:

- ACLs
- QoS configurations that use the following commands:
 - QoS device priority (IP Address) through the CLI using the `qos device-priority` command
 - QoS application port through the CLI using `qos tcp-port` or `qos udp-port`
- Management VLAN configuration
- DHCP snooping
- Dynamic ARP protection
- Jumbo IP-MTU

Resource usage on ACLs or QoS are configured per-port through RADIUS authentication.

Displaying current resource usage

To display current resource usage in the switch, enter the `show qos resources` or `show access-list resources` command.

Syntax:

```
show <qos | access-list>
resources
```

Displays the resource usage of the policy enforcement engine on the switch by software feature. For each type of resource, the amount still available and the amount used by each software feature is shown.

The `qos` and `access-list` parameters display the same command output.

The `show <qos|access-list> resources` command output allows you to view current resource usage and, if necessary, help prioritize and reconfigure software features to free resources reserved for less important features.

Resources are used dynamically, that is, resources are reallocated depending on usage. If IDM ACLs are not configured, then the resources are available for other ACLs, or QoS.

An IDM ACL group uses 4 resources for every ACE (2 address matches times 2 entries per address). When the initial IDM ACE is configured, 256 entries are reserved for IDM ACLs, leaving 128 entries free for allocation. When nothing is configured, all 384 entries display as "free".

[Example 154 "Displaying current resource usage with 31 ACE in 1 ACL"](#) shows the resource usage on a switch configured for ACLs and QoS.

Example 154 Displaying current resource usage with 31 ACE in 1 ACL

```
HP Switch(config)# show access-list resources
```

```
Policy Engine Resource Usage
```

Group	Rules Allocated	Rules Used	Group Number
QoS	0	0	1
CLI-ACL	0	0	2
IDM-ACL	256	126	3
Free	128		

Usage notes for show resources output

- A 1:1 mapping of internal rules to configured policies in the switch does not necessarily exist. As a result, displaying current resource usage is the most reliable method for keeping track of available resources. Also, because some internal resources are used by multiple features, deleting a feature configuration may not increase the amount of available resources.
- Resource usage includes resources actually in use or reserved for future use by the listed features.
- "Internal dedicated-purpose resources" include the following features:
 - Per-port ingress and egress rate limiting through the CLI using `rate-limit in`
 - Per-port or per-VLAN priority or DSCP through the CLI using `qos priority` or `qos dscp`
 - Per protocol priority through the CLI using `qos protocol`

When insufficient resources are available

The switch has ample resources for configuring features and supporting RADIUS-authenticated clients (with or without the optional IDMap application).

If the resources supporting these features become fully subscribed:

- The current feature configuration continues to operate normally.
- The switch generates an event log notice to say that current resources are fully subscribed.
- Currently engaged resources must be released before any of the following actions are supported:
 - Configuration of new entries for ACL, QoS, IDM, and other features (Management VLAN, DHCP snooping, dynamic ARP protection).
 - Acceptance of new RADIUS-based client authentication requests (displayed as a new resource entry for IDM).

NOTE: Failure to authenticate a client that presents valid credentials may indicate that insufficient resources are available for the features configured for the client in the RADIUS server. To troubleshoot, check the event log.

G Daylight Saving Time on HP Switches

This information applies to the following HP switches:

- 212M
- 224M
- 1600M
- 2400M
- 2424M
- 4000M
- 8000M
- Series 2500
- Series 2510
- Series 2600
- Series 2610
- Series 2620
- Series 2800
- Switch 2910
- Series 3400cl
- Series 3500
- Series 3500yl
- Series 3600
- Series 3800
- Series 4100gl
- Series 4200vl
- Series 5300xl
- Series 5400zl
- Switch 6108
- Switch 6200yl
- Series 6400cl
- Switch 6600
- Series 8200zl
- HP AdvanceStack Switches
- HP AdvanceStack Routers

HP Switches provide a way to automatically adjust the system clock for Daylight Saving Time (DST) changes. To use this feature, define the month and date to begin and to end the change from standard time. In addition to the value "none" (no time changes), there are five pre-defined settings, named:

- Alaska
- Canada and Continental US
- Middle Europe and Portugal
- Southern Hemisphere
- Western Europe

The pre-defined settings follow these rules:

Alaska:

- Begin DST at 2 am on the second Sunday in March.
- End DST at 2 am on the first Sunday in November.

Canada and Continental US:

- Begin DST at 2 am on the second Sunday in March.
- End DST at 2 am on the first Sunday in November.

Middle Europe and Portugal:

- Begin DST at 2 am the first Sunday on or after March 25th.
- End DST at 2 am the first Sunday on or after September 24th.

Southern Hemisphere:

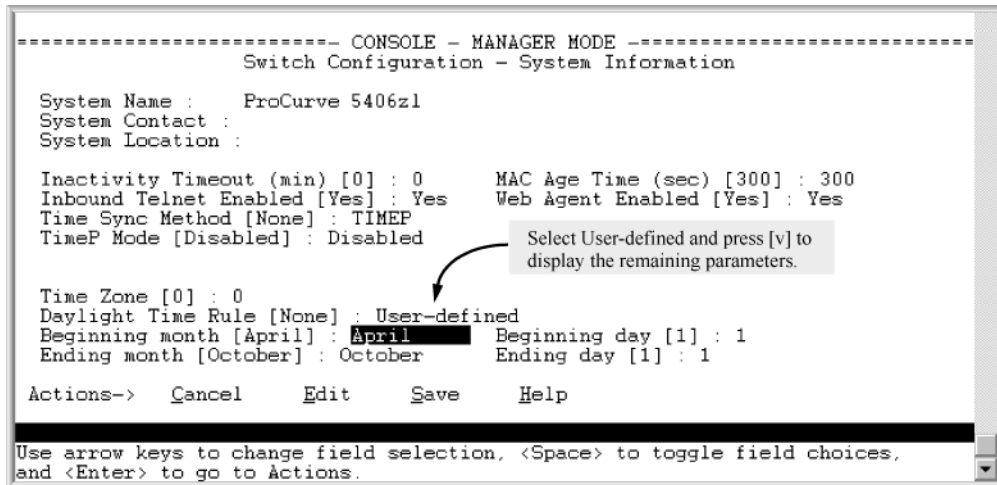
- Begin DST at 2 am the first Sunday on or after October 25th.
- End DST at 2 am the first Sunday on or after March 1st.

Western Europe:

- Begin DST at 2 am the first Sunday on or after March 23rd.
- End DST at 2 am the first Sunday on or after October 23rd.

A sixth option named "User defined" allows you to customize the DST configuration by entering the beginning month and date plus the ending month and date for the time change. The menu interface screen looks like this (all month/date entries are at their default values):

Figure 77 Menu interface with "user-defined" daylight time rule option



Before configuring a "User defined" daylight time rule, it is important to understand how the switch treats the entries. The switch knows which dates are Sundays and uses an algorithm to determine on which date to change the system clock, given the configured "Beginning day" and "Ending day":

- If the configured day is a Sunday, the time changes at 2 am on that day.
- If the configured day is not a Sunday, the time changes at 2 am on the first Sunday after the configured day.

This is true for both the "Beginning day" and the "Ending day."

With that algorithm, you should use the value "1" to represent "first Sunday of the month," and a value equal to "number of days in the month minus 6" to represent "last Sunday of the month." This allows a single configuration for every year, no matter what date is the appropriate Sunday to change the clock.

H Scalability: IP Address, VLAN, and Routing Maximum Values

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	

Subject	Maximum
RIP interfaces	128
IPv6 Routing Protocol	
DHCPv6 Helper Addresses	32 unique addresses; multiple instances of same address counts as 1 towards maximum

-
- ¹ Actual availability depends on combined resource usage on the switch. See [“Monitoring Resources” \(page 298\)](#).
 - 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.

I Power-Saving Features

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.

Example 155 The savepower led command

```
HP 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.

Example 156 Output of the show savepower led command

```
HP Switch(config)# show savepower led
```

```
LED Save Power Information
```

```
Configuration Status : Enabled
```

J Support and Other Resources

Intended audience

This guide is intended for network administrators with intermediate-to-advanced knowledge of computer networking.

Related documentation

The following sources provide related information:

- *Power over Ethernet (PoE/PoE+) Planning and Implementation Guide*
- *HP Management and Configuration Guide*
- *HP Advanced Traffic Management Guide*
- *HP Access Security Guide*
- *HP Multicast and Routing Guide*
- *HP IPv6 Configuration Guide*

You can also find the documents referenced in this guide on the Manuals page of the HP Business Support Center website:

<http://www.hp.com/support/manuals>.

Contacting HP

HP technical support

For worldwide technical support information, see the HP support website:

<http://www.hp.com/networking>.

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Subscription service

HP recommends that you register your product at the Subscriber's choice for business website:

<http://www.hp.com/go/e-updates>

After registering, you will receive email notifications of product enhancements, new driver versions, firmware updates, and other product resources.

Related information

HP websites

- HP:
<http://www.hp.com>
- HP Networking:
<http://www.hp.com/go/networking>

- HP Partner Locator:
http://www.hp.com/service_locator
- HP Software Downloads:
<http://www.hp.com/support/downloads>

Typographical conventions

Table 27 Document conventions

Convention	Element
Blue text: Table 26	Cross-reference links and email addresses
Blue underlined text: http://www.hp.com	Website addresses
Bold text	<ul style="list-style-type: none"> • Keys that are pressed • Text entered into a GUI element, such as a box • Text entered as a CLI command • GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
<i>Italic</i> text	<ul style="list-style-type: none"> • Text emphasis • Variables for which you must supply a value when executing a command
Monospace text	<ul style="list-style-type: none"> • File and directory names • System output • Code • Commands, their arguments, and argument values
<i>Monospace italic</i> text	<ul style="list-style-type: none"> • Code variables • Command variables
Monospace bold text	Emphasized monospace text
. . . .	Indication that example continues

Command syntax statements

Syntax

```
ip default-gateway <ip-addr> | routing
```

Syntax

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

- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([]) indicate optional elements.
- Braces () enclose required elements.
- Braces within square brackets ([]) indicate a required element within an optional choice.
- Boldface indicates use of a CLI command, part of a CLI command syntax, or other displayed element in general text. For example:

"Use the `copy tftp` command to download the key from a TFTP server."

- Italics indicate variables for which you must supply a value when executing the command. For example, in this command syntax, you must provide one or more port numbers:

```
aaa port-access authenticator <port-list>
```

Command prompts

In the default configuration, your switch displays a CLI prompt similar to the following example:

```
HP Switch 2620#
```

To simplify recognition, this guide uses `HP Switch` to represent command prompts for all switch models. For example:

```
HP Switch#
```

(You can use the `hostname` command to change the text in the CLI prompt.)

Screen simulations

Figures containing simulated screen text and command output look similar to this:

Figure 78 Example of a simulated screen

```
HP Switch> show version
Image stamp:      /sw/code/build/info
                  May 1, 2010 13:43:13
                  K.15.01.0031
                  139
Boot Image: Primary
```

In some cases, brief command-output sequences appear without figure identification. For example:

```
HP Switch(config)# clear public-key
HP Switch(config)# show ip client-public-key
show_client_public_key: cannot stat keyfile
```

Configuration and operation examples

Unless otherwise noted, examples using a particular switch model apply to all switch models covered by this guide.

Keys

Simulations of actual keys use a bold, sans-serif typeface with square brackets. For example, the Tab key appears as **[Tab]** and the "Y" key appears as **[Y]**.

To set up and install the switch in your network

Physical installation

Use the *Installation and Getting Started Guide* for the following:

- Notes, cautions, and warnings related to installing and using the switch and its related modules.
- Instructions for physically installing the switch in your network.
- Quickly assigning an IP address and subnet mask, setting a Manager password, and (optionally) configuring other basic features.
- Interpreting LED behavior.

For the latest version of the *Installation and Getting Started Guide* for your switch, see "Getting documentation from the web".

Product warranties

For information about HP Networking product warranties, see the warranty information website: <http://www.hp.com/networking>

Table 28 lists related products and their part numbers.

Table 28 Applicable Products

Product	Part Number
HP 2620 Switches	
HP 2620-24 Switch	J9623A
HP 2620-24-PPoE+ Switch	J9624A
HP 2620-24-PoE+ Switch	J9625A
HP 2620-48 Switch	J9626A
HP 2620-48-PoE+ Switch	J9627A

Online help

Menu interface

If you need information on specific parameters in the menu interface, refer to the online help provided in the interface.

Command-line interface

If you need information on a specific command in the CLI, type the command name followed by the word `help`.

HP customer support services

If you are having trouble with your switch, Hewlett-Packard offers support 24 hours a day, seven days a week through the use of a number of automated electronic services. See the Customer Support/Warranty booklet that came with your switch for information on how to use these services to get technical support. HP provides up-to-date customer care, support, and warranty information at

www.hp.com/networking.

Your HP authorized network reseller can also provide assistance, both with services that they offer and with services offered by HP.

Before calling support

Before calling your networking dealer or HP Support, to make the support process most efficient, first retrieve the following information:

Information item	Information location
Product identification, including mini-GBICs	The front of the switch and on labels on the mini-GBICs
Details about the switch—status including the software (OS) version, a copy of the switch configuration, a copy of the switch Event Log, and a copy of the switch status and counters information	Switch console: <code>show tech</code> command
Copy of your network topology map, including network addresses assigned to the relevant devices	Your network records

Glossary of Terms and Acronyms

ACE	Access control entry
ACL	Access control list
active PoE port	A PoE-enabled port connected to a PD requesting power.
active port	A port linked to another active device (regardless of whether MSTP is blocking the link).
adjacent device	See "Neighbor or Neighbor Device"
advertisement	See LLDPDU
all-traffic rate-limiting	Applies a rate-limit to all traffic (including ICMP traffic) on an interface.
AM	Active management module. A management module that booted successfully and is actively managing the switch.
bps	bits per second
BSD rpc	Berkeley UNIX remote copy
CDP	Cisco discovery protocol. Supports reading CDP packets received from neighbor devices, enabling a switch to learn about adjacent CDP devices.
chassis	Hardware operation, including modules and ports, power supply, fans, transceivers, CPU interrupt errors, switch temperature, and so on.
classifier-based mirroring policy	The service policy applied to a monitored (port or VLAN) interface that specifies the classes of traffic to be copied to preconfigured mirroring destinations.
CoS	Class of service. Provides priority handling of packets traversing the switch, based on the IEEE 802.1p priority carried by each packet.
DCA	Dynamic configuration arbiter. Determines the client-specific parameters that are assigned in an authentication session.
destination	The host device that is connected to an exit port on the local source switch or a remote switch, and associated with a mirrorsession number (1 to 4). See also Exit Port and Host.
DHCP	Dynamic host configuration protocol
direction-based mirroring	On an interface configured for mirroring, the traffic direction (entering or leaving the switch, or both) is used as criteria for selecting the traffic to be mirrored.
DLC	data link layer classification
DLL	data link layer
DMA	Direct access memory. Transmits and receives packets between the CPU and the switch.
DNS	Domain name system
domain suffix	Includes all labels to the right of the unique host name in a fully qualified domain name assigned to an IP address. For example, in the fully qualified domain name "device53.evergreen.trees.org," the domain suffix is "evergreen.trees.org," while "device53" is the unique (host) name assigned to a specific IP address.
DoS	Denial of service
DT	Distributed trunk
DTD	Distributed trunking device
DTE	Data terminal equipment
DTIP	Distributed trunking internet protocol
DTS	Distributed trunking switches
ECS	Emergency call service
EEE	Energy efficient ethernet
ELIN	Emergency location identification number. A valid telephone number in the North American Numbering Plan format and assigned to a multiline telephone system operator by the appropriate

authority. This number calls a public service answering point (PSAP) and relays automatic location identification data to the PSAP.

exit port

The port to which a traffic analyzer or IDS is connected to receive mirrored traffic.

- For local mirroring, an exit port can be any port to which a traffic analyzer or IDS is connected and that is not configured as a monitored interface. You can configure up to four exit ports for local mirroring on a switch, using the command: `mirror session port exit-port`
- For remote mirroring, the destination IP address (`dst-ip`) and exit port in a remote mirroring endpoint can belong to different VLANs. You can configure up to 32 exit ports for remote mirroring on a switch, using the command: `mirror endpoint ip src-ip src-udp-port dst-ip exit-port`



CAUTION: An exit port should be connected only to a network analyzer, IDS, or other network edge device that has no connection to other network resources. Connecting a mirroring exit port to a network can result in serious network performance problems, and is strongly discouraged by HP Switches Networking.

exit switch

The switch with the exit port to which a destination device is connected. See also Exit Port.

failed management module

A management module that did not pass selftest and is not in standby mode.

FFI (event type)

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.

FIB

Forwarding information base.

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.

fully qualified domain name

The sequence of labels in a domain name identifying a specific host (host name) and the domain in which it exists. For example, if a device with an IP address of 10.10.10.101 has a host name of device53 and resides in the evergreen.trees.org domain, the device's fully qualified domain name is device53.evergreen.trees.org and the DNS resolution of this name is 10.10.10.101.

GARP

Generic attribute registration protocol (defined in the IEEE 802.1D-1998 standard).

GMB

Guaranteed minimum bandwidth. Provides a method for ensuring that each of a given port's outbound traffic priority queues has a specified minimum consideration for sending traffic out on the link to another device. This can prevent a condition where applications generating lower-priority traffic in the network are frequently or continually "starved" by high volumes of higher-priority traffic.

GVRP

GARP VLAN registration protocol. 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.

host

Used in traffic mirroring to refer to a traffic analyzer or IDS.

host name

The unique, leftmost label in a domain name assigned to a specific IP address in a DNS server configuration. This enables the server to distinguish a device using that IP address from other devices in the same domain. For example, in the evergreen.trees.org domain, if an IPv4 address of 10.10.100.27 is assigned a host name of accounts015 and another IP address of 10.10.100.33 is assigned a host name of sales021, the switch configured with the domain suffix evergreen.trees.org and a DNS server that resolves addresses in that domain can use the host names to reach the devices with DNS-compatible commands.

For example: `ping accounts015 traceroute accounts015`

ICMP

Internet control message protocol.

ICMP Rate-Limiting

Applies a rate-limit to all inbound ICMP traffic received on an interface, but does not limit other types of inbound traffic.

IDM	Identify-driven management.
IDS	Intrusion detection system.
IGMP	Internet group management protocol.
IP addressing	Internet protocol (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.
ISC	InterSwitch connect. A special interface that connects DTSs.
jumbo frame	An IP frame exceeding 1522 bytes in size. The maximum Jumbo frame size is 9220 bytes. (This size includes 4 bytes for the VLAN tag.)
jumbo VLAN	A VLAN configured to allow inbound jumbo traffic. All ports belonging to a jumbo and operating at 1 Gbps or higher can receive jumbo frames from external devices. If the switch is in a meshed domain, then all meshed ports (operating at 1 Gbps or higher) on the switch will accept jumbo traffic from other devices in the mesh.
KMS	Key management system.
LACP	Link aggregation control protocol.
link test	A test of the connection between the switch and a designated network device on the same LAN (or VLAN, if configured).
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 neighbor	An LLDP device that is either directly connected to another LLDP device or connected to that device by another, non- LLDP Layer 2 device (such as a hub) Note that an 802.1D-compliant switch does not forward LLDP data packets even if it is not LLDP-aware.
LLDP-aware	A device that has LLDP in its operating code, regardless of whether LLDP is enabled or disabled.
LLDP-MED	LLDP-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 standardsbased functionality.
LLDP-MED	Second description for this acronym. The TIA telecommunications standard produced by engineering subcommittee TR41.4, "VoIP Systems - IP Telephony infrastructure and Endpoints" to address needs related to deploying VoIP equipment in IEEE 802-based environments. This standard will be published as ANSI/TIA- 1057.
LLDPDU	LLDP data unit. LLDP data unitLLDP data packet are transmitted on active links and include multiple TLVs containing global and perport switch information. In this guide, LLDPDUs are termed "advertisements" or "packets".
local mirroring	The monitored (source) interface and exit port in a mirroring session are on the same switch.
local mirroring traffic destination	Port on the same switch as the source of the traffic being mirrored. See also remote mirroring traffic destination.
log throttle periods	Used to regulate (throttle) duplicate messages for recurring events.
LSA	Link-state advertisements.
MED	Media endpoint discovery/devices (see LLDPMED).
MIB	Management information base. An internal database the switch maintains for configuration and performance information.
MLD	Multicast listener discovery. IPv6 protocol used by a router to discover the presence of multicast listeners. MLD can also optimize IPv6 multicast traffic flow with the snooping feature.
MLTS	Multiline telephone system/service. A network-based and/or premises-based telephone system having a common interface with the public switched telephone system and having multiple telephone lines, common control units, multiple telephone sets, and control hardware and software.
mm	Management module.
monitored interface	The interface (port, VLAN, trunk, or mesh) on the source switch on which the inbound and/or outbound traffic to be mirrored originates , configured with one of the interface monitor or vlan monitor commands.

MPS	Maintenance power signature. The signal a PD sends to the switch to indicate that the PD is connected and requires power.
MSTP	Multiple spanning tree protocol.
MTM	Multicast traffic manager. Controls and coordinates L3 multicast traffic for upper layer protocols.
MTU	Maximum transmission unit. The maximum size IP frame the switch can receive for Layer 2 frames inbound on a port.
NANP	North American numbering plan. A ten-digit telephone number format where the first three digits are an area code and the last seven-digits are a local telephone number.
Neighbor	See LLDP Neighbor.
non-LLDP device	A device that is not capable of LLDP operation.
nonstop switching	The standby management module is synced continuously with the active management module so that all features and config files are the same on both management modules. The standby management module is ready to become the active management module. The transition is quick and seamless; switching continues without interruption.
offline management module	A management module that is offline because Management Module redundancy is disabled.
OOBM	out-of-band management
OSPF	Open short path first. A routing protocol that uses link-state advertisements (LSA) to update neighboring routers regarding its interfaces and information on those interfaces. Each routing switch maintains an identical database that describes its area topology to help a router determine the shortest path between it and any neighboring router.
oversubscribed	The state where there are more PDs requesting PoE power than can be accommodated.
oversubscribed queue	The condition where there is insufficient bandwidth allocated to a particular outbound priority queue for a given port. If additional, unused bandwidth is not available, the port delays or drops the excess traffic.
PCM	HP Switch Manager. Windows-based network management solutions for managing and monitoring performance of HP devices.
PCM (+)	HP Switch Manager Plus. See PCM.
PD	Powered device. An IEEE 802.3af-compliant or IEEE 802.3at-compliant device that receives its power through a direct connection to a 10/ 100Base-TX PoE RJ-45 port in an HP fixed-port or chassis-based switch. Examples of PDs include Voice-over-IP (VoIP) telephones, wireless access points, and remote video cameras.
PIM	Protocol-independent multicast (routing). Enables IP multicast traffic to be transmitted for multimedia applications throughout a network without being blocked at routed interface (VLAN) boundaries.
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).
PLC	Physical layer classification.
PoE	Power-over-ethernet. The method by which PDs receive power from a PoE module (operates according to the IEEE 802.3af standard).
PoE + (POEP)	Power-over-ethernet plus. The method by which PDs receive power according to the 802.3at standard.
port-number priority	The type of power prioritization where, within a priority class, a PoE module assigns the highest priority to the lowestnumbered port in the module, the secondhighest priority to the second lowestnumbered port in the module, and so on.
primary image	The software version stored in primary flash on each management module.
priority class	The type of power prioritization that uses Low (the default), High, and Critical priority assignments to determine which groups of ports will receive power.
PSAP	Public safety answering point. Typically, emergency telephone facilities established as a first point to receive emergency (911) calls and to dispatch emergency response services such as police, fire and emergency medical services.

PSCP	PuTTY SCP (see SCP).
PSE	Power-sourcing equipment/entity. A PSE, such as a PoE module installed in a switch, provides power to IEEE 802.3af-compliant or IEEE 802.3at-compliant PDs directly connected to the ports on the module.
QoS	Quality-of-service. 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.
RADIUS	Remote authentication dial-in user service.
rapid switchover stale timer	Allows configuration of a timer (in seconds) for Layer 3 forwarding of packets. After failover, the route and neighbor entries in the forwarding information base (FIB) on the active management module are marked as stale. As new routes are added, the stale flag is reset. This continues for the number of seconds indicated by the timer, after which all remaining stale entries are removed.
remote mirroring	The monitored (source) interface and exit port in a mirroring session are on different switches. For remote mirroring, you must always configure the IP destination address and exit port (the remote mirroring endpoint) before you configure the monitored interface, by using the following commands: On the remote (destination) switch: <code>mirror endpoint ip src-ip src-udp-port dst-ip exit-port</code> On the local (source) switch: <code>mirror session remote ip src-ip src-udp-port dst-ip [truncation]</code>
RMON	Remote monitoring.
SA/DA	Source address/destination address.
SCP	Secure copy.
secondary image	The software version stored in secondary flash on each management module.
selftest	A test performed at boot to ensure the management module is functioning correctly. If the module fails selftest, it does not go into active or standby mode. If both modules fail selftest, the switch does not boot.
sFlow	Flow sampling. An industry standard sampling technology, defined by RFC 3176, used to continuously monitor traffic flows on all ports providing network-wide visibility into the use of the network.
sFlow agent	A software process that runs as part of the network management software within a device. The agent packages data into datagrams that are forwarded to a central data collector.
sFlow destination	The central data collector that gathers datagrams from sFlow-enabled switch ports on the network. The data collector decodes the packet headers and other information to present detailed Layer 2 to Layer 7 usage statistics.
SFTP	Secure ftp (file transfer protocol).
SM	Standby management module.
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.
SNTP	Simple network time protocol. Synchronizes and ensures a uniform time among interoperating devices.
source switch	The source switch on which the inbound and/or outbound traffic to be mirrored originates. See also Monitored Interface.
spoofed ping	An ICMP echo request packet intentionally generated with a valid source IP address and an invalid destination IP address. Spoofed pings are often created with the intent to oversubscribe network resources with traffic having invalid destinations.
SSH	Secure shell. Provides remote access to management functions on a switch via encrypted paths between the switch and management station clients capable of SSH operation.
SSL	Secure socket layer.
SSM	System support modules.
standard MTU	An IP frame of 1522 bytes in size. (This size includes 4 bytes for the VLAN tag.)

standby management module	A management module that is ready to become the active management module if the active management module fails.
STP	Spanning tree protocol.
switchover	When the other management module becomes the active management module.
syslog	Debug/system logging feature.
TCP	Transmission control protocol. A transport protocol that runs on IP and is used to set up connections. See also UDP.
TFTP	Trivial file transfer protocol. Supports the download of files to the switch from a TFTP network server.
threshold	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.
TLV	Type-length-value. A data unit that includes a data type field, a data unit length field (in bytes), and a field containing the actual data the unit is designed to carry (as an alphanumeric string, a bitmap, or a subgroup of information). Some TLVs include subelements that occur as separate data points in displays of information maintained by the switch for LLDP advertisements. (That is, some TLVs include multiple data points or subelements.)
ToS	Type of service.
traffic mirroring	Intelligent mirroring.
trap receiver	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.
trunk group	A set of up to eight ports configured as members of the same port trunk.
TTL	Time-to-live.
UDLD	Uni-directional link detection. 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.
UDP	See TCP.
VoIP	Voice-over IP.
VRRP	Virtual router redundancy protocol. Provides dynamic failover support as backup for gateway IP addresses (first-hop routers) so that if a VR's Master router becomes unavailable, the traffic it supports will be transferred to a backup router without major delays or operator intervention, eliminating single-point-of-failure problems.
VT	Virus throttling.
warm reboot	Binary transfer feature that supports the download of software files from a PC or Unix workstation.
warm standby	The active management module does not sync continuously with the standby management module. The standby management module boots to a certain point, syncs basic files, and only finishes booting if the active management module fails or you choose to change which module is the active management module. The transition is not seamless or immediate.
Xmodem	Binary transfer feature that supports the download of software files from a PC or Unix workstation.

Index

Symbols

- 802.1X
 - effect, LLDP, 174
 - LLDP blocked, 147
- 802.1X access control
 - authentication failure, SNMP notification, 134
 - SNMP notification of authentication failure, 134
- =prompt, 287

A

- access
 - manager, 124
 - operator, 124
- ACL
 - debug messages, 258
 - dynamic port ACL, 300, 302
 - gateway fails, 224
 - transferring command files, 200
 - troubleshooting, 223
- ACL, IPv4
 - limit, 307
 - RADIUS-assigned, limit, 307
 - scalability, 307
- ACL, IPv6
 - limit, 307
 - RADIUS-assigned, limit, 307
 - scalability, 307
- address
 - network manager, 118
- address table, port, 212
- address, network manager, 118
- advertise location, 159
- ARP
 - maximums, 307
- ARP protection
 - SNMP notification, 127, 133
- authentication
 - notification messages, 127, 133
 - SNTP, 28
 - SNTP client, 25
- authentication trap, 128
- authorized IP managers
 - SNMP, blocking, 117
- auto MDI/MDI-X
 - configuration, display, 52
 - operation, 51, 52
 - port mode, display, 52
- Auto-10, 82, 83, 91
- auto-TFTP, 187
 - disable, 188, 189
 - disabled, 187
 - download to a redundant management system, 186
 - downloading software images, 187
- autonegotiate, 159

B

- bandwidth
 - displaying port utilization, 46
- blue locator LED, 293
- boot ROM console, 183
- Bootp
 - effect of no reply, 222
- Bootp/DHCP, LLDP, 153
- broadcast
 - limit, 41
- broadcast mode
 - SNTP, 28
- broadcast storm, 81, 230
- broadcast traffic
 - IPX, 41
 - RIP, 41

C

- CDP, 174, 175, 176, 178
- chassislocate
 - LED, 293
- clear
 - statistics
 - global, 212
 - ports, 212
- CLI
 - context level, 48
- command line interface, 48
- Command Syntax
 - logging filter, 181
- communities, SNMP
 - viewing and configuring with the menu, 125
- configuration
 - copying, 198
 - impacts of software download on, 183
 - port, 40
 - port trunk group, 81
 - port, duplex, 47
 - port, speed, 47
 - restoring factory defaults, 286
 - SNMP, 118, 123
 - SNMP communities, 125
 - traffic mirroring, 205
 - transferring, 198
 - trap receivers, 128
- configuration file
 - browsing for troubleshooting, 278
- configuration file, multiple
 - copy via tftp, 198
 - copy via Xmodem, 199
- console
 - measuring network activity, 222
 - status and counters menu, 205
 - troubleshooting access problems, 221
- copy

- command output, 203
- config
 - oobm, 198
- crash data, 204
- event log output, 203
- show tech, 281
- tftp
 - show-tech, 199
 - tftp show-tech, 199
- CPU utilization, 206
- cpu utilization data, 207
- customizing, show command output, 44

D

- date format, events, 245
- debug
 - acl messages, 258
 - compared to event log, 257
 - destination, logging, 258
 - displaying debug configuration, 261
 - filtering messages, 265
 - forwarding IPv4 messages, 259
 - lldp messages, 259
 - message filtering, 265
 - overview, 257
 - packet messages, 259
 - sending event log messages, 257
 - standard event log messages, 258
 - using CLI session, 258
- debug command
 - all, 258, 264
 - cdp, 258, 264
 - configuring debug/Syslog operation, 260
 - destination, 264
 - destinations, 258, 266
 - event log, 272
 - event log as default, 258
 - event types supported, 257
 - ip, 258
 - ip fib, 258
 - ip pim, 259
 - ipv6 dhcpv6-client, 259
 - ipv6 dhcpv6-relay, 259
 - ipv6 forwarding, 259
 - ipv6 nd, 259
 - lldp, 265
 - operating notes, 272
 - rip, 264
 - security, 259
 - services, 259
 - show debug, 261
 - snmp, 259, 265
 - support for "debug" severity on Syslog servers, 267, 273
 - syntax, 258, 263
 - using CLI session, 266
- default settings, 16, 48, 49, 52, 59, 75, 155, 187
- ping, 274
- default trunk type, 86

- DHCP
 - address problems, 222
 - effect of no reply, 222
- DHCP snooping
 - SNMP notification, 127, 133
- DHCP/Bootp, LLDP, 153
- DHCPv6
 - client, 259
 - debug messages, 259
- diagnostics tools, 273
 - browsing the configuration file, 278
 - displaying switch operation, 279, 281
 - ping and link tests, 273
 - traceroute, 276
 - viewing switch operation, 278
- DNS
 - configuration, 290, 291
 - configuration, viewing, 292
 - DNS-compatible commands, 290
 - domain name, fully qualified, 289, 292
 - event log messages, 293
 - example, 291
 - IPv6 DNS resolution, 288
 - operating notes, 293
 - ping, 290
 - resolver, 288
 - resolver operation, 289
 - secure management VLAN, 293
 - server address, DHCP not used, 293
 - server IP address, 289, 293
 - three entries supported, 290
 - traceroute, 290
 - VLAN, best route selection, 293
- documentation
 - latest versions, 15
 - release notes, 15
- dot11TlvEnable, 155
- download, 195
 - software using TFTP, 183
 - switch-to-switch, 195
 - TFTP, 184
 - troubleshooting, 185
 - Xmodem, 194
- duplex advertisements, 155
- duplex information, displaying, 169
- duplicate MAC address, 233
- Dyn1, 84

E

- edge ports, 100
- Emergency Location Id Number, 164
- event log, 15
 - compared to debug/Syslog operation, 257
 - debugging by severity level, 258, 267
 - debugging by system module, 258, 267
 - generated by system module, 245
 - how to read entries, 244
 - listing entries, 254
 - losing messages, 244

- navigation, [252](#)
- not affected by debug configuration, [272](#)
- security levels, [129](#)
- sending event log messages as traps, [129](#)
- sending messages to Syslog server, [258](#)
- severity level, [245](#), [271](#)
- system module, [272](#)
- time format, [245](#)
- used for debugging, [258](#)
- used for troubleshooting, [244](#)

excessive frames, [116](#)

F

facility

- logging, [258](#)

factory default configuration

- restoring, [286](#)

failover, locator LED, [294](#)

failover, management module, locator LED, [294](#)

failure, switch software download, [185](#)

fan failure, [235](#)

fault-finder

- transceiver link-flap, [235](#)
- transceiver sensitivities, [235](#)
- warn and disable, [236](#)
- Web interface, [236](#)

fault-tolerance, [82](#)

fiber optics, monitoring links, [58](#)

filter, source-port

- jumbo VLANs, [115](#)

firmware version, [206](#)

flow control

- constraints, [41](#), [48](#)
- effect on rate-limiting, [102](#), [107](#)
- global, [48](#)
- global requirement, [41](#)
- jumbo frames, [115](#)
- per-port, [41](#), [48](#)

flow sampling, [117](#)

friendly port names, [53](#)

G

gateway

- routing fails, [224](#)

giant frames, [116](#)

H

Help

- for CLI, [313](#)
- for menu interface, [313](#)

HP

- Auto-MDIX feature, [51](#)

I

ICMP rate-limiting, [104](#)

- all-traffic, [104](#)
- caution, [103](#)
- configuring, [104](#)
- current rate-limit configuration, [106](#)

- effect of flow control, [107](#)
- effect on port trunks, [106](#)
- effects of, [104](#)
- event log messages, [108](#)
- interface support, [106](#)
- monitoring/mirroring, [106](#)
- network application, [104](#)
- no meshing, [106](#)
- note on testing, [107](#)
- operating notes, [106](#)
- operation, [103](#), [104](#)
- optimum packet size, [107](#)
- spoofed ping, [104](#)

IDM

- resources, [301](#), [303](#)

IEEE 802.1d, [230](#)

IEEE P802.1AB/D9, [146](#)

IGMP

- host not receiving, [225](#)
- not working, [225](#)
- statistics, [215](#)

Inbound Telnet Enabled parameter, [221](#)

include-credentials, SNMP, [30](#)

informs

- sending to trap receiver, [129](#)
- SNMP, [130](#)

IP

- address maximums, [307](#)
- duplicate address, [222](#)
- duplicate address, DHCP network, [222](#)
- time server address, [20](#), [34](#)

IP address

- for SNMP management, [117](#)

IP routing

- debug messages, [258](#)

IPv4

- static route, maximum, [307](#)

IPv6

- debug dhcpv6 messages, [259](#)
- static route, maximum, [307](#)

IPX

- broadcast traffic, [41](#)
- network number, [208](#)

J

jumbo frames

- configuration, [111](#)
- excessive inbound, [115](#)
- flow control, [115](#)
- GVRP operation, [111](#)
- management VLAN, [115](#)
- maximum size, [110](#), [113](#)
- MTU, [110](#)
- port adds and moves, [111](#)
- port speed, [111](#)
- security concerns, [115](#)
- standard MTU, [111](#)
- through non-jumbo ports, [115](#)
- traffic sources, [111](#)

- troubleshooting, 116
- VLAN tag, 110
- voice VLAN, 115

L

LACP

- 802.1X not allowed, 93
- active, 90
- blocked ports, 95
- default port operation, 93
- described, 83, 91
- Dyn1, 84
- dynamic, 91
- full-duplex required, 82, 91
- IGMP, 94
- no half-duplex, 95
- operation not allowed, 226
- overview of port mode settings, 82
- passive, 90
- removing port from active trunk, 90
- restrictions, 93
- standby link, 92
- status, terms, 93
- STP, 94
- trunk limit, 91
- VLANs, 94
- with 802.1X, 94
- with port security, 94

Layer-3

- scalability, 307

link failures

- detecting, 58

- link speed, port trunk, 82

- link test, 273

- link-change traps, 127, 135

LLDP

- 802.1D-compliant switch, 173
- 802.1X blocking, 147
- 802.1X effect, 174
- advertisement content, 153
- advertisement, mandatory data, 153
- advertisement, optional data, 154
- advertisements, delay interval, 150
- CDP neighbor data, 175
- chassis ID, 153
- chassis type, 153
- clear statistics counters, 171
- comparison with CDP data fields, 175
- configuration options, 144
- configuring optional data, 154
- data options, 145
- data read options, 146
- debug messages, 258, 259
- default configuration, 144
- DHCP/Bootp operation, 147
- display neighbor data, 170
- enable/disable, global, 149
- general operation, 143
- global counters, 172

- holdtime multiplier, 150
- hub, packet-forwarding, 144
- IEEE P802.1AB/D9, 146
- inconsistent value, 151
- information options, 145
- invalid frames, 172
- IP address advertisement, 147, 173
- IP address subelement, 153
- IP address, DHCP/Bootp, 153
- IP address, options, 153
- IP address, version advertised, 153
- mandatory TLVs, 174
- MIB, 143, 146
- neighbor data remaining, 174
- neighbor data, displaying, 170
- neighbor statistics, 172
- neighbor, maximum, 173
- operating rules, 146
- operation, 143
- outbound packet options, 145
- packet boundaries, 144
- packet dropped, 144
- packet time-to-live, 146
- packet-forwarding, 144
- per-port counters, 172
- port description, 154
- port ID, 153
- port speed, 155
- port trunks, 147
- port type, 153
- refresh interval, 149
- reinitialization delay, 151
- remote management address, 146
- remote manager address, 153
- reset counters, 171
- setmib, delay interval, 150
- setmib, reinit delay, 151
- show commands, 147, 148
- show outbound advertisement, 167
- SNMP notification, 145
- SNMP traps, 145
- spanning-tree blocking, 147
- standards compatibility, 146
- statistics, 171
- statistics, displaying, 171
- system capabilities, 154
- system description, 154
- system name, 154
- Time-to-Live, 144
- time-to-live, 144, 150
- transmission frequency, 144
- transmission interval, change, 149
- transmit and receive, 144
- transmit/receive modes, 144
- trap notice interval, 152
- trap notification, 152
- trap receiver, data change notice, 152
- TTL, 144, 146
- VLAN, untagged, 174

- walkmib, 146
- with PoE, 69
- lldp
 - port vlan ID support, 155
- LLDP-MED
 - displaying speed, 169
 - ELIN, 164
 - enable or disable, 144
 - endpoint support, 159
 - fast start control, 161
 - location data, 164
 - medTlvenable, 162
 - Neighbors MIB, 170
 - topology change notification, 160
 - Voice over IP, 158
- load balancing, 81, 98
- logging
 - facility, 258
 - priority-desc, 258
 - udp, 269
- logging command, 263
 - syntax, 258, 267
- logical port, 85
- loop, network, 81

M

- MAC address, 206
 - displaying detected devices, 295
 - duplicate, 230, 233
 - port, 295
 - VLAN, 296
- Management Information Base, 117
- management module failover, locator LED, 294
- management VLAN, 117
 - DNS, 293
- manager access, 124
- manager password
 - SNMP notification, 127, 133
- max frame size, jumbo, 113
- maximums, 308
- MDI/MDI-X
 - configuration, display, 52
 - operation, 51
 - port mode, display, 52
- media type, port trunk, 82
- MIB
 - HP proprietary, 117
 - listing, 117
 - standard, 117
- monitoring
 - links between ports, 58
 - locator LED, 293
- multicast, 110
- multiple forwarding database, 208, 212
- multiple VLAN, 117

N

- navigation, event log, 253
- network management functions, 118, 124

- network manager address, 118
- network slow, 222
- notifications
 - authentication messages, 127, 133
 - configuring trap receivers, 128
 - link-change traps, 127
 - network security, 133

O

- oobm
 - copy config to remote host, 198
- operating system, 183
- operation not allowed, LACP, 226
- operator access, 124
- OS, 183
 - version, 196

P

- packet
 - debug messages, 259
- password
 - SNMP notification, 133
 - SNMP notification for invalid login, 127
- pattern matching, show command output, 283
- ping, 274, 288, 290
- ping test, 273
- PoE
 - advertisement of power, 71
 - advertisements, 163
 - allocate-by, 64
 - benefit of LLDP-MED, 158
 - configuration options, 63
 - detection status, 76
 - displaying power status, 74
 - DLC, 70
 - enable or disable operation, 65
 - EPS, defined, 63
 - Event Log messages, 80
 - fault, 68
 - IEEE 802.3at stdn, 70
 - LLDP detection, enabling or disabling, 70, 71
 - lldp negotiation, 70
 - manually configuring power levels, 67
- MPS
 - absent cnt, 78
 - needed power for PoE+, 64
 - other fault, 77
 - over current cnt, 77
 - overview of status, 75
 - PD support, 64
 - poe-lldp-detect command, 69
 - power denied cnt, 77
 - priority critical, 66
 - priority high, 66
 - priority low, 66
 - priority, port, 64
 - RPS, defined, 63
 - setting allocation, 67
 - short cnt, 78

- status, 161
- status on specific ports, 77
- terminology, 63
- threshold, power, 69
- TLVs, 70, 71
- usage, 64
- using LLDP, 69
- policy enforcement engine
 - described, 300, 302
 - displaying resource usage, 300, 302
- poll interval, 17
- port
 - address table, 212
 - blocked by UDLD, 59
 - configuration, 40
 - configuring UDLD, 59
 - context level, 48
 - counters, 209
 - counters, reset, 210
 - enabling UDLD, 60
 - fiber-optic, 41
 - MAC address, 295, 296
 - menu access, 41
 - traffic patterns, 209
 - transceiver status, 46
 - trunk, 83
 - utilization, 46
 - CLI, 46
- port configuration, 81
- port names, friendly
 - configuring, 54
 - displaying, 55
 - summary, 53
- port security
 - port trunk restriction, 82
 - trunk restriction, 85
- port trunk, 81, 82
 - bandwidth capacity, 81
 - caution, 81, 85, 90
 - default trunk type, 86
 - enabling UDLD, 60
 - IGMP, 85
 - limit, 81
 - limit, combined, 91
 - link requirements, 82
 - logical port, 85
 - media requirements, 83
 - media type, 82
 - monitor port restrictions, 85
 - nonconsecutive ports, 81
 - port security restriction, 85
 - removing port from static trunk, 89
 - requirements, 84
 - spanning tree protocol, 84
 - static trunk, 84
 - static trunk, overview, 82
 - static/dynamic limit, 91
 - STP, 84
 - STP operation, 84
 - traffic distribution, 84
 - Trk1, 84
 - trunk (non-protocol) option, 83
 - trunk option described, 96
 - types, 83
 - UDLD configuration, 59
 - VLAN, 85
 - VLAN operation, 84
- port trunk group
 - interface access, 81
- port-based access control
 - event log, 226
 - LACP not allowed, 93
 - troubleshooting, 226
- port-utilization and status displays, 46
- power levels, configuring, 67
- power-over-ethernet, 63
- ProCurve
 - Auto-MDIX feature, 51
 - HP, URL, 117
 - switch documentation, 15
- ProCurve Manager
 - security concerns when deleting public community, 118
 - SNMP and network management, 117
- prompt, =, 287
- public SNMP community, 118, 124

Q

- QoS, 300, 302

R

- RADIUS-assigned ACLs
 - resources, 300, 302
- rate display for ports, 46
- rate-limiting, 110
 - bcast command, 109
 - broadcast traffic, 109
 - caution, 100
 - configuration, 100, 104
 - disabling multicast, 110
 - displaying configuration, 101, 106
 - edge ports, 100
 - effect of flow control, 102, 107
 - effect on port trunks, 102, 106
 - ICMP, 100
 - ICMP operation, 103
 - intended use, 100
 - mcast command, 109
 - multicast traffic, 109
 - note on testing, 103, 107
 - operating notes, 102
 - optimum packet size, 102, 107
 - per-port only, 100
 - purpose, 100
 - traffic filters, 102
- redundancy
 - locator LED, 294
- reset
 - port counters, 210

- resetting the switch
 - factory default reset, 286
- resource monitor
 - event log, 301, 304
- resource usage
 - displaying, 298, 302
 - insufficient resources, 301, 303
- restricted write access, 124
- RFCs, 117
 - RFC 1493, 117
 - RFC 1515, 117
 - RFC 2737, 146
 - RFC 2863, 146
 - RFC 2922, 146
- RIP
 - broadcast traffic, 41
 - debug messages, 259
- RMON, 117
- RMON groups supported, 140
- router
 - maximum routes, 307
 - OSPF area maximum, 307
 - OSPF interface maximum, 307
 - RIP interface maximum, 307
 - supported routes, 307
- routing
 - gateway fails, 224
 - traceroute, 276
- S**
- savepower
 - led option, 309
- scalability, 307
- SCP/SFTP
 - enabling, 189
 - session limit, 192, 193
 - troubleshooting, 192
- secure copy, 189
- secure FTP, 189
- secure management VLAN, DNS, 293
- security
 - enabling network security notifications, 133
- Self Test LED
 - behavior during factory default reset, 287
- serial number, 206
- setmib
 - delay interval, 150
 - reinit delay, 151
- severity level
 - event log, 245
 - selecting Event Log messages for debugging, 271
- sFlow, 117
 - CLI-owned versus SNMP-owned configurations, 141
 - configuring via the CLI, 141
 - sampling-polling information, 142
 - show commands, 141
- show
 - custom option, 44
 - displaying specific output, 283
 - pattern matching with, 283
 - show cpu, 207
 - show debug, 261
 - show interfaces
 - dynamic display, 43
 - show interfaces display, 283
 - show management, 34
 - show tech, 199, 279
 - slow network, 222
- SNMP, 117
 - ARP protection events, 127
 - authentication notification, 127, 133
 - CLI commands, 124
 - communities, 118, 124
 - configuring with the menu, 125
 - mapping, 123
 - configure, 118
 - configuring security groups, 131
 - configuring SNMPv3 notification, 131
 - configuring SNMPv3 users, 131
 - configuring trap receivers, 128
 - configuring trap receivers, 128
 - DHCP snooping events, 127
 - different versions, 127
 - enabling informs, 130
 - enabling SNMPv3, 131
 - fixed traps, 128
 - invalid password in login, 127
 - IP, 117
 - link-change traps, 127, 135
 - manager password change, 127
 - network security notification, 133
 - notification, LLDP
 - SNMP notification, 145
 - public community, 118, 124
 - supported notifications, 127
 - system thresholds, 128
 - traps, 59, 117, 127
 - walkmib, 297
 - well-known traps, 128
- SNMP trap, LLDP, 152
- SNMPv3
 - "public" community access caution, 119
 - access, 118
 - assigning users to groups, 120
 - communities, 123
 - enable command, 119
 - enabling, 118
 - group access levels, 122, 123
 - groups, 122
 - network management problems with snmpv3 only, 119
 - restricted-access option, 119
 - set up, 118
 - users, 118
- SNTP
 - authentication command, 28
 - authentication mode, 26
 - broadcast
 - mode, 28

- broadcast mode, [16](#), [20](#)
- broadcast mode, requirement, [16](#)
- client authentication, [25](#)
- configuration, [17](#)
- disabling, [23](#)
- display config information, [29](#)
- display statistics, [29](#)
- event log messages, [39](#)
- include-credentials, [30](#)
- key-id, [26](#), [27](#)
- key-value, [26](#)
- menu interface operation, [38](#)
- operating modes, [16](#)
- poll interval, [24](#)
- priority, [24](#), [28](#)
- server priority, [24](#)
- show authentication, [29](#)
- trusted key, [27](#)
- unicast mode, [16](#), [22](#), [28](#)
- unicast time polling, [37](#)
- unicast, replacing servers, [37](#)
- viewing, [17](#), [19](#)
- software, [183](#)
- software image, [183](#)
- software version, [206](#)
- source port filters
 - jumbo VLANs, [115](#)
- spanning tree
 - fast-uplink, troubleshooting, [230](#)
 - problems related to, [230](#)
 - show tech, copy output, [280](#)
 - using with port trunking, [84](#)
- SSH
 - file transfer, [188](#)
 - TACACS exclusion, [191](#)
 - troubleshooting, [192](#), [231](#)
- standard MIB, [117](#)
- static route, maximum, [307](#)
- statistics
 - clearing, [212](#)
 - SNTP, [29](#)
- switch software, [183](#)
 - download using TFTP, [183](#)
 - download, failure indication, [185](#)
 - download, troubleshooting, [185](#)
 - download, using TFTP, [183](#)
 - software image, [183](#)
 - version, [185](#), [194](#)
- Syslog, [257](#)
 - "debug" severity level as default, [273](#)
 - adding priority description, [271](#)
 - compared to event log, [257](#)
 - config friendly descriptions, [270](#)
 - configuring for debugging, [260](#)
 - configuring server address, [258](#)
 - configuring server IP address, [263](#)
 - configuring Syslog servers and debug destinations, [258](#)
 - control-desc, [270](#)
 - displaying Syslog configuration, [261](#)
 - logging command, [263](#), [266](#)
 - operating notes, [272](#)
 - overview, [257](#)
 - priority-descr, [271](#)
 - sending event log messages, [257](#)
 - server configuration, [268](#)
 - severity, "debug", [267](#)
 - specifying severity level events for debugging, [271](#)
 - specifying system module events for debugging, [272](#)
 - user facility as default, [273](#)
 - using event log for debugging, [258](#), [267](#)
- system module
 - selecting event log messages for debugging, [272](#)

T

- TACACS
 - SSH exclusion, [191](#)
- task monitor, [207](#)
- taskusage -d, [207](#)
- taskUsageShow, [207](#)
- Telnet
 - troubleshooting access, [221](#)
- TFTP
 - auto-TFTP, [187](#)
 - auto-TFTP feature, [187](#)
 - auto-TFTP, disable, [188](#), [189](#)
 - copy command output, [203](#)
 - copy crash data, [204](#)
 - copy event log output, [203](#)
 - copying a configuration file, [198](#)
 - copying software image, [197](#)
 - disable, [189](#)
 - disabled, [187](#)
 - download software using CLI, [186](#)
 - downloading software using console, [184](#)
 - enable client or server, [187](#)
 - enabling client functionality, [187](#)
 - enabling server functionality, [187](#)
 - switch-to-switch transfer, [195](#)
 - troubleshooting download failures, [185](#)
 - uploading an ACL command file, [201](#)
 - using to download switch software, [183](#)
- threshold setting, [118](#), [124](#)
- thresholds, SNMP, [128](#)
- throttling, broadcast/multicast traffic, [109](#)
- time format, events, [245](#)
- time protocol
 - selecting, [17](#)
- Time-to-Live
 - LLDP, [144](#)
- TimeP
 - assignment methods, [16](#)
 - disabling, [36](#)
 - poll interval, [36](#)
 - server address listing, [20](#), [33](#)
 - show management, [34](#)
 - viewing and configuring, menu, [33](#)
 - viewing, CLI, [33](#)
- timesync, disabling, [36](#)

- TLV advertisement, 155
- TLVs, mandatory, 174
- traceroute, 290
 - blocked route, 278
 - fails, 277
- traffic
 - broadcast rate-limiting, 109
 - multicast rate-limiting, 109
 - port, 209
- traffic monitoring, 117, 118, 124
- transceiver
 - error messages, 47
 - fault sensitivities, 235
 - fault-finder, 235
 - fiber-optic, 41
 - flapping, 235
 - link-flap, 235
 - view status, 46
- trap
 - CLI access, 128
 - configuring trap receivers, 128
 - security levels, 129
- trap notification, 152
- trap receiver, 118
 - configuring, 128
 - sending event log messages, 129
 - sending SNMPv2 informs, 129
 - SNMP, 128
 - up to ten supported, 128
- traps, 128
 - arp-protect, 134
 - auth-server-fail, 134
 - dhcp-snooping, 134
 - dynamic-ip-lockdown, 134
 - fixed, 128
 - link-change, 134, 135
 - login-failure-mgr, 134
 - password-change-mgr, 134
 - port-security, 134
 - snmp-authentication, 134
 - threshold, 128
- troubleshooting
 - ACL, 223
 - approaches, 220
 - browsing the configuration file, 278
 - configuring debug destinations, 258
 - console access problems, 221
 - diagnosing unusual network activity, 222
 - diagnostics tools, 273
 - displaying switch operation, 279, 281
 - DNS, 288
 - fast-uplink, 230
 - ping and link tests, 273
 - resource usage, 299, 302
 - restoring factory default configuration, 286
 - spanning tree, 230
 - SSH, 231
 - SSH, SFTP, and SCP Operations, 192
 - switch software download, 185
 - switch won't reboot, shows = prompt, 287
 - traceroute, 290
 - unusual network activity, 222
 - using CLI session, 258
 - using debug and Syslog messaging, 257
 - using the event log, 244
 - viewing switch operation, 278
 - web browser access problems, 221
- trunk, 81
 - L4 load balancing, 98
 - load balancing, 98
 - number supported, 82
- TTL
 - LLDP, 144
- U
- UDLD
 - changing the keepalive interval, 60
 - changing the keepalive retries, 60
 - configuring for tagged ports, 60
 - enabling on a port, 60
 - operation, 59
 - overview, 58
 - viewing configuration, 61
- UDP
 - logging messages, 269
- undersize frames, 116
- Uni-directional Link Detection, 58
- unicast mode
 - SNTP, 28
- unrestricted write access, 124
- unusual network activity, 222
- up time, 206
- URL
 - ProCurve, 117
- USB
 - copy command output, 203
 - copy event log output, 203
- users, SNMPv3, 120
- utilization, port, 46
- V
- version, OS, 196
- version, switch software, 185, 194
- view
 - transceiver status, 46
- VLAN
 - address, 117
 - configuring UDLD for tagged ports, 60
 - device not seen, 233
 - event log entries, 245
 - IP address maximum, 307
 - jumbo max frame size, 113
 - link blocked, 230
 - MAC address, 296
 - management and jumbo frames, 115
 - management VLAN, SNMP block, 117
 - maximums, 307
 - multiple, 117

- port configuration, [233](#)
- secure management VLAN, with DNS, [293](#)
- switch software download, [183](#)
- tagging broadcast, multicast, and unicast traffic, [233](#)

VoIP

- LLDP-MED support, [158](#)

W

- walkmib, [146](#), [297](#)

- web browser interface

 - troubleshooting access problems, [221](#)

- web site, HP, [117](#)

- write access, [124](#)

X

Xmodem

- copy command output, [203](#)

- copy crash data, [204](#)

- copy event log output, [203](#)

- copying a configuration file, [199](#)

- copying a software image, [197](#)

- uploading an ACL command file, [202](#)

- using to download switch software, [194](#)