

HP Switch Software Management and Configuration Guide K/KA.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 3500 Series
HP Switch 3500yl Series
HP Switch 3800 Series
HP Switch 5406zl Series
HP Switch 5412zl Series

HP Switch 6200yl-24G (J8992A)
HP Switch 6600 Series
HP Switch 8206zl (J9475A)
HP Switch 8212zl (J8715A/B)



© Copyright 2008, 2013 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. UNIX is a registered trademark of The Open Group.

Acknowledgments

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

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/support.

Contents

1 Time Protocols	21
General steps for running a time protocol on the switch	21
Disabling time synchronization	21
Viewing and configuring SNTP (CLI)	21
Enabling or disabling the SNTP mode (CLI)	23
Disabling the SNTP Mode	23
Configuring (enabling or disabling) the SNTP mode	23
Enabling SNTP in Broadcast Mode	24
Configuring (enabling or disabling) in Broadcast mode	24
Enabling SNTP in unicast mode (CLI)	25
SNTP unicast time polling with multiple SNTP servers	26
Changing the SNTP poll interval (CLI)	27
Changing the SNTP server priority (CLI)	27
Disabling time synchronization without changing the SNTP configuration (CLI)	27
Displaying all SNTP server addresses configured on the switch (CLI)	28
Adding and deleting SNTP server addresses	28
Configuring the key-identifier, authentication mode, and key-value (CLI)	29
Configuring a key-id as trusted (CLI)	29
Associating a key with an SNTP server (CLI)	30
Enabling SNTP client authentication (CLI)	30
Requirements to enable SNTP client authentication	30
Enabling SNTP Client Authentication	31
Configuring unicast and broadcast mode for authentication	31
Viewing SNTP authentication configuration information (CLI)	32
Viewing all SNTP authentication keys that have been configured on the switch (CLI)	32
Viewing statistical information for each SNTP server (CLI)	33
Storing security information in the running-config file (CLI)	33
Viewing and configuring SNTP (Menu)	33
Viewing the current TimeP configuration (CLI)	35
Configuring (Enabling or Disabling) the TimeP mode	36
Enabling TimeP in manual mode (CLI)	37
Enabling TimeP in DHCP Mode	37
Enabling TimeP in Manual Mode	38
Changing from one TimeP server to another (CLI)	39
Changing the TimeP poll interval (CLI)	39
Disabling time synchronization without changing the TimeP configuration (CLI)	39
Disabling the TimeP mode	40
Viewing, enabling, and modifying the TimeP protocol (Menu)	41
About SNTP time synchronization	42
About SNTP: Selecting and configuring	42
About SNTP unicast time polling with multiple SNTP servers	43
About operating with multiple SNTP server addresses configured (Menu)	44
About SNTP client authentication	44
Requirements	44
About configuring a trusted key	45
About saving configuration files and the include-credentials command	45
SNTP messages in the Event Log	46
About TimeP time synchronization	46
About TimeP: Selecting, and configuring	46
2 Port Status and Configuration	48
Viewing port status and configuration (CLI)	48

Internal port names.....	48
Services.....	49
Operator services.....	49
Manager services.....	49
Services in configure.....	50
Services boot.....	50
No services.....	51
Show services.....	51
No parameters.....	51
Show services locator.....	52
Displaying the port VLAN tagged status.....	52
Dynamically updating the show interfaces command (CLI/Menu).....	53
Customizing the show interfaces command (CLI).....	53
Viewing port utilization statistics (CLI).....	54
Operating notes for viewing port utilization statistics.....	55
Viewing transceiver status (CLI).....	55
Operating notes.....	56
Enabling or disabling ports and configuring port mode (CLI).....	56
Enabling or disabling the USB port (CLI/SNMP).....	57
Software versions K.13.XX operation.....	58
Software Version K.14.XX Operation.....	58
Enabling or disabling flow control (CLI).....	58
Configuring auto-MDIX (CLI).....	60
Viewing port configuration (Menu).....	62
Configuring ports (Menu).....	63
Configuring friendly port names (CLI).....	63
Configuring a single port name (CLI).....	64
Configuring the same name for multiple ports (CLI).....	64
Displaying friendly port names with other port data (CLI).....	64
Listing all ports or selected ports with their friendly port names (CLI).....	65
Including friendly port names in per-port statistics listings (CLI).....	65
Searching the configuration for ports with friendly port names (CLI).....	66
Configuring the type of a module.....	67
Clearing the module configuration.....	67
Configuring uni-directional link detection (UDLD) (CLI).....	67
Enabling UDLD (CLI).....	68
Changing the keepalive interval (CLI).....	68
Changing the keepalive retries (CLI).....	69
Configuring UDLD for tagged ports.....	69
Viewing UDLD information (CLI).....	69
Viewing summary information on all UDLD-enabled ports (CLI).....	70
Viewing detailed UDLD information for specific ports (CLI).....	70
Clearing UDLD statistics (CLI).....	70
Configuring UFD.....	71
About viewing port status and configuring port parameters.....	71
Connecting transceivers to fixed-configuration devices.....	71
Error messages associated with the show interfaces command.....	73
Note on using pattern matching with the show interfaces custom command.....	73
About configuring auto-MDIX.....	73
Manual override.....	74
About using friendly (optional) port names.....	74
Configuring and operating rules for friendly port names.....	74
About configuring transceivers and modules that have not been inserted.....	75
Transceivers.....	75
Modules.....	75

About clearing the module configuration.....	75
Restrictions.....	75
Uni-directional link detection (UDLD).....	76
About Configuring UDLD.....	77
About Uplink failure detection.....	77
UFD operating notes.....	78
Viewing UFD configuration.....	79
3 Power Over Ethernet (PoE/PoE+) Operation.....	80
Introduction to PoE.....	80
PoE.....	80
Disabling or re-enabling PoE port operation.....	80
Enabling support for pre-standard devices.....	80
Configuring the PoE port priority.....	80
Controlling PoE allocation.....	81
Manually configuring PoE power levels.....	82
Configuring PoE redundancy (chassis switches only).....	82
Changing the threshold for generating a power notice.....	83
Enabling or disabling ports for allocating power using LLDP.....	83
Enabling PoE detection via LLDP TLV advertisement.....	84
Negotiating power using the DLL.....	84
Initiating advertisement of PoE+ TLVs.....	85
Viewing PoE when using LLDP information.....	85
Viewing the global PoE power status of the switch.....	87
Viewing PoE status on all ports.....	88
Viewing the PoE status on specific ports.....	90
Planning and implementing a PoE configuration.....	91
Power requirements.....	91
Assigning PoE ports to VLANs.....	92
Applying security features to PoE configurations.....	92
Assigning priority policies to PoE traffic.....	92
About PoE operation.....	93
Configuration options.....	93
PD support.....	93
Power priority operation.....	94
About assigning PoE priority with two or more modules.....	94
About configuring PoE operation.....	95
About configuring thresholds for generating a power notice.....	96
About PoE/PoE+ allocation using LLDP information.....	97
LLDP with PoE.....	97
LLDP with PoE+.....	98
Overview.....	98
PoE allocation.....	98
Operation Note.....	98
4 Port Trunking.....	99
Viewing and configuring port trunk groups (CLI).....	99
Viewing static trunk type and group for all ports or for selected ports.....	99
Viewing static LACP and dynamic LACP trunk data.....	100
Configuring a static trunk or static LACP trunk group.....	100
Removing ports from a static trunk group.....	101
Enabling dynamic LACP trunk groups.....	101
Removing ports from a dynamic LACP trunk group.....	101
Setting the LACP key.....	102
Viewing existing port trunk groups (WebAgent).....	102
Viewing and configuring a static trunk group (Menu).....	102

Enabling L4-based trunk load balancing (CLI).....	104
Displaying information about trunk load balancing.....	105
Operating notes.....	106
Distributed trunking.....	106
Configuring ISC ports.....	106
Configuring distributed trunking ports.....	106
Configuring peer-keepalive links.....	107
Viewing distributed trunking information.....	107
Viewing information about the peer-keepalive configuration.....	108
Viewing the switch interconnect information.....	109
Overview of port trunking.....	109
Port connections and configuration.....	110
Port trunk features and operation.....	110
Fault tolerance	111
Trunk configuration methods.....	111
Dynamic LACP trunk.....	111
Dynamic LACP Standby Links.....	111
Displaying LACP Local Information.....	112
Displaying LACP Peer Information.....	112
Displaying LACP Counters.....	112
Using keys to control dynamic LACP trunk configuration.....	112
Static trunk.....	113
General operating rules for port trunks.....	114
About configuring a static or dynamic trunk group.....	116
About enabling a dynamic LACP trunk group.....	116
Dynamic LACP standby links.....	117
Displaying LACP local information.....	118
Displaying LACP peer information.....	118
Displaying LACP counters.....	118
Trunk group operation using LACP.....	119
Default port operation.....	121
LACP notes and restrictions.....	122
802.1X (Port-based access control) configured on a port.....	122
Port security configured on a port.....	123
Changing trunking methods.....	123
Static LACP trunks.....	123
Dynamic LACP trunks.....	123
VLANs and dynamic LACP.....	123
Blocked ports with older devices.....	123
Spanning Tree and IGMP.....	124
Half-duplex, different port speeds, or both not allowed in LACP trunks.....	124
Dynamic/static LACP interoperation.....	125
Trunk group operation using the "trunk" option.....	125
How the switch lists trunk data.....	125
Outbound traffic distribution across trunked links.....	125
Trunk load balancing using Layer 4 ports.....	126
Distributed trunking overview.....	127
Distributed trunking interconnect protocol.....	128
About configuring distributed trunking.....	129
About configuring peer-keepalive links.....	129
Maximum DT trunks and links supported.....	130
Forwarding traffic with distributed trunking and spanning tree.....	131
Forwarding unicast traffic.....	131
Forwarding broadcast, multicast, and unknown traffic	132
IP routing and distributed trunking.....	133

Distributed trunking restrictions.....	135
DT operating notes when updating software versions.....	136
5 Port Traffic Controls.....	138
Rate-limiting.....	138
Configuring rate-limiting.....	138
Displaying the current rate-limit configuration.....	138
Configuring ICMP rate-limiting.....	140
Viewing the current ICMP rate-limit configuration.....	141
Resetting the ICMP trap function of the port.....	141
Determining the switch port number used in ICMP port reset commands.....	142
Configuring a broadcast limit on the switch.....	142
Configuring inbound rate-limiting for broadcast and multicast traffic.....	143
Operating Notes.....	144
Configuring Guaranteed Minimum Bandwidth (GMB) for outbound traffic.....	144
Viewing the current GMB configuration.....	147
Configuring jumbo frame operation.....	148
Overview.....	148
Viewing the current jumbo configuration.....	148
Enabling or disabling jumbo traffic on a VLAN.....	150
Configuring a maximum frame size.....	150
Configuring IP MTU.....	150
Displaying the maximum frame size.....	150
Operating notes for maximum frame size.....	151
All traffic rate-limiting.....	151
Operating notes for rate-limiting.....	152
ICMP rate-limiting.....	153
Guidelines for configuring ICMP rate-limiting.....	154
Using both ICMP rate-limiting and all-traffic rate-limiting on the same interface.....	155
Operating notes for ICMP rate-limiting.....	156
Notes on testing ICMP rate-limiting.....	156
ICMP rate-limiting trap and Event Log messages.....	157
Guaranteed minimum bandwidth (GMB).....	157
GMB operation.....	157
Impacts of QoS queue configuration on GMB operation.....	158
Impact of QoS queue configuration on GMB commands.....	159
Jumbo frames.....	159
Operating rules.....	159
SNMP implementation.....	160
Jumbo maximum frame size.....	160
Jumbo IP MTU.....	160
Operating notes for jumbo traffic-handling.....	160
Troubleshooting.....	162
A VLAN is configured to allow jumbo frames, but one or more ports drops all inbound jumbo frames.....	162
A non-jumbo port is generating "Excessive undersize/giant frames" messages in the Event Log.....	162
6 Configuring for Network Management Applications.....	163
Command Summary.....	163
Enabling SNMPv3.....	163
Configuring users (CLI).....	163
Enabling and disabling switch for access from SNMPv3 agents.....	164
Enabling or disabling restrictions to access from only SNMPv3 agents.....	164
Enabling or disabling restrictions from all non-SNMPv3 agents to read-only access.....	164
Viewing the operating status of SNMPv3.....	164

Viewing status of message reception of non-SNMPv3 messages.....	164
Viewing status of write messages of non-SNMPv3 messages.....	164
Viewing management stations configured to access the switch and view authentication and privacy protocols of each station (CLI).....	165
Assigning users to groups (CLI).....	165
Mapping SNMPv3 communities (CLI).....	166
Listing community names and values (CLI).....	167
Configuring community names and values (CLI).....	167
Viewing and configuring non-version-3 SNMP communities (Menu).....	168
Configuring an SNMP trap receiver (CLI).....	169
Enabling SNMPv2c informs (CLI).....	170
Configuring SNMPv3 notifications (CLI).....	171
Enabling or disabling notification/traps for network security failures and other security events (CLI).....	173
Viewing the current configuration for network security notifications (CLI).....	174
Enabling Link-Change Traps (CLI).....	175
Enabling SNMP traps on running configuration changes (CLI).....	175
Enabling SNMP traps on Startup Configuration changes.....	175
Viewing configuration file change information (CLI).....	176
Configuring the source IP address for SNMP notifications (CLI).....	178
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).....	179
Viewing SNMP notification configuration (CLI).....	179
Configuring the MAC address count option.....	180
Displaying information about the mac-count-notify option.....	180
Configuring the MAC address table change option.....	181
Additional mac-notify options for per-port Mac changes.....	182
Configuring the mac-notify option at the interface context level.....	183
Viewing mac-notify traps configuration information.....	183
Configuring listening mode (CLI).....	184
Configuring sFlow (CLI).....	184
Viewing sFlow Configuration and Status (CLI).....	186
Displaying the global LLDP, port admin, and SNMP notification status (CLI).....	187
Viewing port configuration details (CLI).....	188
Enabling or disabling LLDP operation on the switch (CLI).....	189
Changing the packet transmission interval (CLI).....	189
Changing the time-to-live for transmitted advertisements (CLI).....	189
Changing the delay interval between advertisements generated by value or status changes to the LLDP MIB (CLI).....	190
Changing the reinitialization delay interval (CLI).....	190
Enabling LLDP data change notification for SNMP trap receivers (CLI).....	191
Changing the minimum interval for successive data change notifications for the same neighbor.....	191
Configuring per-port transmit and receive modes (CLI).....	191
Configuring a remote management address for outbound LLDP advertisements (CLI).....	192
Configuring support for port speed and duplex advertisements (CLI).....	193
Configuring the VLAN ID TLV.....	193
Viewing the TLVs advertised.....	194
Tracking LLDP-MED connects and disconnects—topology change notification.....	195
LLDP-MED fast start control.....	195
Enabling or Disabling medTlvEnable.....	196
Viewing PoE advertisements.....	197
Configuring location data for LLDP-MED devices.....	197
Viewing switch information available for outbound advertisements.....	200
Viewing the current port speed and duplex configuration on a switch port.....	201
Viewing advertisements currently in the neighbors MIB.....	201
Viewing LLDP statistics.....	202

Configuring CDP mode.....	204
Viewing the current CDP configuration of the switch.....	204
Viewing the current CDP neighbors table of the switch.....	205
Enabling and Disabling CDP Operation.....	206
Enabling or disabling CDP operation on individual ports.....	206
Configuring CDPv2 for voice transmission.....	207
Filtering CDP information.....	208
Filtering PVID Mismatch Log Messages.....	209
Configuring the switch to filter untagged traffic.....	209
Displaying the configuration.....	209
About using SNMP tools to manage the switch.....	210
SNMP management features.....	210
SNMPv1 and v2c access to the switch.....	211
SNMPv3 access to the switch.....	211
SNMPv3 users.....	212
About adding users.....	212
Group access levels.....	212
SNMPv3 communities.....	213
SNMP community features.....	213
SNMP notifications.....	214
Supported Notifications.....	214
General steps for configuring SNMP notifications.....	214
SNMPv1 and SNMPv2c Traps.....	215
SNMP trap receivers.....	215
SNMP trap when MAC address table changes.....	215
Show command.....	216
SNMP trap when power supply is inserted or removed.....	216
About configuring SNMP notification support.....	217
SNMPv2c informs.....	217
Network security notifications.....	217
SNMP traps on running configuration changes.....	218
Source IP address for SNMP notifications.....	218
Listening mode.....	218
Advanced management: RMON.....	219
CLI-configured sFlow with multiple instances.....	221
LLDP.....	221
General LLDP operation.....	221
LLDP-MED.....	222
Packet boundaries in a network topology.....	222
LLDP operation configuration options.....	222
Enable or disable LLDP on the switch.....	222
Enable or disable LLDP-MED.....	222
Change the frequency of LLDP packet transmission to neighbor devices.....	222
Change the Time-To-Live for LLDP packets sent to neighbors.....	222
Transmit and receive mode.....	222
SNMP notification.....	223
Per-port (outbound) data options.....	223
Remote management address.....	224
Debug logging.....	224
Options for reading LLDP information collected by the switch.....	224
LLDP and LLDP-MED standards compatibility.....	225
LLDP operating rules.....	225
Port trunking.....	225
IP address advertisements.....	225
Spanning-tree blocking.....	225

802.1X blocking.....	225
LLDP operation on the switch.....	226
Time-to-Live for transmitted advertisements.....	226
Delay interval between advertisements generated by value or status changes to the LLDP MIB.....	226
Reinitialization delay interval.....	226
SNMP notification support.....	226
Changing the minimum interval for successive data change notifications for the same neighbor.....	226
Basic LLDP per-port advertisement content.....	226
Mandatory Data.....	226
Optional Data.....	227
Support for port speed and duplex advertisements.....	227
Port VLAN ID TLV support on LLDP.....	227
SNMP support.....	227
LLDP-MED (media-endpoint-discovery).....	228
LLDP-MED endpoint support.....	228
LLDP-MED endpoint device classes.....	229
LLDP-MED operational support.....	229
Advertising device capability, network policy, PoE status and location data.....	230
Network policy advertisements.....	230
VLAN operating rules.....	230
Policy elements.....	230
PoE advertisements.....	231
Location data for LLDP-MED devices.....	231
About configuring coordinate-based locations.....	231
About displaying the current port speed and duplex configuration on a switch port.....	233
Displaying LLDP statistics.....	233
LLDP Operating Notes.....	233
Neighbor maximum.....	233
LLDP packet forwarding.....	234
One IP address advertisement per port.....	234
802.1Q VLAN Information.....	234
Effect of 802.1X Operation.....	234
Neighbor data can remain in the neighbor database after the neighbor is disconnected.....	234
Mandatory TLVs.....	234
About determining the switch port number included in topology change notification traps.....	234
LLDP and CDP data management.....	234
LLDP and CDP neighbor data.....	235
CDP operation and commands.....	236
7 Chassis Redundancy (8200zl Switches).....	237
Command Summary.....	237
Viewing management module redundancy status (CLI).....	237
Enabling or disabling redundant management (CLI).....	237
Transitioning from no redundancy to nonstop switching.....	241
Setting the Rapid Switchover Stale Timer.....	241
Directing the standby module to become active.....	241
Setting the rapid switchover stale timer (CLI).....	242
Directing the standby module to become active (CLI).....	243
Setting the active management module for next boot (CLI).....	243
Hotswapping out the active management module.....	245
Resetting the management module.....	246
Viewing management information.....	246
Viewing information about the management and fabric modules.....	247

Viewing information about the redundancy role of each management module.....	248
Viewing which software version is in each flash image.....	248
Viewing system software image information for both management modules.....	249
Viewing the status of the switch and its management modules.....	249
Standby management module commands.....	250
Viewing redundancy status on the standby module (CLI).....	250
Viewing the flash information on the standby module (CLI).....	251
Viewing the version information on the standby module (CLI).....	251
Setting the default flash for boot (CLI).....	251
Booting the active management module from the current default flash (CLI).....	252
Displaying module events.....	253
Viewing log events (CLI).....	253
Copying crash file information to another file (CLI).....	254
Viewing saved crash information.....	255
Enabling and disabling fabric modules (CLI).....	255
Overview of chassis redundancy (8200zl switches).....	256
Nonstop switching with redundant management modules.....	256
How the management modules interact.....	256
About using redundant management.....	257
Transition from no redundancy to nonstop switching.....	257
About setting the rapid switchover stale timer.....	257
About directing the standby module to become active.....	257
Nonstop switching with VRRP.....	257
Example nonstop routing configuration.....	259
Nonstop forwarding with RIP.....	260
Nonstop forwarding with OSPFv2 and OSPFv3.....	260
Enabling nonstop forwarding for OSPFv2.....	261
Configuring restart parameters for OSPFv2.....	261
Viewing OSPFv2 nonstop forwarding information.....	261
Enabling nonstop forwarding for OSPFv3.....	262
Configuring restart parameters for OSPFv3.....	262
Viewing OSPFv3 nonstop forwarding information.....	263
Hotswapping management modules.....	263
Management module switchover.....	263
Events that cause a switchover.....	263
What happens when switchover occurs.....	263
When switchover will not occur.....	264
When a management module crashes while the other management module is rebooting.....	264
About hotswapping out the active management module.....	264
When the standby module is not available.....	264
Hotswapping in a management module.....	264
Software version mismatch between active and hotswapped module.....	265
Other software version mismatch conditions.....	265
About downloading a new software version.....	266
File synchronization after downloading.....	266
Potential software version mismatches after downloading.....	266
Downloading a software version serially if the management module is corrupted.....	268
About turning off redundant management.....	268
Disable management module redundancy with two modules present.....	268
Disable management module redundancy with only one module present.....	269
Active management module commands.....	269
Show modules.....	269
CLI commands affected by redundant management.....	269
boot command.....	270
Boot and reload commands with OSPFv2 or OSPFv3 enabled.....	271

Modules operating in nonstop mode.....	271
Additional commands affected by redundant management.....	272
Using the WebAgent for redundant management.....	273
Notes on how the active module is determined.....	274
Diagram of the decision process.....	275
CLI commands and syncing.....	275
Management module redundancy features.....	276
Nonstop switching features.....	276
Unsupported zl modules.....	276
Hot swapping of management modules.....	276
Rapid routing switchover and stale timer.....	277
8 Link Aggregation Control Protocol-Multi-Active Detection (LACP-MAD).....	278
Configuration command.....	278
show commands.....	278
clear command.....	278
LACP-MAC Overview.....	278
9 Product Documentation.....	280
About your switch manual set.....	280
Printed publications.....	280
Electronic publications.....	280
A File transfers.....	281
Overview.....	281
Using TFTP to download software from a server.....	281
Downloading from a server to flash using TFTP (CLI).....	281
Enabling TFTP (CLI).....	282
Configuring the switch to download software automatically from a TFTP server using auto-TFTP (CLI).....	283
Downloading from a server to primary flash using TFTP (Menu).....	283
Enabling SCP and SFTP.....	285
Using Xmodem to download switch software from a PC or UNIX workstation.....	285
Downloading to primary or secondary flash using Xmodem and a terminal emulator (CLI).....	285
Downloading to primary flash using Xmodem (Menu).....	286
Downloading switch software using USB (CLI).....	287
Switch-to-switch download.....	288
Downloading the OS from another switch (CLI).....	288
Downloading from primary only (CLI).....	288
Downloading from either flash in the source switch to either flash in the destination switch (CLI).....	289
Switch-to-switch download to primary flash (Menu).....	289
Copying software images.....	290
TFTP: Copying a software image to a remote host (CLI).....	290
Xmodem: Copying a software image from the switch to a serially connected PC or UNIX workstation (CLI).....	290
USB: Copying a software image to a USB device (CLI).....	290
Transferring switch configurations.....	291
TFTP: Copying a configuration file to a remote host (CLI).....	291
TFTP: Copying a configuration file from a remote host (CLI).....	291
TFTP: Copying a customized command file to a switch (CLI).....	292
Xmodem: Copying a configuration file to a serially connected PC or UNIX workstation (CLI).....	292
Xmodem: Copying a configuration file from a serially connected PC or UNIX workstation (CLI).....	293
USB: Copying a configuration file to a USB device (CLI).....	293
USB: Copying a configuration file from a USB device (CLI).....	294
Transferring ACL command files.....	294
TFTP: Uploading an ACL command file from a TFTP server (CLI).....	294

Xmodem: Uploading an ACL command file from a serially connected PC or UNIX workstation (CLI).....	296
USB: Uploading an ACL command file from a USB device (CLI).....	296
Copying diagnostic data to a remote host, USB device, PC, or UNIX workstation.....	297
Copying command output to a destination device (CLI).....	297
Copying Event Log output to a destination device (CLI).....	297
Copying crash data content to a destination device (CLI).....	298
Copying crash data with redundant management (CLI).....	298
Copying crash log data content to a destination device (CLI).....	299
Copying crash logs with redundant management (8200zl) (CLI).....	299
Copying core dumps from the standby management module (8200zl switches).....	300
Flight data recorder.....	301
Enabling or disabling the USB port with the CLI.....	301
Viewing the status of the USB port (CLI).....	301
Using USB autorun.....	302
Configuring autorun on the switch (CLI).....	302
Viewing autorun configuration information.....	303
About downloading switch software.....	303
General software download rules.....	303
Troubleshooting TFTP download failures.....	303
Using SCP and SFTP.....	304
Disabling TFTP and auto-TFTP for enhanced security.....	305
Enabling SSH V2 (required for SFTP).....	307
Confirming that SSH is enabled.....	307
Disabling secure file transfer.....	307
Authentication.....	307
SCP/SFTP operating notes.....	307
Troubleshooting SSH, SFTP, and SCP operations.....	308
Broken SSH connection.....	308
Attempt to start a session during a flash write.....	309
Failure to exit from a previous session.....	309
Attempt to start a second session.....	309
About using USB to transfer files to and from the switch.....	309
Using PCM+ to update switch software.....	310
About transferring switch configurations.....	310
About transferring ACL command files.....	310
About copying diagnostic data to a remote host, USB device, PC or UNIX workstation.....	310
Behavior of autorun when USB port is disabled.....	311
Software versions K.13.XX operation.....	311
Software version K.14.XX operation.....	311
About using USB autorun.....	311
Security considerations.....	311
Troubleshooting autorun operations.....	312
USB auxiliary port LEDs.....	312
AutoRun status files.....	312
Event log or syslog.....	312
Autorun secure mode.....	313
Operating notes and restrictions.....	313
Autorun and configuring passwords.....	313
B Monitoring and Analyzing Switch Operation.....	314
Status and counters data.....	314
Accessing status and counters (Menu).....	314
Accessing system information (CLI).....	314
Locating a switch.....	315

Chassislocate at Boot.....	316
Show command.....	316
Collecting processor data with the task monitor (CLI).....	316
Accessing system information (Menu).....	317
Accessing switch management address information (CLI).....	317
Accessing switch management address information (Menu).....	317
Viewing additional component information (CLI).....	318
Viewing port status (Menu).....	319
Enabling and Disabling Compatibility Mode for v2 zl and zl modules.....	319
Viewing port status (CLI).....	320
Viewing port status (Menu).....	320
Accessing port and trunk group statistics (CLI).....	320
Viewing the port counter summary report.....	320
Viewing a detailed traffic summary for specific ports.....	320
Resetting the port counters.....	321
Viewing port and trunk group statistics (WebAgent).....	321
Accessing port and trunk statistics (Menu).....	321
Viewing the switch's MAC address tables.....	322
Accessing MAC address views and searches (CLI).....	322
Listing all learned MAC addresses on the switch, with the port number on which each MAC address was learned.....	322
Listing all learned MAC addresses on one or more ports, with their corresponding port numbers.....	322
Listing all learned MAC addresses on a VLAN, with their port numbers.....	322
Finding the port on which the switch learned a specific MAC address.....	322
Accessing MAC address views and searches (Menu).....	323
Viewing and searching per-VLAN MAC-addresses.....	323
Finding the port connection for a specific device on a VLAN.....	323
Viewing and searching port-level MAC addresses.....	324
Determining whether a specific device is connected to the selected port.....	324
Accessing MSTP Data (CLI).....	325
Viewing internet IGMP status (CLI).....	326
Viewing VLAN information (CLI).....	327
WebAgent status information.....	328
Configuring local mirroring (CLI).....	329
Local mirroring overview.....	329
Setting up local mirroring.....	330
Configuring a local mirroring session.....	330
Configuring traffic-direction criteria to select traffic.....	330
Configuring ACL criteria to select inbound traffic — deprecated.....	330
Configuring a mirroring policy to select inbound traffic.....	330
Configuring MAC-based criteria to select traffic.....	331
Configuring a remote mirroring destination (CLI).....	331
On the remote switch.....	331
On the local switch.....	331
Configuring a local mirroring destination on the local switch (CLI).....	331
Configuring monitored traffic (CLI).....	331
Configuring local mirroring (Menu).....	332
Configuring the mirroring destination on a remote switch (CLI).....	334
Configuring the mirroring source on the local switch (CLI).....	335
Configuring traffic-direction criteria to select traffic (CLI).....	335
Configuring ACL criteria to select inbound traffic.....	335
Configuring a mirroring policy to select inbound traffic (CLI).....	335
Configuring the MAC-based criteria to select traffic (CLI).....	336
Configuring a destination switch in a remote mirroring session (CLI).....	336

Configuring a source switch in a local mirroring session (CLI).....	337
Configuring a source switch in a remote mirroring session (CLI).....	337
Selecting all traffic on a port interface for mirroring according to traffic direction (CLI).....	339
Selecting all traffic on a VLAN interface for mirroring according to traffic direction (CLI).....	340
Configuring a MAC address to filter mirrored traffic on an interface (CLI).....	340
Configuring classifier-based mirroring (CLI).....	342
Applying a mirroring policy on a port or VLAN interface.....	343
Viewing a classifier-based mirroring configuration (CLI).....	343
Viewing all mirroring sessions configured on the switch (CLI).....	344
Viewing the remote endpoints configured on the switch (CLI).....	345
Viewing the mirroring configuration for a specific session (CLI).....	346
Viewing a remote mirroring session (CLI).....	346
Viewing a MAC-based mirroring session (CLI).....	347
Viewing a local mirroring session (CLI).....	347
Viewing information on a classifier-based mirroring session (CLI).....	348
Viewing information about a classifier-based mirroring configuration (CLI).....	348
Viewing information about a classifier-based mirroring configuration (CLI).....	349
Viewing information about statistics on one or more mirroring policies (CLI).....	349
Viewing resource usage for mirroring policies (CLI).....	350
Viewing the mirroring configurations in the running configuration file (CLI).....	351
Overview.....	352
Compatibility mode.....	352
Port and trunk group statistics and flow control status.....	353
About traffic mirroring.....	353
Mirroring.....	354
Mirroring destinations.....	354
Mirroring sources and sessions.....	354
Mirroring sessions.....	355
Mirroring session limits.....	355
Selecting mirrored traffic.....	355
Mirrored traffic destinations.....	356
Local destinations.....	356
Remote destinations.....	356
Monitored traffic sources.....	357
Criteria for selecting mirrored traffic.....	357
Mirroring configuration.....	357
Remote mirroring endpoint and intermediate devices.....	358
Migration to release K.12.xx.....	359
Bootimg from software versions earlier than K.12.xx.....	359
Maximum supported frame size.....	359
Frame truncation.....	359
Migration to release K.14.01 or greater.....	359
Using the Menu to configure local mirroring.....	360
Menu and WebAgent limits.....	360
Remote mirroring overview.....	360
Quick reference to remote mirroring setup.....	361
High-level overview of the mirror configuration process.....	361
Determine the mirroring session and destination.....	361
For a local mirroring session.....	361
For a remote mirroring session.....	362
Configure a mirroring destination on a remote switch.....	362
Configure a destination switch in a remote mirroring session.....	362
Configure a mirroring session on the source switch.....	362
Configure a source switch in a remote mirroring session.....	363
Configure the monitored traffic in a mirror session.....	363

Traffic selection options.....	363
Mirroring-source restrictions.....	363
About selecting all inbound/outbound traffic to mirror.....	364
Untagged mirrored packets.....	364
About using SNMP to configure no-tag-added.....	364
Operating notes.....	365
About selecting inbound traffic using an ACL (deprecated).....	365
About selecting inbound/outbound traffic using a MAC address.....	366
About selecting inbound traffic using advanced classifier-based mirroring.....	367
Classifier-based mirroring configuration.....	368
Classifier-based mirroring restrictions.....	369
About applying multiple mirroring sessions to an interface.....	370
Mirroring configuration examples.....	371
Example: Local mirroring using traffic-direction criteria.....	371
Example: Remote mirroring using a classifier-based policy.....	371
Example: Remote mirroring using traffic-direction criteria.....	373
Maximum supported frame size.....	375
Enabling jumbo frames to increase the mirroring path MTU.....	375
Effect of downstream VLAN tagging on untagged, mirrored traffic.....	376
Operating notes for traffic mirroring.....	376
Troubleshooting traffic mirroring.....	378
C Troubleshooting.....	379
Overview.....	379
Troubleshooting approaches.....	379
Browser or Telnet access problems.....	380
Cannot access the WebAgent.....	380
Cannot Telnet into the switch console from a station on the network.....	380
Unusual network activity.....	381
General problems.....	381
The network runs slow; processes fail; users cannot access servers or other devices.....	381
Duplicate IP addresses.....	381
Duplicate IP addresses in a DHCP network.....	381
The switch has been configured for DHCP/Bootp operation, but has not received a DHCP or Bootp reply.....	381
802.1Q Prioritization problems.....	382
Ports configured for non-default prioritization (level 1 to 7) are not performing the specified action.....	382
Addressing ACL problems.....	382
ACLs are properly configured and assigned to VLANs, but the switch is not using the ACLs to filter IP layer 3 packets.....	382
The switch does not allow management access from a device on the same VLAN.....	382
Error (Invalid input) when entering an IP address.....	383
Apparent failure to log all "deny" matches.....	383
The switch does not allow any routed access from a specific host, group of hosts, or subnet...	383
The switch is not performing routing functions on a VLAN.....	383
Routing through a gateway on the switch fails.....	383
Remote gateway case.....	383
Local gateway case.....	384
IGMP-related problems.....	384
IP multicast (IGMP) traffic that is directed by IGMP does not reach IGMP hosts or a multicast router connected to a port.....	384
IP multicast traffic floods out all ports; IGMP does not appear to filter traffic.....	385
LACP-related problems.....	385

Unable to enable LACP on a port with the interface <i>port-number</i> lacp command. In this case, the switch displays the following message.....	385
Mesh-related problems.....	385
Traffic on a dynamic VLAN does not get through the switch mesh.....	385
Port-based access control (802.1X)-related problems.....	385
The switch does not receive a response to RADIUS authentication requests.....	385
The switch does not authenticate a client even though the RADIUS server is properly configured and providing a response to the authentication request.....	385
During RADIUS-authenticated client sessions, access to a VLAN on the port used for the client sessions is lost.....	386
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.....	386
The supplicant statistics listing shows multiple ports with the same authenticator MAC address.....	386
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.....	386
RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch.....	386
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.....	387
A trunked port configured for 802.1X is blocked.....	387
QoS-related problems.....	387
Loss of communication when using VLAN-tagged traffic.....	387
Radius-related problems.....	387
The switch does not receive a response to RADIUS authentication requests.....	387
RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch.....	388
MSTP and fast-uplink problems.....	388
Broadcast storms appearing in the network.....	388
STP blocks a link in a VLAN even though there are no redundant links in that VLAN.....	388
Fast-uplink troubleshooting.....	388
SSH-related problems.....	389
Switch access refused to a client.....	389
Executing IP SSH does not enable SSH on the switch.....	389
Switch does not detect a client's public key that does appear in the switch's public key file (show ip client-public-key).....	389
An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages.....	389
Client ceases to respond ("hangs") during connection phase.....	389
TACACS-related problems.....	390
Event Log.....	390
All users are locked out of access to the switch.....	390
No communication between the switch and the TACACS+ server application.....	390
Access is denied even though the username/password pair is correct.....	390
Unknown users allowed to login to the switch.....	391
System allows fewer login attempts than specified in the switch configuration.....	391
TimeP, SNTP, or Gateway problems.....	391
The switch cannot find the time server or the configured gateway.....	391
VLAN-related problems.....	391
Monitor port.....	391
None of the devices assigned to one or more VLANs on an 802.1Q-compliant switch are being recognized.....	391
Link configured for multiple VLANs does not support traffic for one or more VLANs.....	391
Duplicate MAC addresses across VLANs.....	392

Fan failure.....	392
Mitigating flapping transceivers.....	392
Viewing transceiver information.....	394
Viewing information about transceivers (CLI).....	394
MIB support.....	394
Viewing transceiver information.....	394
Information displayed with the detail parameter.....	395
Testing the Cable.....	398
Viewing transceiver information for copper transceivers with VCT support.....	400
Using the Event Log for troubleshooting switch problems.....	400
Event Log entries.....	401
Displaying and navigating in the Event Log.....	409
Displaying the Event Log.....	409
Navigating in the Event Log.....	411
Clearing Event Log entries.....	411
Turning event numbering on.....	412
Using log throttling to reduce duplicate Event Log and SNMP messages.....	412
Log throttle periods.....	412
Example of event counter operation.....	413
Debug/syslog operation.....	413
Debug/syslog messaging.....	414
Debug/syslog destination devices.....	414
Debug/syslog configuration commands.....	414
Configuring debug/syslog operation.....	417
Viewing a debug/syslog configuration.....	418
Debug command.....	420
Configuring the types of debug messages that the switch can send to configured debug destinations.....	421
Filtering debug messages by debug type.....	425
Enabling or disabling syslog messaging.....	426
Logging command.....	427
Configuring a syslog server.....	428
Deleting syslog addresses in the startup configuration.....	429
Verifying the deletion of a syslog server address.....	429
Blocking the messages sent to configured syslog servers from the currently configured debug message type.....	429
Disabling syslog logging on the switch without deleting configured server addresses.....	429
Sending logging messages using TCP.....	429
Disable LinkUp/Down Syslog Messages Based on Port	431
Creating a Filter.....	432
Enabling a Filter after Creation.....	432
Clearing a Filter.....	432
Displaying Filter Configuration Information.....	432
Adding a description for a Syslog server.....	436
Adding a priority description.....	436
Configuring the severity level for Event Log messages sent to a syslog server.....	437
Configuring the system module used to select the Event Log messages sent to a syslog server.....	437
Operating notes for debug and Syslog.....	438
Diagnostic tools.....	439
Port auto-negotiation.....	439
Ping and link tests.....	439
Ping test.....	439
Link test.....	439
Executing ping or link tests (WebAgent).....	439

Testing the path between the switch and another device on an IP network.....	440
Halting a ping test before it concludes.....	442
Issuing single or multiple link tests.....	442
Tracing the route from the switch to a host address.....	443
Halting an ongoing traceroute search.....	444
A low maxttl causes traceroute to halt before reaching the destination address.....	444
If a network condition prevents traceroute from reaching the destination.....	445
Viewing Switch Configuration and Operation.....	445
Viewing the startup or running configuration file.....	445
Viewing the configuration file (WebAgent).....	446
Viewing a summary of switch operational data.....	446
Saving show tech command output to a text file.....	447
Customizing show tech command output.....	448
Viewing more information on switch operation.....	450
Searching for text using pattern matching with show command.....	450
Displaying the information you need to diagnose problems.....	453
Restoring the factory-default configuration.....	454
Resetting to the factory-default configuration.....	454
Using the CLI.....	454
Using Clear/ Reset.....	454
Zeroization.....	454
Zeroizing from the ROM console.....	455
Zeroizing with HA.....	455
Restoring a flash image.....	456
Recovering from an empty or corrupted flash state.....	456
DNS resolver.....	457
Basic operation.....	457
Configuring and using DNS resolution with DNS-compatible commands.....	458
Configuring a DNS entry.....	459
Using DNS names with ping and traceroute: Example.....	460
Viewing the current DNS configuration.....	461
Operating notes.....	461
Event Log messages.....	462
Locating a switch (Locator LED).....	462
D MAC Address Management.....	463
Overview.....	463
Determining MAC addresses.....	463
Viewing the MAC addresses of connected devices.....	463
Viewing the switch's MAC address assignments for VLANs configured on the switch.....	463
Viewing the port and VLAN MAC addresses.....	464
E Monitoring Resources.....	466
Displaying current resource usage.....	466
Viewing information on resource usage.....	467
Policy enforcement engine.....	468
Usage notes for show resources output.....	468
When insufficient resources are available.....	469
F Daylight Saving Time on HP Switches.....	470
G Scalability: IP Address, VLAN, and Routing Maximum Values.....	472
H Switch Licensing.....	474
General procedure.....	474
I Power-Saving Features.....	476
Overview.....	476

Configuring the power-saving options.....	476
Configuring the savepower module option.....	477
Configuring the savepower LED option.....	478
Configuring the savepower port-low-pwr option.....	479
show savepower commands.....	479
show savepower module.....	479
show savepower port-low-pwr.....	480
show savepower led	480
Enabling energy efficient ethernet.....	481
LLDP support for EEE.....	482
J Network Out-of-Band Management (OOBM) for the 6600 Switch.....	484
OOBM Configuration.....	484
Entering the OOBM configuration context from the general configuration context.....	484
Enabling and disabling OOBM.....	484
Enabling and disabling the OOBM port.....	484
Setting the OOBM port speed.....	485
Configuring an OOBM IPv4 address.....	485
Configuring an OOBM IPv4 default gateway.....	486
OOBM show commands.....	486
Showing the global OOBM and OOBM port configuration.....	486
Showing OOBM IP configuration.....	487
Showing OOBM ARP information.....	487
Application server commands.....	487
Application client commands.....	488
Concepts.....	490
Example.....	491
OOBM and switch applications.....	492
K Support and Other Resources.....	493
Intended audience.....	493
Related documentation.....	493
Contacting HP.....	493
HP technical support.....	493
Subscription service.....	493
Related information.....	494
HP websites.....	494
Typographical conventions.....	494
Command syntax statements.....	494
Command prompts.....	495
Screen simulations.....	495
Configuration and operation examples.....	495
Keys.....	495
To set up and install the switch in your network.....	496
Physical installation.....	496
Product warranties.....	496
Online help.....	496
Menu interface.....	496
Command-line interface.....	497
HP customer support services.....	497
Before calling support.....	497
Index.....	498

1 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.

-
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**.)

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.

Example

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:

Figure 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.

Figure 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
```

Even though, in this example, TimeP is the current time synchronous method, the switch maintains the SNTP configuration.

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.

Figure 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
```

Enabling or disabling the SNTP mode (CLI)

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.

Figure 4 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
```

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

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
```

```
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:

Figure 5 Enabling SNTP operation in Broadcast Mode

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

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

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.

Configuring (enabling or disabling) in Broadcast mode

The switch provides an SNTP polling interval (default:720 seconds.) You need the two following 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 time synchronization is in the factory-default configuration (TimeP is the currently selected time synchronization method.)

You want to:

1. View the current time synchronization: `show sntp` displays the SNTP configuration and also shows that TimeP is the currently active time synchronization mode.
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 appear as follows:

Figure 6 `show sntp` configuration output

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

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

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.

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 [“About SNTP unicast time polling with multiple SNTP servers” \(page 43\)](#)

Syntax

```
timesync sntp
```

Selects SNTP as the time synchronization method.

```
sntp unicast
```

Configures the SNTP mode for unicast operation.

Syntax

```
[no] sntp server priority [ 1-3 ] ip-address [ oobm ] [ 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.

oobm For switches that have a separate out-of-band management port, specifies that SNTP traffic goes through that port. (By default, SNTP traffic goes through the data ports.)

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.)

Figure 7 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

Figure 8 Specifying the SNTP protocol version number

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

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

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.

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*.

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 1 SNTP with time synchronization disabled

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

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 2 Example of 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 3 Example of 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 priority 1-3 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
```

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.

For detailed information on client authentication, see [“About SNMP: Selecting and configuring” \(page 42\)](#).

Syntax

```
sntp authentication key-id key-id authentication-mode md5 key-value  
key-string trusted [encrypted-key key-string]  
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 SNMP packet.
<code>key-value <i>key-string</i></code>	The secret key that is used to generate the message digest. Up to 32 characters are allowed for <i>key-string</i> .

NOTE: For the 5400zl, 3800, and 8200zl switches, when the switch is in enhanced secure mode, commands that take a secret key as a parameter have the echo of the secret typing replaced with asterisks. The input for `<key-string>` is prompted for interactively. For more information, see the *Access Security Guide*.

<code>encrypted-key <i>key-string</i></code>	Set the SNMP authentication key value using a base64-encoded aes-256 encrypted string.
----------------------------------------------	----------------------------------------------------------------------------------------

Example 4 Example of setting parameters for SNMP authentication

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

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 “About configuring a trusted key” (page 45)

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 5 Example of 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 (CLI)

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 to enable SNTP client authentication

The following must be configured 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” on page 1-19.
- 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 will be 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

The following must be performed 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 will fail on the clients otherwise, and the SNTP packets will be dropped.

NOTE: SNTP server is not supported on HP products.

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 6 Example of 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

Example 7 show sntp authentication command output

To display all the SNTP authentication keys that have been configured on the switch, enter the `show sntp authentication` command.

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

```
SNTP Authentication Information
```

```
SNTP Authentication: Enabled
```

Key-ID	Auth Mode	Trusted
55	MD5	YES
10	MD5	NO

To display the statistical information for each SNTP server, enter the `sntp statistics` command. The number of SNTP packets that have failed authentication is displayed for each SNTP server address.

```
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

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

Enter the `show sntp authentication` command, as shown in [Example 8](#).

Example 8 Example of show snmp authentication command output

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

SNMP Authentication Information

SNMP Authentication : Enabled

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

Viewing statistical information for each SNMP server (CLI)

To display the statistical information for each SNMP server, enter the `show snmp statistics` command.

The number of SNMP packets that have failed authentication is displayed for each SNMP server address, as shown in [Example 9](#).

Example 9 Example of SNMP authentication statistical information

```
HP Switch(config)# show snmp statistics
SNMP Statistics

Received Packets : 0
Sent Packets : 3
Dropped Packets : 0

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

Storing security information in the running-config file (CLI)

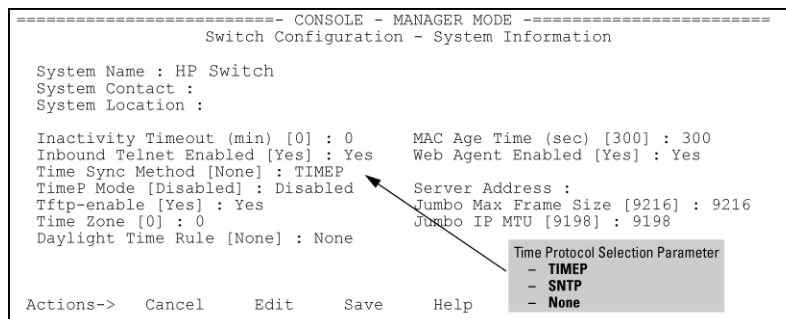
Enter the `include-credentials` command.

For more information and examples, see [“About saving configuration files and the include-credentials command”](#) (page 45).

Viewing and configuring SNMP (Menu)

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

Figure 9 System Information screen (default values)



2. Press **[E]** (for **Edit**.)
The cursor moves to the **System Name** field.
3. Use **â** to move the cursor to the **Time Sync Method** field.
4. Use the **Space** bar to select **SNTP**, then press **â** once to display and 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. Press **â** to move the cursor to the **Poll Interval** field.
- c. Go to step 6 (page 34). (For Broadcast mode details, see “About SNTP time synchronization” (page 42))

Figure 10 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. Press **â** to 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 “About SNTP unicast time polling with multiple SNTP servers” (page 43).

- g. Press **â** to 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 “About SNTP unicast time polling with multiple SNTP servers” (page 43).

- h. Press **â** to move the cursor to the **Poll Interval** field, then go to step 6.

Figure 11 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 42), on “SNTP parameters” (page 42).)

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 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.)

Example

Example 10 TimeP configuration when TimeP is the selected Time synchronization method

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:

```
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
```

Example 11 TimeP configuration when TimeP is not the selected time synchronization method

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 (see data in bold below):

```
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 12 Example of 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.

Example

Suppose time synchronization is configured for SNTP. You want to:

1. View the current time synchronization.
`show timep` displays the TimeP configuration and also shows that SNTP is the currently active time synchronization mode.
2. Select TimeP as the time synchronization mode.
3. Enable TimeP for DHCP mode.

4. View the TimeP configuration.

`show timep` again displays the TimeP configuration and shows that TimeP is now the currently active time synchronization mode.

The commands and output appear as follows:

Enabling TimeP operation in DHCP mode

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

```
Timep Configuration
```

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

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

```
HP Switch(config)# ip timep dhcp
```

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

```
Timep Configuration
```

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

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 [ oobm ]
```

Activates TimeP in manual mode with a specified TimeP server.

For switches that have a separate out-of-band management port, *oobm* specifies that SNTP traffic goes through that port. (By default, SNTP traffic goes through the data ports.)

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.

Example 13 TimeP synchronization method

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.

The commands and output would appear as follows:

Figure 12 Enabling TimeP operation in DHCP mode

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

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

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.

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> [oobm]
```

Activates TimeP in manual mode with a specified TimeP server.

For switches that have a separate out-of-band management port, oobm specifies that SNTP traffic goes through that port. (By default, SNTP traffic goes through the data ports.)

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, then 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 14 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
```

```
Poll Interval (min) : 720
```

```
Server Address : 10.28.227.141
```

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 15 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 16 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
```

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

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

Figure 13 System Information screen (default values)

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

System Name : HP Switch
System Contact :
System Location :

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

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

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

Actions->  Cancel      Edit      Save      Help
```

2. Press **[E]** (for **Edit**.)
The cursor moves to the **System Name** field.
3. Use **â** to move the cursor to the **Time Sync Method** field.
4. If **TIMEP** is not already selected, use the **Space** bar to select **TIMEP**, then press **â** once to display and move to the **TIMEP Mode** field.
5. Do one of the following:
 - Use the **Space** bar to select the **DHCP** mode.
 - Press **â** to 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.
 - Press **à** to 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.

- Press **à** to 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.

About 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.

About SNTP: Selecting and configuring

Table 1 (page 42) 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

Table 1 SNTP parameters (continued)

SNTP parameter	Operation
	the original server, the switch accepts a broadcast time update from the next server it detects.
Poll Interval (seconds)	In Unicast Mode: Specifies how often the switch polls the designated SNTP server for a time update. In Broadcast Mode: Specifies how often the switch polls the network broadcast address for a time update. Value is between 30 to 720 seconds.
Server Address	Used only when the SNTP Mode is set to Unicast . Specifies the IP address of the SNTP server that the switch accesses for time synchronization updates. You can configure up to three servers; one using the menu or CLI, and two more using the CLI.
Server Version	Specifies the SNTP software version to use and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3. Default: 3; range: 1 to 7.
Priority	Specifies the order in which the configured servers are polled for getting the time. Value is between 1 and 3.

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 (“[Viewing and configuring SNTP \(CLI\)](#)” (page 21).)

```
sntp [ broadcast | unicast ]
```

Enables the SNTP mode (page A-7 and “[Enabling SNTP in unicast mode \(CLI\)](#)” (page 25).)

Syntax

```
sntp server ip-addr
```

Required only for unicast mode (“[Enabling SNTP in unicast mode \(CLI\)](#)” (page 25).)

Syntax

```
sntp poll-interval [ 30 - 720 ]
```

Enabling the SNTP mode also enables the SNTP poll interval (default: 720 seconds; “[Changing the SNTP poll interval \(CLI\)](#)” (page 27).)

Syntax

```
sntp server priority [1 - 3 ]
```

Specifies the order in which the configured servers are polled for getting the time (page A-11.)

About 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.

About 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.

About 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” \(page 31\)](#).
- 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.

About 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" \(page 31\)](#).

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.

About 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 the *Access Security Guide*.

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.

Configuration file with SNTP authentication information example

```
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 figure.

Retrieved configuration file when `include-credentials` is not configured

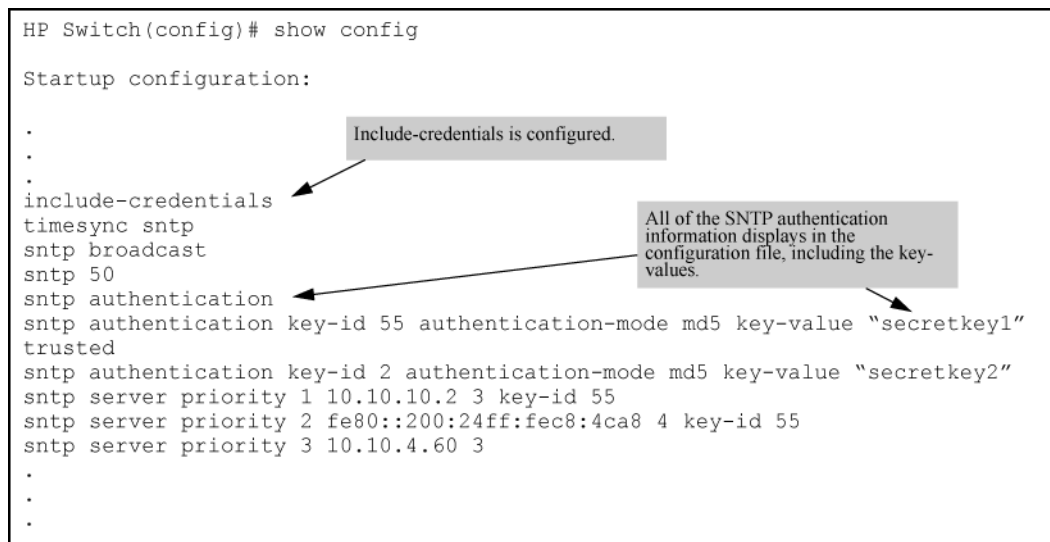
```
HP Switch (config) # copy tftp startup-config 10.2.3.44 config1
.
.
.
Switch reboots ...
.
Startup configuration
.
.
.
timesync sntp
sntp broadcast
sntp 50 sntp server priority 1 10.10.10.2.3
sntp server priority 2 fe80::200:24ff:fec8:4ca8 4
.
.
.
```

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 14 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
.
.
.
```



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.

About 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.

About TimeP: Selecting, and configuring

Table 2 (page 47) 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	
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.

Enabling the TimeP mode means to configure it for either broadcast or unicast mode. Remember that 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 protocol ([“Viewing the current TimeP configuration \(CLI\)” \(page 35\).](#))

Syntax

```
ip timep dhcp | manual
```

Enables the selected TimeP mode ([“Configuring \(Enabling or Disabling\) the TimeP mode” \(page 36\)](#) and [“Enabling TimeP in manual mode \(CLI\)” \(page 37\).](#))

Syntax

```
no ip timep
```

Disables the TimeP mode ([“Disabling the TimeP mode” \(page 40\).](#))

Syntax

```
no timesync
```

Disables the time protocol ([“Disabling time synchronization without changing the TimeP configuration \(CLI\)” \(page 39\).](#))

2 Port Status and Configuration

Viewing port status and configuration (CLI)

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

Syntax

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

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

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

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

Examples

Example 17 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 18 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

Internal port names

Both external and internal ports are supported on the same module. Internal ports have an “i” suffix to indicate that they are internal ports.

Internal-port type strings

- “10GbE-INT” – Internal 10G data-plane ports (1i-2i, 4i-5i)
- “1GbE-INT” – Internal 1G control-plane port (3i)

Port 3i always shows as link-down.

Example

```
HP-8212zl# show interfaces brief d1i-d3i
```

Status and Counters		Port Status				MDI	Flow	Bcast
Port	Type	Intrusion		Status	Mode	Mode	Ctrl	Limit
		Alert	Enabled					
D1i	10GbE-INT	No	Yes	Up	10GigFD	NA	off	0
D2i	10GbE-INT	No	Yes	Up	10GigFD	NA	off	0
D3i	1GbE-INT	No	Yes	Down	1000FDx	NA	off	0

Example

```
HP-8212zl# show interfaces brief b1-b3i
```

Status and Counters		Port Status				MDI	Flow	Bcast
Port	Type	Intrusion		Status	Mode	Mode	Ctrl	Limit
		Alert	Enabled					
B1	100/1000T	No	Yes	Down	1000FDx	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
B7	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B8	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B9	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B10	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B11	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B12	100/1000T	No	Yes	Down	1000FDx	Auto	off	0
B1i	10GbE-INT	No	Yes	Up	10GigFD	NA	off	0
B2i	10GbE-INT	No	Yes	Up	10GigFD	NA	off	0
B3i	1GbE-INT	No	Yes	Up	1000FDx	NA	off	0

Services

The `services` command requires a slot-name parameter followed a option. Options permitted in this command depend on the context (operator, manager, or configure).

Operator services

Displays parameters for the services module. Options are permitted in this command for the Operator.

Syntax

```
services <d><SLOT-ID>[index/locator/name]
```

d	Display parameters for the services module.
SLOT-ID	Device slot identifier for the services module.
<INDEX>	Index of the services CLI to access.
locator	Control services module locator LED.
name <NAME>	Name of the services CLI to access.

Manager services

Displays parameters pr changes to the services module. Options (boot, reload, serial and shutdown) are permitted in this command for the Manager.

Syntax

```
services <d><SLOT-ID>[<index>/locator/name  
<name>/boot/reload/serial/shutdown]  
d                Display parameters for the services module.  
SLOT-ID          Device slot identifier for the services module.  
<INDEX>          Index of the services CLI to access.  
locator          Control services module locator LED.  
name <NAME>      Name of the services CLI to access.  
boot             Reboot the services module.  
reload           Reset the services module.  
serial           Connect to application via serial port.  
shutdown        Shutdown (halt) the services module.
```

Services in configure

Configures the parameters for the services module or changes the module's state (reload or shutdown).

Syntax

```
services  
<SLOT-ID>[<INDEX>/boot/locator/name<NAME>/reload/serial/shutdown]  
[no] services <SLOT-ID>[device [shutdown | usb]  
d                Display parameters for the services module.  
SLOT-ID          Device slot identifier for the services module.  
<INDEX>          Index of the services CLI to access.  
locator          Control services module locator LED.  
name <NAME>      Name of the services CLI to access.  
boot             Reboot the services module.  
reload           Reset the services module.  
serial           Connect to application via serial port.  
shutdown        Shutdown (halt) the services module.  
device           Enable or disable devices.
```

Services boot

A reboot of the x86. If no parameters are giving, the command will boot to the same OS (product, service, or USB) that was enabled before.

Syntax

```
services d <boot>[product/PXE/service/USB]  
product          Boot to the Product OS.  
PXE              Boot to the PXE or Product OS (if supported).  
service          Boot to the Service OS.  
USB              Boot to the USB or Product OS (if supported).
```

Example

```
HP-8212z1# services d shutdown
```

NOTE: The `services d shutdown` command is similar to `services <slot> boot` command as it prepares a graceful shutdown of the x86. This command does not restart the x86. If the graceful-shutdown attempt fails there is no follow-up attempt made.

Error message

If the `services <slot> boot product|usb` command is used on a non-permitted module, a similar error message is returned:

```
HP-8212z1# services b boot product - Command not supported
for the Services module in slot B.
```

```
HP-8212z1# services b boot pxe - Command not supported for
the Services module in slot B.
```

```
HP-8212z1# services b boot usb - Command not supported for
the Services module in slot B.
```

No services

Enables or disables devices. This command must be run from configure context.

Syntax

```
[no]services <slot-ID> device [shutdown|USB|PXE|cf|SD2]
PXE          Enable or disable booting from PXE (if supported).
shutdown    Enable or disable the front panel shutdown or reset button.
USB         Enable or disable the module USB port after boot.
cf          Enable or disable the Compact Flash or SD card.
sd2         Enable or disable the SD2 card.
```

Show services

Show services modules information.

Syntax

```
show services [slot-id] <details|device>
```

Slot-ID	Display summary table for the specified slot.
<SLOT-ID> details	Display application information for the specified slot.
<SLOT-ID> device	Display the current configuration of the devices.
ASCII-STR	Enter an ASCII string.

```
HP-8212z1# show services
```

Slot	Installed Services Index Description	Name
H,L	1. Services z1 Module	services-module
L	2. HP ProCurve MSM765 z1 Int-Ctrlr	msm765-applicati
H	3. Threat Management Services z1 Module	tms-module

No parameters

This command with no parameters lists only installed modules which have applications running that provide a pass-through CLI feature.

Syntax

```
show services
```

Show services of only installed modules.

Example

```
HP-8212z1# show services
```

```
Installed Services

Slot      Index Description                               Name
H,L       1. Services z1 Module                       services-module
L         2. HP ProCurve MSM765 z1 Int-Ctrl          msm765-applicati
H         3.Threat Management Services z1 Module    tms-module
```

Show services locator

This command sets the Module Locator LED to either solid-on, off or slow-blink for a specified duration of time. Options are permitted in this command for the Operator.

Syntax

```
show services <locator>[blink/off/on] >
blink      Blink the locate led. Default 30 mins. Range <1-1440>.
off        Turn the locate led off.
on         Turn the locate led on.
```

Example

```
HP-8212z1# show services d locator blink
```

Displaying the port VLAN tagged status

The `show interfaces status` command displays port status, configuration mode, speed, type and tagged or untagged information.

Tagged values can be:

- VLAN ID: When the VLAN number is displayed, the port is a member of a single tagged VLAN.
- multi: When "multi" is displayed, the port is a member of multiple tagged VLANs.
- no: When "no" is displayed, the port is not a member of any tagged VLAN.

Untagged values can be:

- VLAN-ID: When the VLAN number is displayed, the port is a member of a single untagged VLAN.
- multi: When "multi" is displayed, the port is added to multiple untagged VLANs.
- no: When "no" is displayed, the port is not a member of any tagged VLAN.

If the port is part of a trunk, then the `trunk_VLAN` membership is display in the Tagged and Untagged columns.

Example

```
HP-Switch(config)# show interfaces status
Port Name Status Config-mode Speed Type Tagged Untagged
```

```

-----
A1 Up Auto 1000FDx 100/1000T 2 1
A2 Down 10HDx 10HDx 100/1000T multi 2
A3 Down 100HDx 100HDx 100/1000T multi 3
A4 Down 10FDx 10FDx 100/1000T 5 4
A5-Trk1 Down 100FDx 100FDx 100/1000T No No
A6 Down Auto 1000FDx 100/1000T No 6
A7 Down Auto-10 10HDx 100/1000T No 7

```

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 15 show interfaces display command with dynamically updating output

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

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

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

Customizing the show interfaces command (CLI)

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

Syntax

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

Select the information that you want to display. Supported columns are shown in [Table 3 \(page 54\)](#).

Table 3 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
intrusion	Intrusion alert status	no
bcast	Broadcast limit	0

Example 19 Example of 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 19 \(page 54\)](#), 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 \(page 73\)](#).

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. [Figure 16 \(page 55\)](#) shows a sample output from this command.

Figure 16 Example of 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.

Figure 17 (page 56) shows sample output from the `show tech transceivers` command.

NOTE: Part # column in Figure 17 (page 56) enables you to determine the manufacturer for a specified transceiver and revision number.

Figure 17 Example of show tech transceivers command

```
HP Switch# show tech transceivers

Transceiver Technical Information:
Port # | Type      | Prod # | Serial #      | Part #
-----+-----+-----+-----+-----
21     | 1000SX   | J4858B | CN605MP23K   |
22     | 1000LX   | J4859C | H117E7X      | 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 [Figure 17 \(page 56\)](#)), 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 4 \(page 71\)](#) ([Broadcast limit \(page 73\)](#)).

Syntax

```
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 18 Two methods for changing a port configuration



For more on flow control, see [“Enabling or disabling flow control \(CLI\)”](#) (page 58).

Enabling or disabling the USB port (CLI/SNMP)

This feature allows configuration of the USB port with either the CLI or SNMP.

To enable/disable the USB port with the CLI:

Syntax

```
usb-port
no usb-port
```

Enables the USB port. The `no` form of the command disables the USB port and any access to the device.

To display the status of the USB port:

Syntax

```
show usb-port
```

Displays the status of the USB port. It can be enabled, disabled, or not present.

Figure 19 Example of `show usb-port` command output on version K.13.59 and later

```
HP Switch(config)# show usb-port
```

```
USB port status: enabled
USB port power status: power on (USB device detected in port)
USB port reseal status: USB reseal not required
```

Figure 20 Example of `show usb-port` command output on version K.14.XX

```
HP Switch(config)# show usb-port
```

```
USB port status: enabled
USB port power status: power on      (USB device detected in port)
```

One of the following messages indicates the presence or absence of the USB device:

- Not able to sense device in USB port
- USB device detected in port
- no USB device detected in port

The reseal status messages can be one of the following (K.13.XX only):

- undetermined USB reseal requirement
- USB reseal not required
- USB device reseal required for USB autorun

The autorun feature works only when a USB device is inserted and the USB port is enabled..

Software versions K.13.XX operation

When using software version K.13.58, if the USB port is disabled (no `usb-port` command), the USB autorun function does not work in the USB port until the USB port is enabled, the config file is saved, and the switch is rebooted. The 5 volt power to the USB port remains on even after the USB port has been disabled. For software versions after K.13.58, the 5 volt power applied to the USB port is synchronized with the enabling of the USB port, that is, when the USB port is enabled, the 5 volts are supplied; when the USB port is disabled, the 5 volts are not supplied. For previous software versions the power was supplied continuously. The autorun function does not require a switch reboot, but the USB device must be inserted at least once after the port is enabled so that the switch recognizes that the device is present. If the USB device is inserted and then the USB port is enabled, the switch does not recognize that a USB device is present.

Software Version K.14.XX Operation.

For software versions K.14.XX, the USB port can be disabled and enabled without affecting the autorun feature. When the USB port is enabled, the autorun feature activates if a USB device is already inserted in the USB port. Power is synchronized with the enabling and disabling of USB ports as described above for K.13.59 and later software.

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 17 \(page 48\)](#).) 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 4 \(page 71\)](#).)

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 20 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 21 Example continued from Example 20 (page 60)

```
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 22 Example continued from Example 21 (page 60)

```
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 (CLI)

The auto-MDIX features apply only to copper port switches using twisted-pair copper Ethernet cables. For information about auto-MDIX, see [“About configuring auto-MDIX” \(page 73\)](#).

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.

Example

`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 23 Example of 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 24 Example of displaying the current MDI operating mode

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

Status and Counters - Port Status

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

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

Figure 21 Example of a switch port status screen

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

  Port      Type      Intrusion  Enabled  Status  Mode      MDI  Flow
  -----  -----  -
  E1        10/100TX  No         Yes      Down    10FDx     Auto off
  E2        10/100TX  No         Yes      Down    10FDx     Auto off
  E3        10/100TX  No         Yes      Down    10FDx     Auto off
  E4        10/100TX  No         Yes      Down    10FDx     Auto off
  E5        10/100TX  No         Yes      Down    10FDx     Auto off
  E6        10/100TX  No         Yes      Down    10FDx     Auto off
  B7-Trk2   10/100TX  No         Yes      Down    10FDx     Auto off
  B8-Trk2   10/100TX  No         Yes      Down    10FDx     Auto off
  B9        10/100TX  No         Yes      Down    10FDx     Auto off
  B10       10/100TX  No         Yes      Down    10FDx     Auto off
  B11       10/100TX  No         Yes      Down    10FDx     Auto 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.
```

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 Chapter 12, "Port Trunking".

1. From the Main Menu, select:

2. Switch Configuration...

2. Port/Trunk Settings

Figure 22 Example of 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
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 ac-
tion.
```

2. Press [E] (for Edit.)

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

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

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

Configuring friendly port names (CLI)

For detailed information about friendly port names, see ["About using friendly \(optional\) port names"](#) (page 74).

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)

Example

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 25 Example of 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
Name : Bill_Smith@10.25.101.73
```

Configuring the same name for multiple ports (CLI)

Example

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 26 Example of 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)

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

Example 27 Example of 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 28 Example of 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.

Example

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 29 Example of 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.

Example

If you configure port A1 with a friendly port name:

Figure 23 Example 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 K.15.05.xxxx
hostname "HPSwitch"
interface AQ
  name "Print_Server@10.25.101.43"
exit

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

This 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**.

Configuring the type of a module

For detailed information about configuring transceivers and modules, see [“About configuring transceivers and modules that have not been inserted”](#) (page 75).

Syntax

```
module module-num type module-type
```

Allows you to configure the type of the module.

Clearing the module configuration

For information about clearing module configuration and operational restrictions, see [“About clearing the module configuration”](#) (page 75).

Syntax

```
[no] module slot
```

Allows removal of the module configuration in the configuration file after the module has been removed. Enter an integer between 1 and 12 for *slot*.

Example

```
HP Switch(config)# no module 3
```

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

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

Syntax

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

Enables UDLD on a port or range of ports.

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

Default: UDLD disabled

Syntax

```
link-keepalive interval interval
```

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

Default: 50 (5 seconds)

Syntax

```
link-keepalive retries num
```

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

Default: 5

Syntax

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

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

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

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

Enabling UDLD (CLI)

UDLD is enabled on a per-port basis.

Example

To enable UDLD on port a1, enter:

```
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 4 \(page 71\)](#) 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 24 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 25 Example of `show link-keepalive statistics` command

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

Port:	1		
Current State:	up	Neighbor MAC Addr:	0000a1-b1c1d1
Udld Packets Sent:	1000	Neighbor Port:	5
Udld Packets Received:	1000	State Transitions:	2
Port Blocking:	no	Link-vlan:	1
Port:	2		
Current State:	up	Neighbor MAC Addr:	000102-030405
Udld Packets Sent:	500	Neighbor Port:	6
Udld Packets Received:	450	State Transitions:	3
Port Blocking:	no	Link-vlan:	200
Port:	3		
Current State:	off line	Neighbor MAC Addr:	n/a
Udld Packets Sent:	0	Neighbor Port:	n/a
Udld Packets Received:	0	State Transitions:	0
Port Blocking:	no	Link-vlan:	1
Port:	4		
Current State:	failure	Neighbor MAC Addr:	n/a
Udld Packets Sent:	128	Neighbor Port:	n/a
Udld 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 25 \(page 70\)](#) for an example.)

Configuring UFD

Syntax

```
[no] uplink-failure-detection
```

Globally enables UFD.

The no form of the command globally disables UFD.

Syntax

```
[no] uplink-failure-detection track track-id links-to-monitor [[lacp-key]  
| [port-list]]  
links-to-disable [[lacp-key] | [port-list]]
```

Configures ports as LtM ports and LtD ports for the specified track. Trunk interfaces are also configurable. The no form of the command removes any track data associated with the specified track.

<code>track <i>track-id</i></code>	Range is 0–63.
<code>links-to-monitor [<i>lacp-key</i> <i>port-list</i>]</code>	Specifies the lacp key or port list to monitor.
<code>links-to-disable [<i>lacp-key</i> <i>port-list</i>]</code>	Specifies the lacp key or port list to disable.

About 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 \(CLI\)” \(page 48\)](#).)
- To display information about the transceivers installed on a switch, enter the `show tech receivers` command in the CLI ([Figure 17 \(page 56\)](#).)

Table 4 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.

Table 4 Status and parameters for each port type *(continued)*

Status or parameter	Description
	<p>10/100/1000Base-T Ports:</p> <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. • 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.

Table 4 Status and parameters for each port type *(continued)*

Status or parameter	Description
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 <code>Auto</code> (the default) and flow control enabled, the switch negotiates flow control on the indicated port. If the port mode is not set to <code>Auto</code>, 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>

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: <i>custom</i>

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`.

About 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.

HP Switch 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*.

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 74) 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.

About 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 64).)
- 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.)

About configuring transceivers and modules that have not been inserted

Transceivers

Previously, a port had to be valid and verified for the switch to allow it to be configured. Transceivers are removable ports and considered invalid when not present in the switch, so they cannot be configured unless they are already in the switch. For HP switches, the verification for allowable port configurations performed by the CLI is removed and configuration of transceivers is allowed even if they are not yet inserted in the switch.

Modules

You can create or edit configuration files (as text files) that can be uploaded to the switch without the modules having been installed yet. Additionally, you can pre-configure the modules with the CLI `module` command.

The same `module` command used in an uploaded configuration file is used to define a module that is being pre-configured. The validation performed when issued through the CLI is still performed just as if the command was executed on the switch, in other words, as if the module were actually present in the switch.

NOTE: You cannot use this method to change the configuration of a module that has already been configured. The slot must be empty and the configuration file must not have a configuration associated with it.

About clearing the module configuration

Because of the hot-swap capabilities of the modules, when a module is removed from the chassis, the module configuration remains in the configuration file. `[no] module slot` allows you to remove the module configuration information from the configuration file.

NOTE: This does not change how hot-swap works.

Restrictions

The following restrictions apply:

- The slot being cleared must be empty
- There was no module present in the slot since the last boot
- If there was a module present after the switch was booted, the switch will have to be rebooted before any module (new or same) can be used in the slot.
- This does not clear the configuration of a module still in use by the switch.

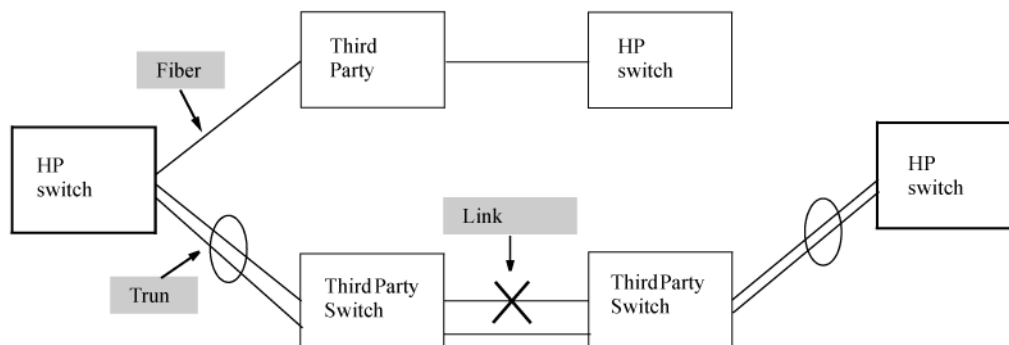
Uni-directional link detection (UDLD)

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 26 \(page 76\)](#) shows an example.

Figure 26 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.

About 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.

NOTE: UDLD interoperates with the following HP switches: 2600, 2800, 3400, 3500, 4200, 5300, 5400, 6200, 6400, 6600, 8212, and 9300. Consult the release notes and current manuals for required software versions.

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.

About Uplink failure detection

Uplink Failure Detection (UFD) is a network path redundancy feature that works in conjunction with NIC teaming functionality. UFD continuously monitors the link state of the ports configured as links-to-monitor (LtM), and when these ports lose link with their partners, UFD will disable the set of ports configured as links-to-disable (LtD.) When an uplink port goes down, UFD enables the switch to auto-disable the specific downlinks connected to the NICs. This allows the NIC teaming software to detect link failure on the primary NIC port and fail over to the secondary NIC in the team.

NIC teams must be configured for switch redundancy when used with UFD, that is, the team spans ports on both Switch A and Switch B. The switch automatically enables the downlink ports when the uplink returns to service. For an example of teamed NICs in conjunction with UFD, see [Figure 27 \(page 78\)](#).) For an example of teamed NICs with a failed uplink, see [Figure 28 \(page 78\)](#).

NOTE: For UFD functionality to work as expected, the NIC teaming must be in Network Fault Tolerance (NFT) mode.

Figure 27 Example of teamed NICs in conjunction with UFD

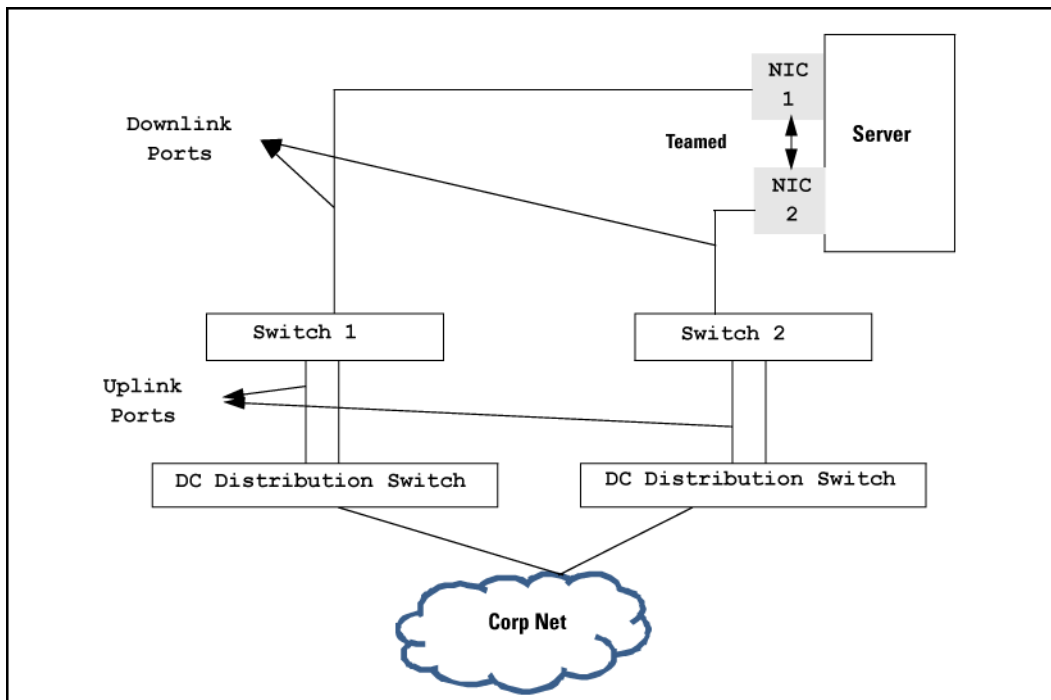
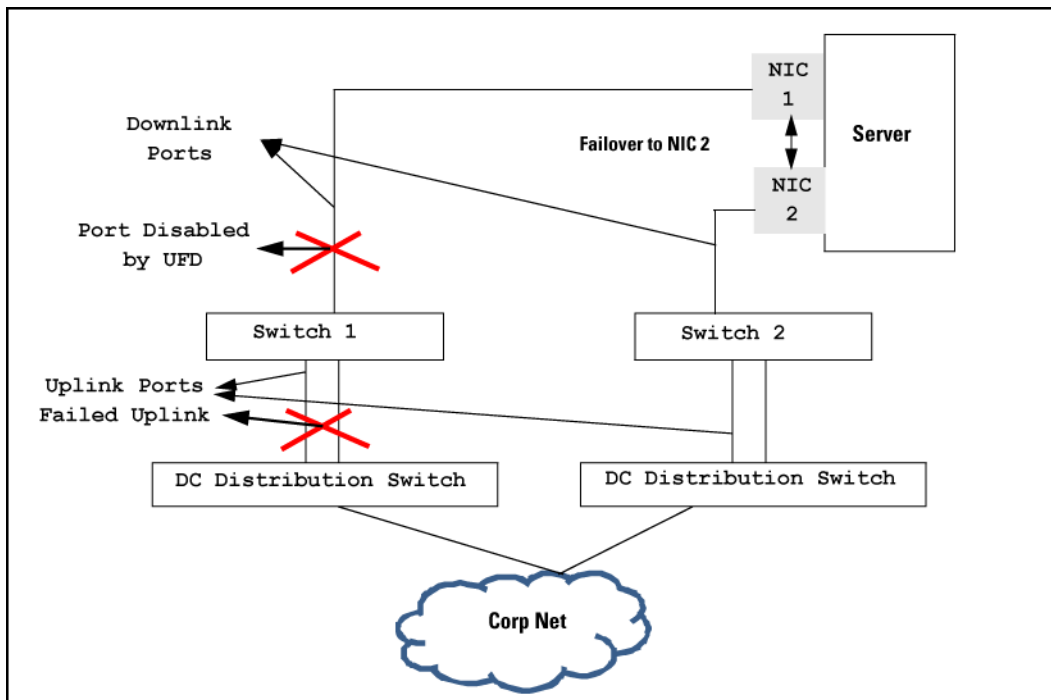


Figure 28 Example of teamed NICs with a failed uplink



UFD operating notes

- A port cannot be added to a trunk group if it already belongs to an LtM or LtD.
- Ports that are already members of a trunk group cannot be assigned to an LtM or LtD.
- Trunks that are configured as LtM or LtD cannot be deleted.

Example 30 Configuring ports as LtM and LtD for track 3

```
HP Switch(config)# uplink-failure-detection track 3 links-to-monitor 5,6,7
links-to-disable 8,9,10
```

Example 31 Removing a LtM port and an LtD port for track 3

```
HP Switch(config)# no uplink-failure-detection track 3 links-to-monitor 5
links-to-disable 8
```

Viewing UFD configuration

Enter the `show uplink-failure-detection` command to display information about the UFD configuration.

Example 32 Displaying informationn for UFD configuration

```
HP Switch(config)# show uplink-failure-detection
UFD State: Enabled
```

Failure Detection Pairs:

TRACK ID	Monitored Links	Links to Disable	LtM State	LtD State
3	5,6	D3	Up	Up
4	13	D4,D5	Down	Auto-Disabled

3 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 Planning and Implementation Guide*, which is available on the HP Networking website at

www.hp.com/networking/support.

PoE

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

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

Syntax

The HP switches covered in this guide also support some pre-802.3af devices.

```
[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. In earlier software, the default setting is "on." In K.15.02 and later software, the default setting is "off."

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

Critical	Specifies the highest-priority PoE support for <i>port-list</i> . 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 <i>port-list</i> . 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 <i>port-list</i> . 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.

For more information configuring PoE port priority and on assigning PoE priority on more than one module, see [“About assigning PoE priority with two or more modules”](#) (page 94).

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.

usage (Default) The automatic allocation by a PD.

class 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 6](#) (page 81).

value 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 6 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.

1. To configure a port by value, first 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. Then select a value:

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

or in interface context:

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

3. To view the settings, enter the `show power-over-ethernet` command, shown in Figure Figure 29 (page 82).

Figure 29 Example displaying 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 A6

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

Over Current Cnt : 0
Power Denied Cnt : 0
Voltage          : 55.1 V
Power            : 8.4 W
MPS Absent Cnt  : 0
Short Cnt        : 0
Current          : 154 mA
```

If you set the PoE maximum value to less than what the PD requires, a **fault** occurs, as shown in Figure Figure 30 (page 82).

Figure 30 Example showing 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
LLDP Detect      : enabled
Configured Type  : 4 W
Value            : 4 W
Power Class      : 2
Maximum power delivered

Over Current Cnt : 1
Power Denied Cnt : 2
Voltage          : 55.1 V
Power            : 8.4 W
MPS Absent Cnt  : 0
Short Cnt        : 0
Current          : 154 mA
```

Configuring PoE redundancy (chassis switches only)

PoE redundancy occurs automatically when enabled. The switch keeps track of power use and does not supply PoE power to additional PoE devices trying to connect if that results in the switch not having enough power in reserve for redundancy.

Syntax

```
[no] power-over-ethernet redundancy [ n+1 | full ]
```

Allows you to set the amount of power held in reserve for redundancy.

no	Means that all available power can be allocated to PDs. Default: No PoE redundancy enforced.
n+1	One of the power supplies is held in reserve for redundancy. If a single power supply fails, no powered devices are shut down. If power supplies with different ratings are used, the highest-rated power supply is held in reserve to ensure full redundancy.
full	Half of the available power supply is held in reserve for redundancy. If power supplies with different ratings are used, the highest-rated power supply is held in reserve to ensure full redundancy.

For more information about PoE redundancy and power supplies, see the *PoE Planning and Implementation Guide*, available on the HP website at

www.hp.com/networking/support.

Changing the threshold for generating a power notice

Syntax

```
power-over-ethernet [ slot slot-id-range ] threshold 1 - 99
```

Specifies the PoE usage level (as a percentage of the PoE power available on a module) 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 a PoE module's power consumption crosses the configured threshold value. That is, the switch generates a notice whenever the power consumption on a module either exceeds or drops below the specified percentage of the total PoE power available on the module.

This command configures the notification threshold for PoE power usage on either a global or per-module (slot) basis.

Without the [slot PoE slot-id-range] option, the switch applies one power threshold setting on all PoE modules installed in the switch.

For more information on configuring thresholds, see [“About configuring thresholds for generating a power notice” \(page 96\)](#).

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 [“About PoE/PoE+ allocation using LLDP information” \(page 97\)](#).

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 [“About PoE/PoE+ allocation using LLDP information” \(page 97\)](#).

Negotiating power using the DLL

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 increases. PoE and LLDP interact to meet the current power demands.

Syntax

```
int port-list PoE-lldp-detect [ enabled | disabled ]
```

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

Default: Disabled

Example

You can enter this command to enable LLDP detection:

```
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 [Figure 31 \(page 85\)](#).

Figure 31 Example of port with LLDP configuration information obtained from the device

```
HPswitch(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   | Power  Power   Alloc Alloc Actual Configured  Detection  Power
Port  | Enable Priority By   Power Power  Type      Status     Class
-----+-----
A1    | Yes    low    usage 17 W  5.0 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  0.0 W                    Searching  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 08/04/10 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 PoE when using LLDP information

Syntax

```
show lldp config port-list
```

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

Examples

Figure 32 LLDP port configuration information with PoE

```
HPSwitch(config)# show lldp config 4

LLDP Port Configuration Detail

Port : 4
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
* poeplus_config

IpAddress Advertised:
```

Figure [Figure 33 \(page 86\)](#) shows an example of the local device power information using the `show lldp info local-device port-list` command.

Figure 33 Local power information

```
HPSwitch(config) show lldp info local-device A1

LLDP 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
```

Figure [Figure 34 \(page 87\)](#) shows an example of the remote device power information using the `show lldp info remote-device port-list` command.

Figure 34 Remote power information

```
HPswitch(config) show lldp info remote-device A3

LLDP 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 3500-24, revision K.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
```

Viewing the global PoE power status of the switch

Syntax

```
show power-over-ethernet [ brief | [ethernet]port-list | [ slot slot-id-range | 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 88).
<code>port-list</code>	Displays PoE information for the ports in port-list. See “Viewing the PoE status on specific ports” (page 90).
<code>slot-id-range</code>	Displays PoE information for the selected slots. (See figure Figure 37 (page 89).) Enter the <code>all</code> option to display the PoE information for all slots.

Example

show power-over-ethernet displays data similar to that shown in Figure [Figure 35 \(page 88\)](#).

Figure 35 show power-over-ethernet command output

```
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, including:

- **PoE Port**
Lists all PoE-capable ports on the switch.
- **Power Enable**
Shows **Yes** for ports enabled to support PoE (the default) and **No** for ports on which PoE is disabled.
- **Power Priority**
Lists the power priority (**Low**, **High**, and **Critical**) configured on ports enabled for PoE. (For more information on this topic, see "Configuring PoE Operation" on page K-6.)
- **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:**
 - 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** Shows the 802.3af power class of the PD detected on the indicated port. Classes include:

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

Examples

show power-over-ethernet brief displays this output:

Figure 36 show power-over-ethernet brief command output

```

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: 600 W  Used: 9 W  Remaining: 591 W

Module A Power
Available: 408 W  Used: 9 W  Remaining: 399 W

PoE   | Power  Power  Alloc Alloc  Actual Configured  Detection  Power
Port  | Enable Priority By   Power Power  Type          Status     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:

Figure 37 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 <i>port-list</i></code> .										
Priority	Lists the power priority (Low , High , and Critical) configured on ports enabled for PoE. (For more on this topic, see "Configuring PoE Operation" on page K-6.)										
Allocate by	How PoE is allocated (usage , class , value .)										
Detection Status	<table><tr><td>Searching</td><td>The port is available to support a PD.</td></tr><tr><td>Delivering</td><td>The port is delivering power to a PD.</td></tr><tr><td>Disabled</td><td>PoE power is enabled on the port but the PoE module does not have enough power available to supply the port's power needs.</td></tr><tr><td>Fault</td><td>The switch detects a problem with the connected PD.</td></tr><tr><td>Other Fault</td><td>The switch has detected an internal fault that prevents it from supplying power on that port.</td></tr></table>	Searching	The port is available to support a PD.	Delivering	The port is delivering power to a PD.	Disabled	PoE power is enabled on the port 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.
Searching	The port is available to support a PD.										
Delivering	The port is delivering power to a PD.										
Disabled	PoE power is enabled on the port 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.

Example

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:

Figure 38 Example of show power-over-ethernet port-list output

```

HP Switch(config)# show power-over-ethernet 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 Planning and Implementation Guide* which is available on the HP Networking web site at www.hp.com/networking/support.

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.

By installing a second power supply in the 5406zl/8206zl or a third power supply in a 5412zl/8212zl chassis, depending on how many PoE ports are being supplied with power, the switch can have redundant power if one power supply fails. A Power Supply Shelf (external power supply) can also be connected to the 5400zl/8200zl switches to provide extra or redundant PoE power.

For example, if the 5406zl has two 24-port PoE modules (J8702A) installed, and all ports are using 15.4 watts, then the total wattage used is 739.2 watts (48 x 15.4.) To supply the necessary PoE wattage a J8713A power supply is installed in one of the power supply slots.

To gain redundant power, a second J8713A must be installed in the second power supply slot. If the first power supply fails, then the second power supply can supply all necessary power.

See the *HP PoE Planning and Implementation Guide* for detailed information about the PoE/PoE+ power requirements.

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*.

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 7 \(page 92\)](#).

Table 7 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, see the *Advanced Traffic Management Guide*.

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. A slot in a zl chassis can provide a maximum of 370 watts of PoE/PoE+ power to a module.

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” \(page 94\)](#).)

For PoE+, there must be 33 watts available for the module to begin supplying power to a port with a PD connected. A slot in a zl chassis can provide a maximum of 370 watts of PoE/PoE+ power to a module.

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.

About assigning PoE priority with two or more modules

Ports across two or more modules can be assigned a class priority of low (the default), high, or critical. For example, A5, B7, and C10 could all be assigned a priority class of **Critical**. When power is allocated to the ports on a priority basis, the **Critical** priority power requests are allocated to module A first, then Module B, then C, and so on. Next, the **High** priority power requests are allocated, starting with module A, then B, then C, and the remaining modules in order. Any remaining power is allocated in the same manner for the **Low** priority ports, beginning with module A though the remaining modules. If there is not enough PoE power for all the PDs connected to PoE modules in the switch, power is allocated according to priority class across modules.

Example

All ports on module C are prioritized as **Critical**.

```
HP Switch(config)# interface c1-c24 power-over-ethernet  
critical
```

All ports on module A are prioritized as **Low**.

```
HP Switch(config)# interface a1-a24 power-over-ethernet
low
```

There are 48 PDs attached to all ports of modules A and C (24 ports each module); however, there is enough PoE power for only 32 ports (8.5 watts × 32 ports=273 watts.) The result is that all the **Critical** priority ports on module C receive power, but only 8 ports on module A receive power.

On module A, the port A1 has the highest priority of the ports in that module if all ports are in the same priority class, which is the case for this example. Since a minimum 17 + 5 watts of power is allocated per PoE module for PoE, port A1 will always receive PoE power. If another port on module A had a higher priority class than port A1, that port would be allocated the power before port A1.

For PoE+ modules there must be a minimum of 33 + 5 watts of power allocated per PoE+ module.

About configuring PoE operation

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 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.

For a given level, ports are prioritized by port number in ascending order. For example, if ports A1 to A24 have a priority level of critical, port A1 has priority over ports A2 to A24.

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.)

In chassis switches, you can use one command to set the same priority level on PoE ports in multiple modules. For example, to configure the priority to **High** for ports c5 to c10, C23 to C24, D1 to D10, and D12, you could use this command:

```
HP Switch(config)# interface c5-c10,c23-c24,
d1-d10,d12 power-over-ethernet high
```

Example

Suppose that you configure the PoE priority for a module in slot C as shown in [Table 8 \(page 95\)](#).

Table 8 PoE priority operation on a PoE module

Port	Priority setting	Configuration command ¹ and resulting operation with PDs connected to ports C3 through C24
C3 - C17	Critical	<p>In this example, the following CLI command sets ports C3 to C17 to Critical:</p> <pre>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</p>

Table 8 PoE priority operation on a PoE module *(continued)*

Port	Priority setting	Configuration command ¹ and resulting operation with PDs connected to ports C3 through C24
		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.
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> 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.
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> 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.
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> There is no priority setting for the ports in this example.

¹ For a listing of PoE configuration commands with descriptions, see "Configuring PoE Operation" on page K-6.

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

About configuring thresholds for generating a power notice

You can configure one of the following thresholds:

A global power threshold that applies to all modules on the switch.

This setting acts as a trigger for sending a notice when the PoE power consumption on any PoE module 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%.

A per-slot power threshold that applies to an individual PoE module installed in the designated slot.

This setting acts as a trigger for sending a notice when the module in the specified slot exceeds or goes below a specific level of PoE power consumption.

Example

Suppose slots A, B, and C each have a PoE module installed. In this case, executing the following command sets the global notification threshold to 70% of available PoE power.


```
HP Switch(config)# power-over-ethernet threshold
70
```

With this setting, if module B is allocated 100 watts of PoE power and is using 68 watts, and then another PD is connected to the module in slot B that uses 8 watts, the 70% threshold of 70 watts is exceeded. The switch sends an SNMP trap and generates this Event Log message:

```
Slot B POE usage has exceeded threshold of 70%.
```

If the switch is configured for debug logging, it also sends the Event Log message to the configured debug destinations.

On any PoE module, if an increasing PoE power load (1) exceeds the configured power threshold—which triggers the log message and SNMP trap—and then (2) later decreases and drops below the threshold again, the switch generates another SNMP trap, plus a message to the Event Log and any configured Debug destinations.

Example

To continue the preceding example, if the PoE power usage on the PoE module in slot B drops below 70%, another SNMP trap is generated and you will see this message in the Event Log:

```
Slot B POE usage is below threshold of 70%.
```

By using the `[slot slot-id-range]` option, you can specify different notification thresholds for different PoE modules installed in the switch. For example, you could set the power threshold for a PoE module in slot "A" to 75% and the threshold for the module in slot "B" to 68% by executing the following two commands:

```
HP Switch(config)# power-over-ethernet slot a
threshold 75
HP Switch(config)# power-over-ethernet slot b
threshold 68
```

The last `threshold` command affecting a given slot supersedes the previous `threshold` command affecting the same slot. Thus, executing the following two commands in the order shown sets the threshold for the PoE module in slot "D" to 75%, but leaves the thresholds for any PoE modules in the other slots at 90%:

```
HP Switch(config)# power-over-ethernet
threshold 90
HP Switch(config)# power-over-ethernet slot d
threshold 75
```

(If you reverse the order of the above two commands, all PoE modules in the switch will have a threshold of 90%.)

Without the `[slot slot-id-range]` option, the switch applies one power threshold setting on all PoE modules installed in the switch.

About 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.

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” \(page 85\)](#). 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.

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)
```

4 Port Trunking

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.

- 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.

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 [Figure 39 \(page 99\)](#) and [Example 33 \(page 100\)](#) 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:

Figure 39 Example 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
```

Port 5 appears with an example of a name that you can optionally assign using the Friendly Port Names feature. (Refer to "Using Friendly (Optional) Port Names".)

Port 6 does not appear in this listing because it is not assigned to a static 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 33 \(page 100\)](#), the command does not include a port list, so the switch lists all ports having static trunk membership.

Example 33 Example of 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 12 \(page 121\)](#).)

Example 34 Example of 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 12 \(page 121\)](#).

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".)

Syntax

```
trunk port-list trk1 ... trk144 trunk | lacp
```

Configures the specified static trunk type.

Example

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

```
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 dynamic LACP trunk groups

An individual trunk can have up to eight links, with additional standby links if you are using LACP. You can configure trunk group types as follows:

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.

Setting the LACP key

During dynamic link aggregation using LACP, ports with the same key are aggregated as a single trunk.

Syntax

```
[no]lacp [[active] | [passive] | [key 0-65535]]
```

Example 35 Enabling LACP and configuring an LACP key

```
HP Switch(config)# int A2-A3 lacp active
HP Switch(config)# int A2-A3 lacp key 500
```

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

LACP							
Port	LACP Enabled	Trunk Group	Port Status	Partner	LACP Status	Admin Key	Oper Key
A2	Active	A2	Down	No	Success	500	500
A3	Active	A3	Down	No	Success	500	500

Example 36 Interface configured with a different LACP key

```
HP Switch(config)# int A5 lacp active
HP Switch(config)# int A5 lacp key 250
```

```
HP Switch> show lacp
```

LACP							
Port	LACP Enabled	Trunk Group	Port Status	Partner	LACP Status	Admin Key	Oper Key
A1	Active	Dyn1	Up	Yes	Success	100	100
A2	Active	Dyn1	Up	Yes	Success	100	100
A3	Active	Dyn1	Up	Yes	Success	100	100
A4	Active	Dyn1	Up	Yes	Success	100	100
A5	Active	A5	Up	No	Success	250	250

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.

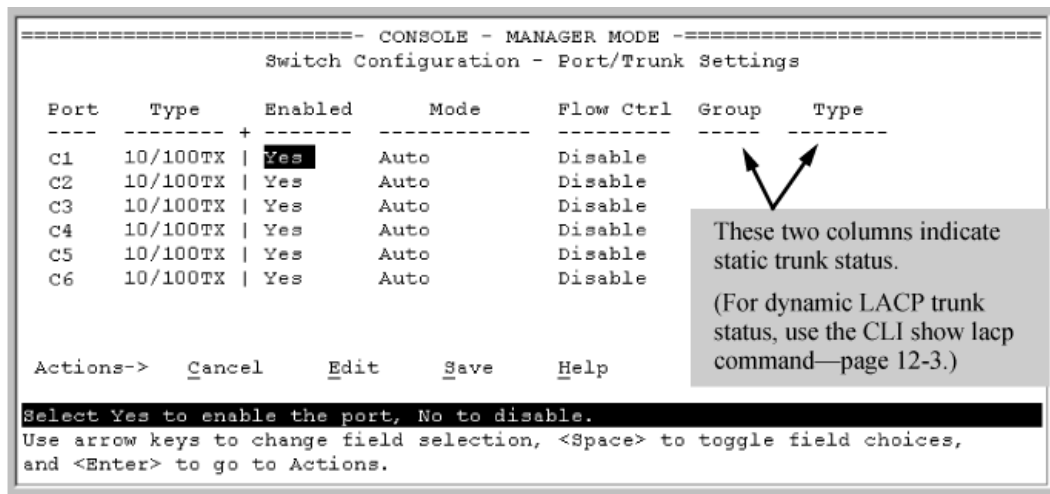
Viewing and configuring a static trunk group (Menu)

- ❗ **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 Edit) and then use the arrow keys to access the port trunk parameters.

Figure 40 Example of the menu screen for configuring a port trunk group



4. In the Group column, move the cursor to the port you want to configure.
5. Use the Space bar to choose a trunk group assignment (Trk1, Trk2, and so on) for the selected port.
 - 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 *Advanced Traffic Management Guide*.

(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 41 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->  Cancell  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" on page C-28) to verify that the trunked ports are operating properly.

Enabling L4-based trunk load balancing (CLI)

For detailed information about trunk load balancing, see ["Trunk load balancing using Layer 4 ports"](#) (page 126).

Enter the following command with the L4-based option to enable load balancing on Layer 4 information when it is present.

Syntax

```
trunk-load-balance [ L3-based | L4-based ]
```

When the L4-based option is configured, enables load balancing based on Layer 4 information if it is present. If it is not present, Layer 3 information is used if present; if Layer 3 information is not present, Layer 2 information is used. The configuration is executed in global configuration context and applies to the entire switch.

L3-based	Load balance 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.

Default: L3-based load balancing

Examples

Figure 42 Enabling L4-based trunk load balancing

```
HPswitch(config)# trunk-load-balance L4-based
```

Figure 43 Output when L4-based trunk load balancing is enabled

```
HPswitch(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
```

Figure 44 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
```

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.

Displaying information about trunk load balancing

The show trunks load-balance interface command displays the port on which the information will be forwarded out for the specified traffic flow with the specified source and destination address.

Syntax

```
show trunks load-balance interface trunk-id mac src-addr dest-addr [ip
src-addr dest-addr [[src-tcp-port | src-udp-port] [[dest-tcp-port |
dest-udp-port]]]
inbound-port port-num
```

Displays the port on which the information will be forwarded out for the specified traffic flow with the specified source and destination address.

<i>trunk-id</i>	The trunk id (trk1, trk2, etc.)
mac <i>src-addr dest-addr</i>	The source MAC address and the destination MAC address.
ip <i>src-addr dest-addr</i>	The source IPv4 /IPv6 address and the destination IPv4/IPv6 address.
[<i>src-tcp-port</i> <i>src-udp-port</i>]	The source TCP port or the source UDP port.

[*dest-tcp-port* | *dest-udp-port*] The destination TCP port or the destination UDP port.

inbound-port port-num the port number of which the traffic is received.

Example 37 Example showing information about the forwarding port

```
HP Switch# show trunks load-balance interface trk1 mac 424521-498421 534516-795463 inbound-port a5
Traffic in this flow will be forwarded out port 23 based on the configured L2 load balancing.
```

Operating notes

The port cannot be determined if:

- All the ports in the trunk are down.
- The MAC address is all zeros.
- The source MAC address is broadcast or multicast.

Distributed trunking

For detailed information on distributed trunking, see “[Distributed trunking overview](#)” (page 127).

Configuring ISC ports

You must configure the ISC ports before you can configure the trunks for distributed trunking.

Syntax

```
switch-interconnect [ port-num | trk1...trkN ]
no switch-interconnect
```

Configures an InterSwitch-Connection (ISC) port. The

```
[ port-num | trk1...trkN ]
```

variable is the interconnect interface that connects two distributed trunking switches. It can be a physical port, manual LACP trunk, or manual non-protocol trunk. You can override an ISC configuration by configuring the command with a different value.

The `no` form of the command removes the ISC interface configuration.

NOTE: A port that is already part of a trunk cannot be configured as an ISC interface.

Configuring distributed trunking ports

Distributed trunking ports must be configured manually.

Syntax

```
trunk port-list trk1 | trk2 ... trkN [ trunk | lACP | dt-lACP | dt-trunk ]
no trunk port-list
```

Configures distributed trunking on a switch. Use either the `dt-lACP` or `dt-trunk` option.

The trunk groups and trunk types must be identical in both switches. For example, if Switch Local is configured with `trk1` and uses the `dt-lACP` option, Switch Remote also must be configured with `trk1` and use the `dt-lACP` option to form a distributed trunk. Similarly, if Switch Local is configured with `trk2` and uses the `dt-trunk` option, Switch Remote must be configured with `trk2` and use the `dt-trunk` option to form the distributed trunk.

The `no` form of the command removes the distributed trunking configuration on the switch.

Example

Figure 45 (page 107) shows an ISC port being configured for the local switch and the remote switch.

Figure 45 Configuring distributed trunking

```
HP Switch Local(config)# switch-interconnect a7
HP Switch Remote(config)# switch-interconnect a8

HP Switch Local(config)# trunk a9-a10 trk10 dt-lacp
HP Switch Remote(config)# trunk a5-a6 trk10 dt-lacp

HP Switch Local(config)# trunk a1-a2 trk20 dt-trunk
HP Switch Remote(config)# trunk a3-a4 trk20 dt-trunk
```

Configuring peer-keepalive links

For detailed information about configuring peer-keepalive links, see [“About configuring peer-keepalive links”](#) (page 129).

Syntax

```
[no] distributed-trunking [hold-timer 3-10] [ peer-keepalive destination
ip-address | vlan vid [interval 400-10000] [ timeout 3-20] [udp-port
1024-49151] ]
```

Distributed trunking uses a VLAN interface between DT peers to transmit periodic peer-keepalive messages. This command configures the peer-keepalive parameters for distributed trunking.

hold-timer 3-10	Configures the hold time in seconds. Default is 3 seconds.
peer-keepalive:	
	destination: The destination IPv4 address to be used by DT switches to send peer-keepalive messages to the peer DT switch when the ISC is down.
	vlan vid : The VLAN used exclusively for sending and receiving peer-keepalive messages.
	interval 400-10000: The interval between peer-keepalive messages (in milliseconds.) Default is 1000 milliseconds.
	timeout 3-20: The peer-keepalive timeout in seconds. Default is 5 seconds.
	udp-port 1024-49151: The source UDP port to be used for transmitting peer-keepalive HELLO messages.

Viewing distributed trunking information

Syntax

```
show lacp [distributed]
```

Displays information about distributed trunks and LACP status.

Example

```
HP Switch Local (config)#: show lacp distributed
                          Distributed LACP
```

```
Local Port Status:
```

Port	LACP Enabled	Trunk Group	Port Status	Partner	LACP Status	Admin Key	Oper Key	Active	Trk10
Up 350	Yes 350	Sucess	350	350A10	Active	Trk10	Up	Yes	Sucess

```
Remote Port Status
```

Port	LACP Enabled	Trunk Group	Port Status	Partner	LACP Status	Oper Key
A5	Active	Trk10	Up	Yes	Sucess	200
A6	Active	Trk10	Up	Yes	Sucess	200

Syntax

```
show distributed trunk consistency parameters global
```

This command displays configured features on VLANs that have dt-lacp or dt-trunk ports as member port. This command also displays VLAN memberships and loop-protect status of a given DT trunk. You can use this command to determine if there is any mismatch in the configuration parameters on VLANs configured for DT ports or on DT interfaces.

Example

```
show distributed trunk consistency parameters global
```

	Local	Peer
Image Version	K.15.XX	K.15.XX
IP Routing	Enabled	Enabled
Peer keepalive interval (ms)	1000	1000
IGMP enabled VLANs on Local	: 1 10, 100 110, 501 ,600 610 ,800	
IGMP enabled VLANs on Peer	: 1 10, 100 110, 501 ,600	
DHCP snooping enabled VLANs on Local	: 1,2	
DHCP snooping enabled VLANs on Peer	: 1	
Loop protect enabled VLANs on Local	: 1,4	
Loop protect enabled VLANs on Peer	: 1,5	
MLD enabled VLANs on Local	: 1 10	
MLD enabled VLANs on Peer	: 1 10	

Example

```
Show distributed trunk
consistency parameters trunk <trk1...trkN>
Allowed VLANs on Local : 1 10, 100 110, 501 ,600
610 ,800
Allowed VLANs on Peer : 1 10, 100 110, 501 ,600
610 ,800
```

Name	Local Value	Peer Value
Loop protect	Enabled	Enabled

Viewing information about the peer-keepalive configuration

Syntax

```
show distributed-trunking peer-keepalive
```

Displays information about peer-keepalive parameters.

Example

Figure 46 Output displaying peer-keepalive settings

```
HPswitch(config)# show distributed-trunking peer-keepalive
Distributed Trunking peer-keepalive parameters
Destination      :10.10.10.2
VLAN             : 2
UDP Port         : 1024
Interval (ms)    : 1000
Timeout (sec)    : 5
```

Viewing the switch interconnect information

Syntax

```
show switch-interconnect
```

Displays information about switch interconnect settings.

Example


Figure 47 Switch-interconnect settings

```
HPswitch(config)# show switch-interconnect
Port           :Trk2
Status         :Up
Active VLANs   :2,3,4,30
```

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 48 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 onemedia type in a port trunk group. Similarly, for proper trunk operation, all links in the same trunk group must have the samespeed, duplex, and flow control.

Port security restriction

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

Port trunk features and operation

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

- LACP: IEEE 802.3ad—[Enabling L4-based trunk load balancing \(CLI\)](#)
- Trunk: Non-Protocol—[Trunk group operation using the "trunk" option](#)

Up to 144 trunk groups are supported on the switches. The actual maximum depends on the number of ports available on the switch and the number of links in each trunk. (Using the link aggregation control protocol—LACP—option, you can include standby trunked ports in addition to the maximum

of eight actively trunking ports.) The trunks do not have to be the same size; for example, 100 two-port trunks and 11 eight-port trunks are supported.

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” \(page 119\)](#).)

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.

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” on page 4-22.) 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”.

```
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

Displaying LACP Local Information

```
HP Switch# show lacp local
LACP Local Information.
System ID: 001871-b98500
```

Port	Trunk	LACP Mode	Aggregated	Tx Timer	Rx Timer Expired
A2	A2	Active	Yes	Fast	No
A3	A3	Active	Yes	Fast	No

Displaying LACP Peer Information

Use the `show lacp peer` command to display information about LACP peers. The System ID represents the MAC address of a partner switch. It will be zero if a partner is not found.

```
HP Switch(config)# show lacp peer
LACP Peer Information.
System ID: 001871-b98500
```

Local Port	Local Trunk	System ID	Port	Port Priority	Oper Key	LACP Mode	Tx Timer
A2	A2	123456-654321	2	0	100	Passive	Fast
A3	A3	234567-456789	3	0	100	Passive	Fast

Displaying LACP Counters

Use the `show lacp counters` command to display statistical information about LACP ports.

Note on the Marker Protocol. Data traffic can be dynamically redistributed in port channels. This may occur when a link is added or removed, or there is a change in load-balancing. Traffic that is redistributed in the middle of a traffic flow could potentially cause mis-ordered data packets.

LACP uses the marker protocol to prevent data packets from being duplicated or reordered due to redistribution. Marker PDUs are sent on each port-channel link. The remote system responds to the marker PDU by sending a marker responder when it has received all the frames received on this link prior to the marker PDU. When the marker responders are received by the local system on all member links of the port channel, the local system can redistribute the packets in the traffic flow correctly.

For the switches covered in this guide, the marker BPDUs are not initiated, only forwarded when received, resulting in the Marker fields in the output usually displaying zeros.

```
HP Switch(config)# show lacp counters
LACP Port Counters.
```

Port	Trunk	LACP PDUs Tx	LACP PDUs Rx	Marker Req. Tx	Marker Req. Rx	Marker Resp. Tx	Marker Resp. Rx	Error
A2	A2	1234	1234	0	0	0	0	0
A3	A3	1234	1234	0	0	0	0	0

Using keys to control dynamic LACP trunk configuration

The `lacp key` option provides the ability to control dynamic trunk configuration. Ports with the same key will be aggregated as a single trunk.

There are two types of keys associated with each port, the Admin key and the Operational key. The Operational key is the key currently in use. The Admin key is used internally to modify the value of the Operational key. The Admin and Operational key are usually the same, but using static LACP can alter the Operational key during runtime, in which case the keys would differ.

The `lacp key` command configures both the Admin and Operational keys when using dynamic LACP trunks. It only configures the Admin key if the trunk is a static LACP trunk. It is executed in the interface context.

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 123).) • You want to use a monitor port on the switch to monitor an LACP trunk. <p>For more information, see “Enabling L4-based trunk load balancing (CLI)” (page 104).</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

Table 10 Trunk configuration protocols (continued)

Protocol	Trunking Options
	<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 125).</p>

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 recommends leaving the port Mode setting at <code>Auto</code> or, in networks using Cat 3 cabling, <code>Auto-10</code> .)
Port Configuration	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.

Figure 49 Recommended port mode setting for LACP

HP Switch(config)# show interfaces config

Port Settings

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

All of the following operate on a per-port basis, regardless of trunk membership:

- Enable/Disable
- Flow control (Flow Ctrl)

LACP is a full-duplex protocol. See [“Enabling L4-based trunk load balancing \(CLI\)”](#) (page 104).

Trunk configuration: All ports in the same trunk group must be the same trunk type (LACP or trunk.) All LACP ports in the same trunk group must be either all static LACP or all dynamic LACP.

A trunk appears as a single port labeled `Dyn1` (for an LACP dynamic trunk) or `Trk1` (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 125)

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.)

Traffic distribution:

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 125).

Spanning Tree:

802.1D (STP) and 802.1w (RSTP) Spanning Tree operate as a global setting on the switch (with one instance of Spanning Tree per switch.) 802.1s (MSTP) Spanning Tree operates on a per-instance basis (with multiple instances allowed per switch.) For each Spanning Tree instance, you can adjust Spanning Tree parameters on a per-port basis.

A static trunk of any type appears in the Spanning Tree configuration display, and you can configure Spanning Tree parameters for a static trunk in the same way that you would configure Spanning Tree parameters on a non-trunked port. (Note that the switch lists the trunk by name—such as Trk1—and does not list the individual ports in the trunk.) For example, if ports C1 and C2 are configured as a static trunk named Trk1, they are listed in the Spanning Tree display as Trk1 and do not appear as individual ports in the Spanning Tree displays. See [Figure 50](#) (page 115).

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.

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.

If you remove a port from a static trunk, the port retains the same Spanning Tree settings that were configured for the trunk.

Figure 50 Example of a port trunk in a Spanning Tree listing

Port	Type	Cost	Priority	State	Designated Bridge
C3	100/1000T	5	128	Forwarding	0020c1-b27ac0
C4	100/1000T	5	128	Forwarding	0060b0-889e00
C5	100/1000T	5	128	Disabled	
C6	100/1000T	5	128	Disabled	
Trk1		1	64	Forwarding	0001e7-a0ec00

IP multicast protocol (IGMP):

A static trunk of any type appears in the IGMP configuration display, and you can configure IGMP for a static trunk in the same way that you would configure IGMP on a non-trunked port. (Note that the switch lists the trunk by name—such as Trk1—and does not list the individual ports in the trunk.) Also, creating a new trunk automatically places the trunk in IGMP Auto status if IGMP is enabled for the default VLAN.

A dynamic LACP trunk operates only with the default IGMP settings and does not appear in the IGMP configuration display or `show ip igmp` listing.

VLANs:

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.

NOTE: For a dynamic LACP trunk to operate in a VLAN other than the default VLAN (DEFAULT_VLAN), GVRP must be enabled. See [“Enabling L4-based trunk load balancing \(CLI\)”](#) (page 104).

Port security:

Trunk groups (and their individual ports) cannot be configured for port security, and the switch excludes trunked ports from the `show port-security` 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:

```
port-list Command cannot operate over a logical port.
```

Monitor port:

NOTE: A trunk cannot be a monitor port. A monitor port can monitor a static trunk but cannot monitor a dynamic LACP trunk.

About configuring a static or dynamic 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 are using LACP. You can configure trunk group types as follows:

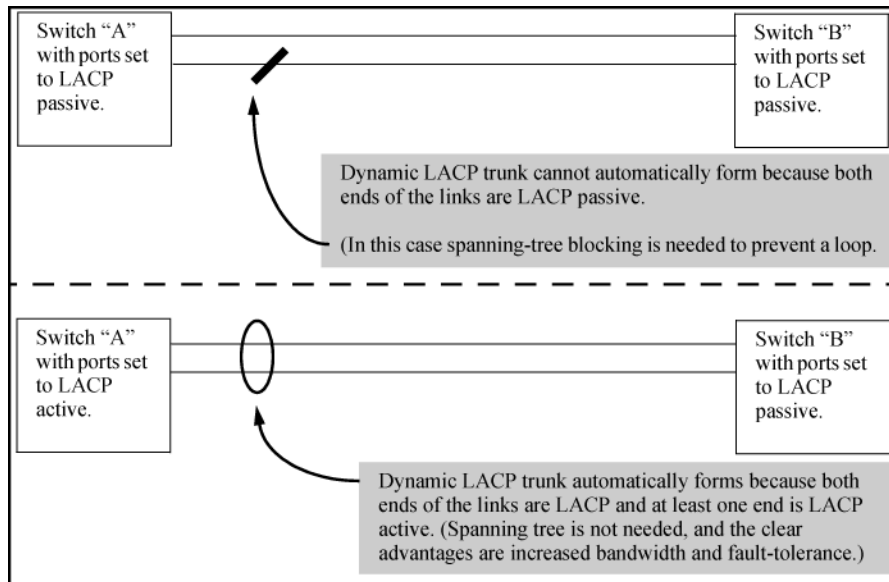
Trunk Type	Trunk Group Membership	
	TrkX (static)	DynX (dynamic)
LACP	Yes	Yes
Trunk	Yes	No

About 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 51 Criteria for automatically forming a dynamic LACP trunk



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 links are held in standby status. If a trunked link that is "Up" fails, it is replaced by a standby link, which maintains your intended bandwidth for the trunk. (See also the "Standby" entry under "Port Status" in [Table 12](#).) In the next example, ports A1 through A9 have been configured for the same LACP trunk. Notice that one of the links shows `Standby` port status, while the remaining eight links show `Up` port status.

Example

Example 38 A dynamic LACP trunk with one standby link

```
HP Switch> show lacp
```

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

Displaying LACP local information

Example 39 Example of LACP local information

```
HP Switch# show lacp local
```

LACP Local Information.

System ID: 001871-b98500

Port	Trunk	LACP Mode	Aggregated	Tx Timer	Rx Timer Expired
A2	A2	Active	Yes	Fast	No
A3	A3	Active	Yes	Fast	No

Displaying LACP peer information

Use the `show lacp peer` command to display information about LACP peers. The System ID represents the MAC address of a partner switch. It will be zero if a partner is not found.

Example 40 Example of LACP peer information

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

LACP Peer Information.

System ID: 001871-b98500

Local Port	Local Trunk	Local System ID	Port	Port Priority	Oper Key	LACP Mode	Tx Timer
A2	A2	123456-654321	2	0	100	Passive	Fast
A3	A3	234567-456789	3	0	100	Passive	Fast

Displaying LACP counters

Use the `show lacp counters` command to display statistical information about LACP ports.

Note on the Marker Protocol. Data traffic can be dynamically redistributed in port channels. This may occur when a link is added or removed, or there is a change in load-balancing. Traffic that is redistributed in the middle of a traffic flow could potentially cause mis-ordered data packets.

LACP uses the marker protocol to prevent data packets from being duplicated or reordered due to redistribution. Marker PDUs are sent on each port-channel link. The remote system responds to the marker PDU by sending a marker responder when it has received all the frames received on this link prior to the marker PDU. When the marker responders are received by the local system on all member links of the port channel, the local system can redistribute the packets in the traffic flow correctly.

For the switches covered in this guide, the marker BPDUs are not initiated, only forwarded when received, resulting in the Marker fields in the output usually displaying zeros.

Example 41 Example of LACP counters output

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

```
LACP Port Counters.
```

LACP Port	LACP Trunk	LACP PDUs Tx	Marker PDUs Rx	Marker Req. Tx	Marker Req. Rx	Marker Resp. Tx	Marker Resp. Rx	Error
A2	A2	1234	1234	0	0	0	0	0
A3	A3	1234	1234	0	0	0	0	0

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 11 \(page 119\)](#).

Table 11 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 144, depending on how many dynamic and static trunks are currently on the switch. (The switch allows a maximum of 144 trunk groups in any combination of static and dynamic trunks.)

Table 11 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 123).</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="879 793 1259 909" data-label="Diagram"> <pre> graph LR subgraph Switch1 [Switch 1] direction TB PX[Port X LACP Enable: Active] PY[Port Y LACP Enable: Active] end subgraph Switch2 [Switch 2] direction TB PA[Port A LACP Enable: Active] PB[Port B LACP Enable: Passive] end PX --- Active-to-Active PA PY --- Active-to-Passive PB </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 123).) • You want to use a monitor port on the switch to monitor an LACP trunk.

Table 11 LACP trunk types (continued)

LACP port trunk configuration	Operation
	<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 113) 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 12 \(page 121\)](#) 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 12 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</p>

Table 12 LACP port status data (continued)

Status name	Meaning
	<p>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. (See "Using the CLI To Configure a Static or Dynamic Trunk Group")

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.

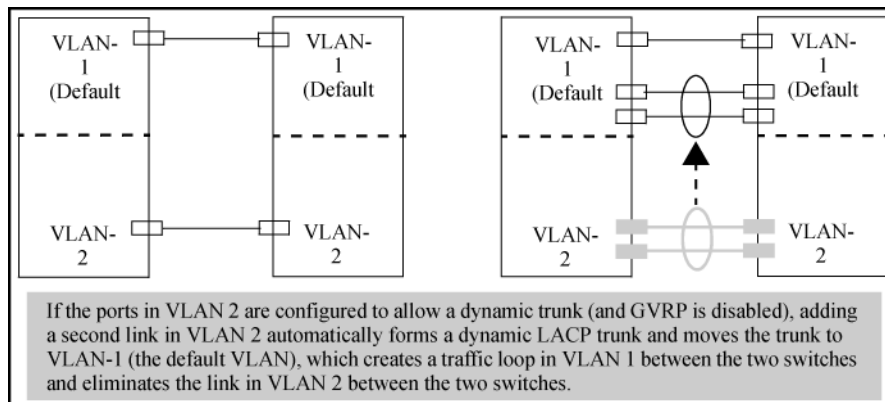
Figure 52 Blocked ports with LACP

```
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 53 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.

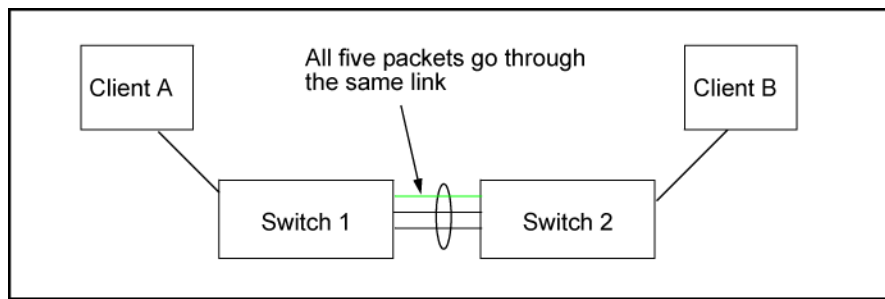
Interface option	Dynamic LACP trunk group	Static LACP trunk group	Static non-protocol
Menu interface	No	Yes	Yes
CLI <code>show trunk</code>	No	Yes	Yes
CLI <code>show interfaces</code>	No	Yes	Yes
CLI <code>show lacp</code>	Yes	Yes	No
CLI <code>show spanning-tree</code>	No	Yes	Yes
CLI <code>show igmp</code>	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 54 \(page 126\)](#). 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 54 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 55 \(page 126\)](#) showing a three-port trunk, traffic could be assigned as shown in [Table 13 \(page 126\)](#).

Figure 55 Example of port-trunked network

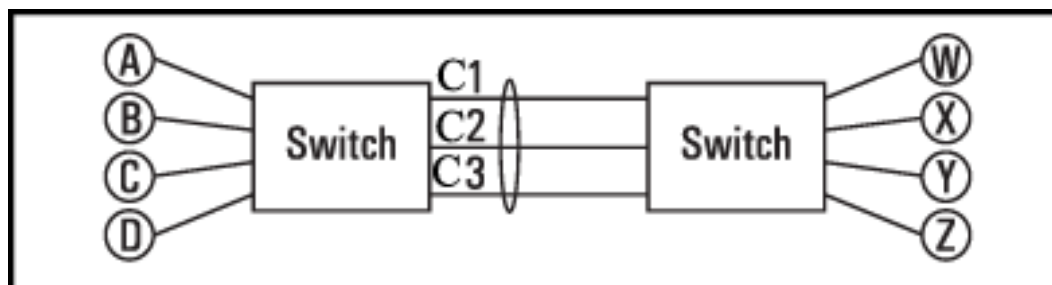


Table 13 Example of link assignments in a trunk group (SA/DA distribution)

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 Layer 4 ports

Trunk load balancing using Layer 4 ports 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 4 packets when this feature is enabled is as follows:

1. If the packet protocol is an IP packet and has Layer 4 port information, use Layer 4.
2. If the packet protocol is an IP packet and does *not* have Layer 4 information, use Layer 3 information.
3. If the packet is *not* an IP packet, use Layer 2 information.

Distributed trunking overview

The IEEE standard 802.3ad requires that all links in a trunk group originate from the same switch. Distributed trunking uses a proprietary protocol that allows two or more port trunk links distributed across two switches to create a trunk group. The grouped links appear to the downstream device as if they are from a single device. This allows third party devices such as switches, servers, or any other networking device that supports trunking to interoperate with the distributed trunking switches (DTSs) seamlessly. Distributed trunking provides device-level redundancy in addition to link failure protection.

DTSs are connected by a special interface called the InterSwitch-Connect (ISC) port. This interface exchanges information so that the DTSs appear as a single switch to a downstream device, as mentioned above. Each distributed trunk (DT) switch in a DT pair must be configured with a separate ISC link and peer-keepalive link. The peer-keepalive link is used to transmit keepalive messages when the ISC link is down to determine if the failure is a link-level failure or the complete failure of the remote peer.

The downstream device is a distributed trunking device (DTD.) The DTD forms a trunk with the DTSs. The connecting links are DT links and the ports are DT ports. A distributed trunk can span a maximum of two switches.

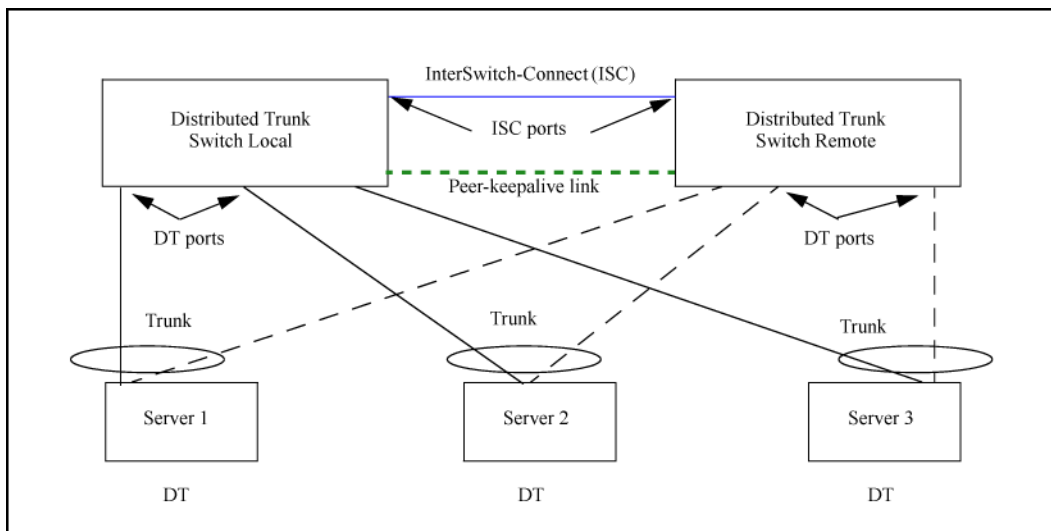
NOTE: All DT linked switches must be running the same software version.

You can group together distributed trunks by configuring two individual dt-lacp/dt-trunk trunks with the same trunk group name in each switch. The DT ports are grouped dynamically after the configuration of distributed trunking.

NOTE: Before you configure the switch, HP recommends that you review the “[Distributed trunking restrictions](#)” (page 135) for a complete list of operating notes and restrictions.

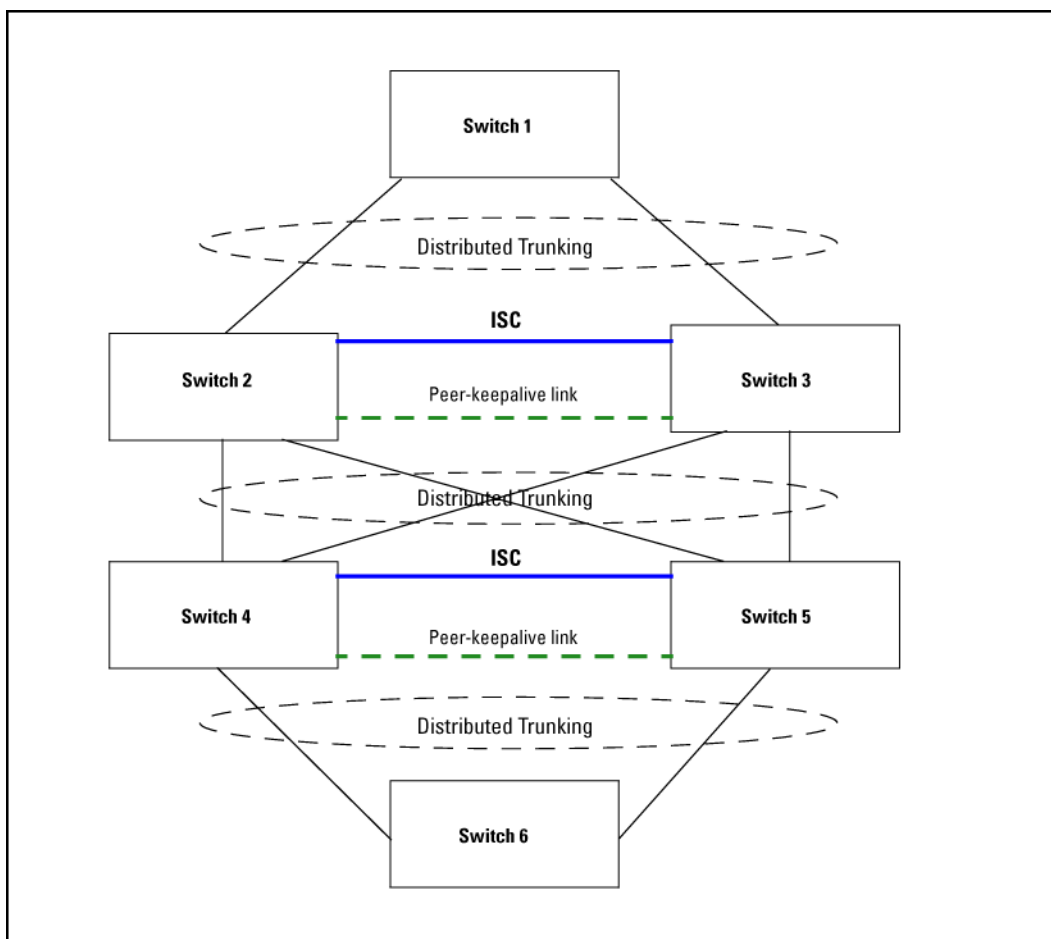
In [Figure 56 \(page 128\)](#), three different distributed trunks with three different servers have one common ISC link. Each trunk spans only two DTSs, which are connected at the ISC ports so they can exchange information that allows them to appear as one device to the server.

Figure 56 Example of distributed trunking with three different distributed trunks with three servers



An example of distributed trunking switch-to-switch in a square topology is shown in [Figure 57](#) (page 128).

Figure 57 Example of a distributed trunking switch-to-switch square topology



Distributed trunking interconnect protocol

Distributed trunking uses the distributed trunking interconnect protocol (DTIP) to transfer DT-specific configuration information for the comparison process and to synchronize MAC and DHCP snooping binding data between the two DT peer switches.

NOTE: For DHCP snooping to function correctly in a DT topology, the system time must be the same on both switches, and the ISC must be trusted for DHCP snooping.

About configuring distributed trunking

The following parameters must be configured identically on the peer devices or undesirable behavior in traffic flow may occur:

- The ISC link must have a VLAN interface configured for the same VLAN on both DT switches.
- VLAN membership for all DT trunk ports should be the same on both DT switches in a DT pair.
- IGMP-snooping or DHCP-snooping configuration on a DT VLAN should be the same on both DT switches. For example, for a DT, if IGMP-snooping or DHCP-snooping is enabled on a VLAN that has a DT port as a member port of the VLAN, the same must be configured on the peer DT on the same VLAN.
- Loop-protection configuration on a DT VLAN should be the same for both DT switches.

About configuring peer-keepalive links

Distributed trunking uses UDP-based peer-keepalive messages to determine if an ISC link failure is at the link level or the peer has completely failed. The following operating rules must be followed to use peer-keepalive links:

- An IP address must be configured for a peer-keepalive VLAN interface and the same IP address must be configured as a peer-keepalive destination on the peer DT switch.
- There must be logical Layer 3 connectivity between the two IP addresses configured for the peer-keepalive VLAN interface.
- Only peer-keepalive messages are sent over the peer-keepalive VLAN (Layer 3 link.) These messages indicate that the DT switch from which the message originates is up and running. No data or synchronization traffic is sent over the peer-keepalive VLAN.
- STP cannot run on peer-keepalive links.
- The peer-keepalive VLAN can have only one member port. If you attempt to assign a second member port to this VLAN, or if you attempt to configure a VLAN that has more than one member port as a peer-keepalive VLAN, this message displays:
A keepalive VLAN can only have one member port.
- A port cannot be a member of a regular VLAN and a peer-keepalive VLAN. An error message displays:
A port cannot simultaneously be a member of a keepalive and a non-keepalive VLAN.
- The DEFAULT VLAN cannot be a peer-keepalive VLAN. An error message displays:
The default VLAN cannot be configured as a keepalive VLAN.

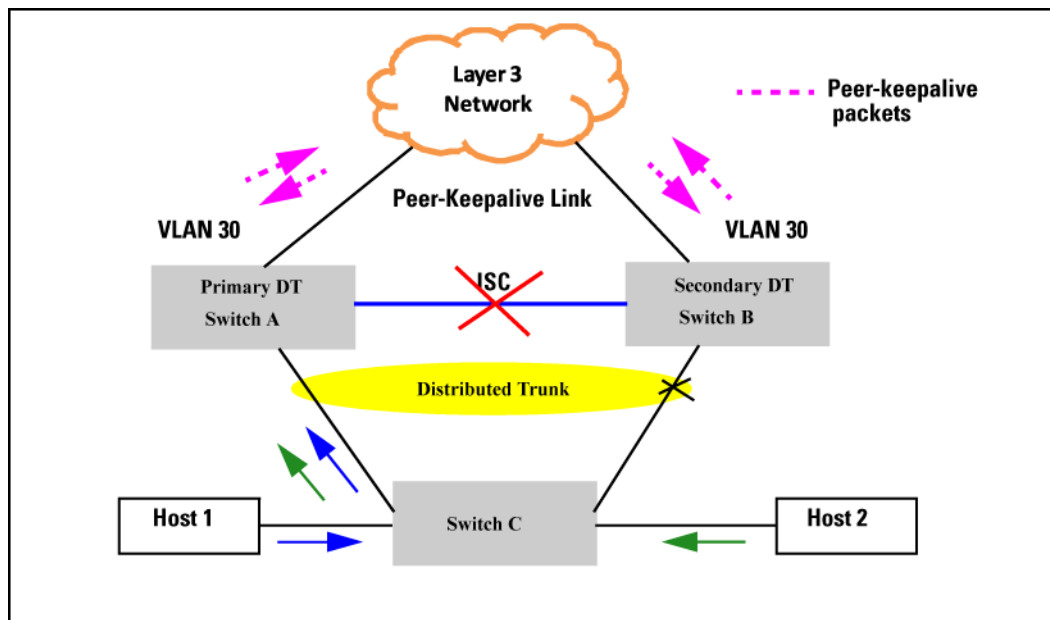
NOTE: If you are upgrading your software from a version prior to K.15.05.xxxxx with a configuration that violates any of the above operating rules, the following message displays:

```
DT: Keepalive mis-configuration detected. Reconfigure the keepalive VLAN.
```

You must then manually correct the configuration.

DT switches have an operational role that depends on the system MAC address. The bridge with the lowest system MAC address acts as the DT primary device; the other device is the DT secondary device. These roles are used to determine which device forwards traffic when the ISC link is down.

Figure 58 Example of ISC link failure with peer-keepalive



Peer-keepalive messages are sent by both the DT switches as soon as the switches detect that the ISC link is down. Peer-keepalive message transmission (sending and receiving) is suspended until the peer-keepalive hold timer expires. When the hold timer expires, the DT switches begin sending peer-keepalive messages periodically while receiving peer-keepalive messages from the peer switch. If the DT switch fails to receive any peer-keepalive messages for the timeout period, it continues to forward traffic, assuming that the DT peer switch has completely failed.

Conversely, if the failure is because the ISC link went down and the secondary DT switch receives even one peer-keepalive message from the primary peer, the secondary switch disables all its DT ports. The primary switch always forwards the traffic on its DT ports even if it receives peer-keepalive messages from the secondary DT switch.

In both situations, if the ISC link or the DT switch becomes operational, both the DT peers sync the MAC addresses learned during the failover and continue to forward traffic normally. The peer-keepalive timers and operation is halted.

Maximum DT trunks and links supported

Table 14 (page 130) shows the maximum number of DT trunks and DT links that are supported.

Table 14 Maximum supported DT trunks and links

Description	Max number
Maximum number of groups (DT trunks) in a DT switch (that is, maximum number of servers supported)	144
Maximum number of switches that can be aggregated	2
Maximum number of physical links that can be aggregated in a single switch from a server (that is, maximum number of ports that can be in a trunk connected to a single switch)	4

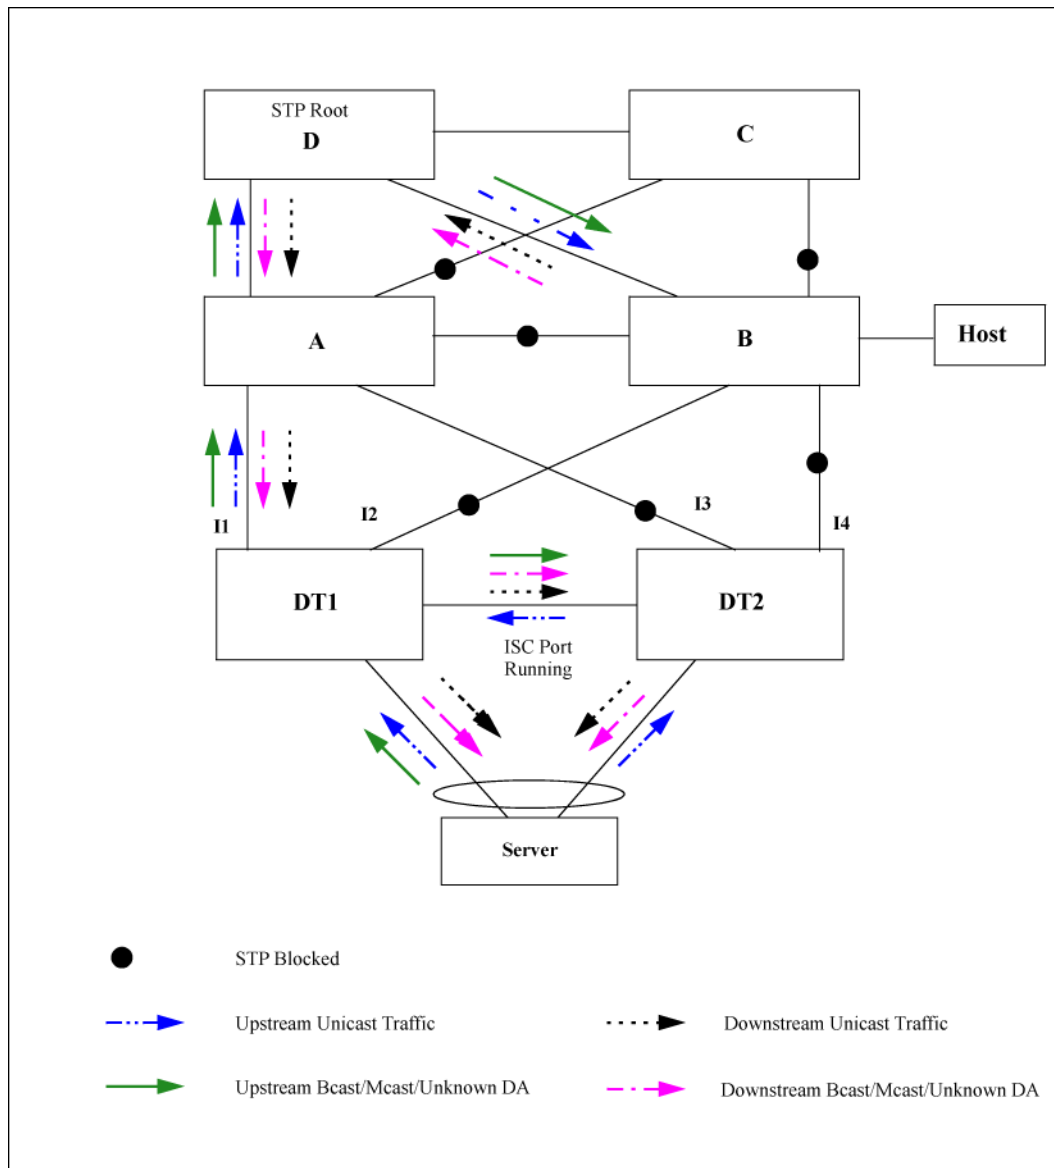
From the server perspective, this means that there could be a maximum total of 60 servers connected to two DT switches. Each server can have up to four physical links aggregated in a single switch, meaning that a single server could have a maximum of eight links (that is, four on each DT switch) in a DT trunk.

Forwarding traffic with distributed trunking and spanning tree

Refer to [Figure 59 \(page 131\)](#) for the following discussion about forwarding traffic when spanning tree is enabled. In this example, it is assumed that traffic is sent from a host off switch B to a server, and from the server back to the host. STP can block any one of the upstream links; in this example, STP has blocked all the links except the I1 link connected to DT1.

NOTE: STP is automatically disabled on the DT ports.

Figure 59 Example of distributed trunking with STP forwarding unicast, broadcast, and multicast traffic




Forwarding unicast traffic

Refer to [Figure 60 \(page 132\)](#) for the following discussion about forwarding traffic with switch-to-switch distributed trunking. Traffic from Host X or Y that is destined for Host F is always forwarded by Switch A over one of its standard 802.1AX trunk links to either Switch B or Switch C. When either Switch B or Switch C receives incoming traffic from Switch A, the traffic is directly forwarded to Switch F without traversing the ISC link.

Traffic from Host Y to Host D may go over the ISC if Switch A sends it to Switch C instead of sending it to Switch B.

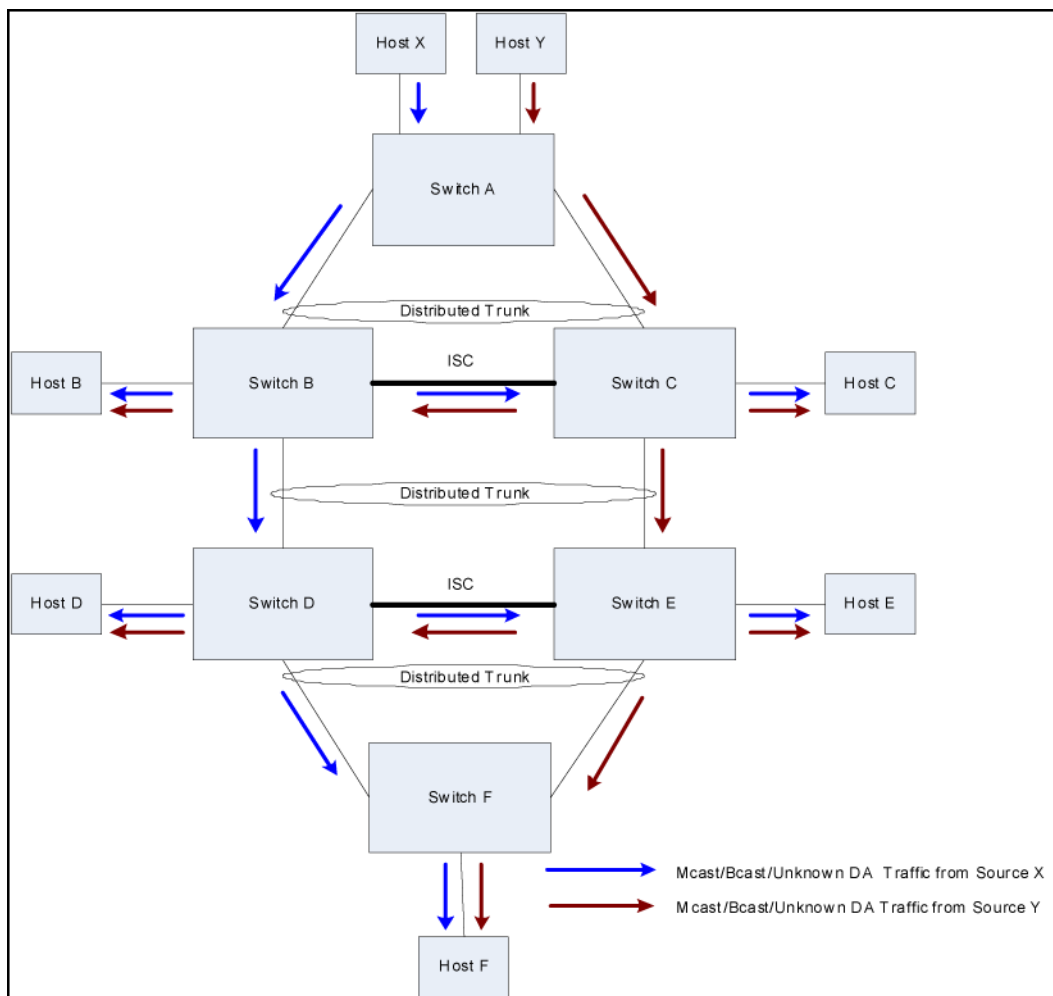
Figure 60 Unicast traffic flow across DT switches



Forwarding broadcast, multicast, and unknown traffic

In the example shown in [Figure 61 \(page 133\)](#), multicast/broadcast/unknown traffic from Host X or Y is always forwarded by Switch A over one of its standard 802.3ad trunk links to either Switch B or C. Switch B or C forwards the traffic on all the links including the ISC port, but not on the port that the traffic was received on. The peer DT switch (B or C) that receives broadcast/multicast/unknown traffic over the ISC port does not forward the packets to any of the DT trunks; the packet is sent only over the non-DT ports. The one exception is if the DT trunk on the peer aggregation device is down, then traffic received over the ISC is forwarded to the corresponding DT trunk.

Figure 61 Broadcast/multicast/unknown traffic flow access DT switches



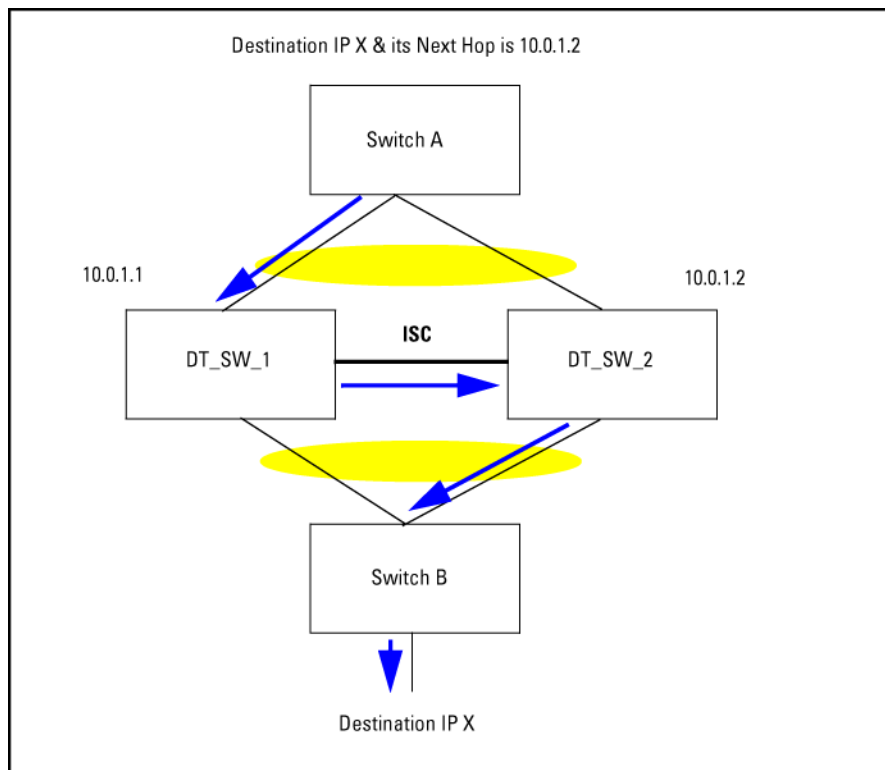
IP routing and distributed trunking

In switch-to-switch distributed trunking, the peer DT switches behave like independent Layer 3 devices with their own IP addresses in each active VLAN. If a DT switch receives a packet destined for the peer DT switch, it switches the packet through the ISC link. Interfaces on a VLAN using DT typically use a single default gateway pointing to only one of the DT switches in a DT pair.

The example in [Figure 62 \(page 134\)](#) shows Layer 3 (IP unicast) forwarding in a DT topology. The packet is sent as follows:

1. Switch A selects the link (using the trunk hash) to the DT pair. The packet is sent to the selected link DT_SW_1.
2. When DT_SW_1 receives the packet, it determines, based on the MAC address, that the packet must be sent over the ISC link to DT_SW_2.
3. When the packet arrives, DT_SW_2 performs a lookup and determines that the packet needs to be sent to Switch B.

Figure 62 Example 1 of layer 3 forwarding (IP unicast) in DT topology

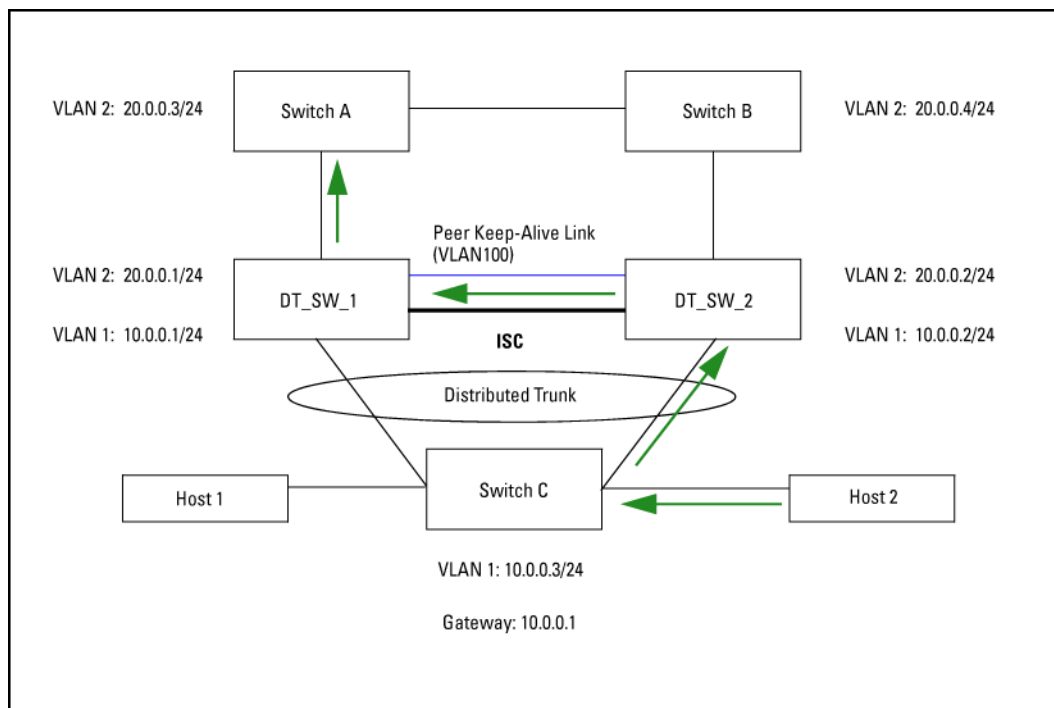


Another example in [Figure 63](#) shows Layer 3 (IP unicast) forwarding in a DT topology. The packet is sent as follows:

1. Host 2 sends a packet to Switch C.
2. Switch C performs a lookup in the routing table and determines that the default gateway IP address is 10.0.0.1.
3. Layer 2 lookup determines that the outgoing interface is the DT port.
4. Hashing determines that the trunk member chosen is DT_SW_2 and the packet is sent there.
5. DT_SW_2 determines that the packet needs to be sent over the ISC link to DT_SW_1 based on the MAC address.
6. DT_SW_1 performs a lookup and determines that the packet goes to Switch A.

The packet is only forwarded if the outgoing interface is not a DT port, or if the outgoing DT port does not have an active interface on the peer switch.

Figure 63 Example 2 of layer 3 forwarding (IP unicast) in DT topology



Distributed trunking restrictions

There are several restrictions with distributed trunking:

Beginning with software version K.15.07, the switch will not allow both Distributed Trunking and MAC-based mirroring to function simultaneously. The switch will respond as follows:

- If the user attempts to configure both, an error message will appear.
- When a switch is updated from older software to K.15.07, if the older config file has both Distributed Trunking and MAC-based mirroring, the switch will automatically remove the MAC-based mirroring lines from the config file, and will give an explanatory error message.
- If a switch is running K.15.07 and an existing config file that has both Distributed Trunking and MAC-based mirroring is loaded onto the switch, the switch will automatically remove the MAC-based mirroring lines from the config file, and will give an explanatory error message.
- All DT linked switches must be running the same software version.
- The port trunk links should be configured manually (using manual LACP or manual trunks.) Dynamic linking across switches is not supported.
- A distributed trunk can span a maximum of two switches.
- A maximum total of 144 servers can be connected to two DT switches. Each server can have up to four physical links aggregated in a single switch, meaning that there can be a maximum of eight ports (four aggregated links for each DT switch) included in a DT trunk.
- Only one ISC link is supported per switch, with a maximum of 60 DT trunks supported on the switch. The ISC link can be configured as a manual LACP trunk, non-protocol trunk, or as an individual link. Dynamic LACP trunks are not supported as ISCs.
- An ISC port becomes a member of all VLANs that are configured on the switch. When a new VLAN is configured, the ISC ports become members of that VLAN.
- Port trunk links can be done only on a maximum of two switches that are connected to a specific server.
- Any VLAN that is in a distributed trunk must be configured on both switches. By default, the distributed trunk belongs to the default VLAN.

- There can be eight links in a distributed trunk grouped across two switches, with a limit of four links per distributed trunking switch.
- The limit of 144 manual trunks per switch includes distributed trunks as well.
- ARP protection is not supported on the distributed trunks.
- Dynamic IP Lockdown protection is not supported on the distributed trunks.
- QinQ in mixed VLAN mode and distributed trunking are mutually exclusive.
- Features not supported include:
 - SVLANs in mixed mode on DT or ISC links
 - Meshing
 - CDP
 - GVRP
 - Multicast routing
 - IPv6 routing
 - MLD/MLD snooping on DT VLAN

DT operating notes when updating software versions

When updating software from a version that does not support DT Keepalive (prior to version K.15.03) to a version that supports shared DT keepalive (K.15.03 and greater), use the following procedure:

1. Disable the ISC interface on both switches, and then upgrade the software. Assume a2 is configured as switch-interconnect.


```
HP Switch(config)# int a2 disable
HP Switch(config)# write mem
```
2. Configure one of the existing uplink VLANs as a keepalive VLAN, and then configure the destination keepalive IP address (peer's keepalive IP address) on both switches at bootup.


```
HP Switch(config)# distributed-trunking
peer-keepalive vlan 2
HP Switch(config)# distributed-trunking
peer-keepalive destination 20.0.0.2
```
3. Ping the keepalive destination address to make sure that there is connectivity between the two DT switches (keepalive VLANs.)
4. Enable the ISC link on both switches and then execute `write memory`. Assume a2 is configured as switch-interconnect.


```
HP Switch(config)# int a2 enable
HP Switch(config)# write mem
```

When updating software from a software version that does not support DT keepalive (prior to version K.15.03) to a version with dedicated point-to-point keepalive (K.15.03 and greater), use the following procedure:

1. Disable the ISC interface on both switches, and then upgrade the software. Assume a2 is configured as switch-interconnect.


```
HP Switch(config)# int a2 disable
HP Switch(config)# write mem
```
2. At switch bootup, create a dedicated VLAN for keepalive, and assign only the keepalive link port as a member port of the VLAN. Configure the keepalive destination IP address.


```
HP Switch(config)# distributed-trunking
peer-keepalive vlan 2
```



```
HP Switch(config)# distributed-trunking
peer-keepalive destination 20.0.0.2
```

3. Ping the keepalive destination address to make sure that there is connectivity between the two DT switches (keepalive VLANs.)
4. Enable the ISC link on both switches, and then execute `write memory`. Assume a2 is configured as switch-interconnect.

```
HP Switch(config)# int a2 enable
HP Switch(config)# write mem
```

When updating software from a software version that does support shared DT keepalive (K.15.03, K.15.04) to a version that supports dedicated point-to-point keepalive (K.15.05), use the following procedure:

1. Disable the ISC interface and undo the keepalive configuration on both switches. Ignore the warning message that is displayed by the keepalive command while undoing the configuration. Upgrade the software. Assume a2 is configured as switch-interconnect.

```
HP Switch(config)# int a2 disable
HP Switch(config)# no distributed-trunking
peer-keepalive vlan
HP Switch(config)# write mem
```

2. At switch bootup, create a dedicated VLAN for keepalive and assign only the keepalive link port as a member port of the VLAN. Configure the keepalive destination IP address.

```
HP Switch(config)# vlan 10 (dedicated point-to-point VLAN interface)
HP Switch(vlan-10)#
HP Switch(vlan-10)# untagged b2 (keepalive link port)
HP Switch(vlan-10)# ip address 10.0.0.1/24
HP Switch(vlan-10)# exit
HP Switch(config)# distributed-trunking
peer-keepalive vlan 10
HP Switch(config)# distributed-trunking
peer-keepalive destination 10.0.0.2
```

3. Ping the keepalive destination address to make sure that there is connectivity between the two DT switches (keepalive VLANs.)
4. Enable the ISC link on both switches, and then execute `write memory`. Assume a2 is configured as switch-interconnect.

```
HP Switch(config)# int a2 enable
HP Switch(config)# write mem
```

5 Port Traffic Controls

Rate-limiting

In earlier releases, all traffic rate-limiting applied to inbound traffic only, and was specified as a percentage of total bandwidth. Beginning with software release K.12.xx or later, it is also possible to configure outbound rate-limiting for all traffic on a port and specify bandwidth usage in terms of kilobits per second (kbps.)

- △ **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.

Configuring rate-limiting

Syntax

```
[no] int port-list rate-limit all [ in | out ] 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 either inbound or outbound traffic.
- Specifying the traffic rate as either a percentage of bandwidth, or in terms of bits per second.

(Default: Disabled.)

<code>in or out</code>	Specifies a traffic rate limit on inbound traffic passing through that port, or on outbound traffic.
<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 [“All traffic rate-limiting” \(page 151\)](#).

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

If you want to view the rate-limiting configuration on the first six ports in the module in slot "A":

Figure 64 Listing the rate-limit configuration

```
HP Switch# show rate-limit all a1-a6
```

Ports A1-A4 are configured with an outbound rate limit of 200 Kbps; Port A5 is configured with an inbound rate limit of 20%. (Port A6 is not configured for rate-

All-Traffic Rate Limit Maximum %						
Port	Inbound			Outbound		
	Limit	Mode	Radius Override	Limit	Mode	Radius Override
A1	Disabled	Disabled	No-override	200	kbps	No-override
A2	Disabled	Disabled	No-override	200	kbps	No-override
A3	Disabled	Disabled	No-override	200	kbps	No-override
A4	Disabled	Disabled	No-override	200	kbps	No-override
A5	20	%	No-override	Disabled	Disabled	No-override
A6	Disabled	Disabled	No-override	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
```

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.)

Figure 65 Example of rate-limit settings listed in the `show config` output

```

HP Switch(config)# show config

Startup configuration:

; J8697A Configuration Editor; Created on release #K.14.01

hostname "HP Switch 8212z1"
module 1 type J8705A
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24
  ip address dhcp-bootp
  exit
interface A1
  rate-limit all out kbps 200
  exit
interface A2
  rate-limit all out kbps 200
  exit
interface A3
  rate-limit all out kbps 200
  exit
interface A4
  rate-limit all out kbps 200
  exit
interface A5
  rate-limit all in percent 200
  exit
interface A6
  rate-limit icmp percent 60
  rate-limit mcast in percent 60
  exit

```

Ports A1-A4 are configured with an outbound rate limit of 200 kbps.

Port A5 is configured with an inbound rate limit of 200 kbps.

Port A6 is configured with an inbound ICMP and multicast rate-limits of 60 percent each.

Configuring ICMP rate-limiting

For detailed information about ICMP rate-limiting, see [“ICMP rate-limiting” \(page 153\)](#).

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 ]
```

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.)

percent 1-100	Values in this range allow ICMP traffic as a percentage of the bandwidth available on the interface.
kbps 0-10000000	Specifies the rate at which to forward traffic in kilobits per second.
0	Causes an interface to drop all incoming ICMP traffic and is not recommended. See the <i>Caution</i> on 154.

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
```

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 155).

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

Example

If you want to view the rate-limiting configuration on the first six ports in the module in slot "B":

Figure 66 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   | kbps      100
B3   | %         5
B4   | %         1
B5   | %         1
B6   | Disabled  Disabled
```

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 156).

Resetting the ICMP trap function of the port

The port ICMP trap function can be reset through SNMP from a network management station or through the CLI with the `setmib` command. For information on ICMP rate-limiting trap and Event Log messages, see [“ICMP rate-limiting trap and Event Log messages”](#) (page 157).

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:

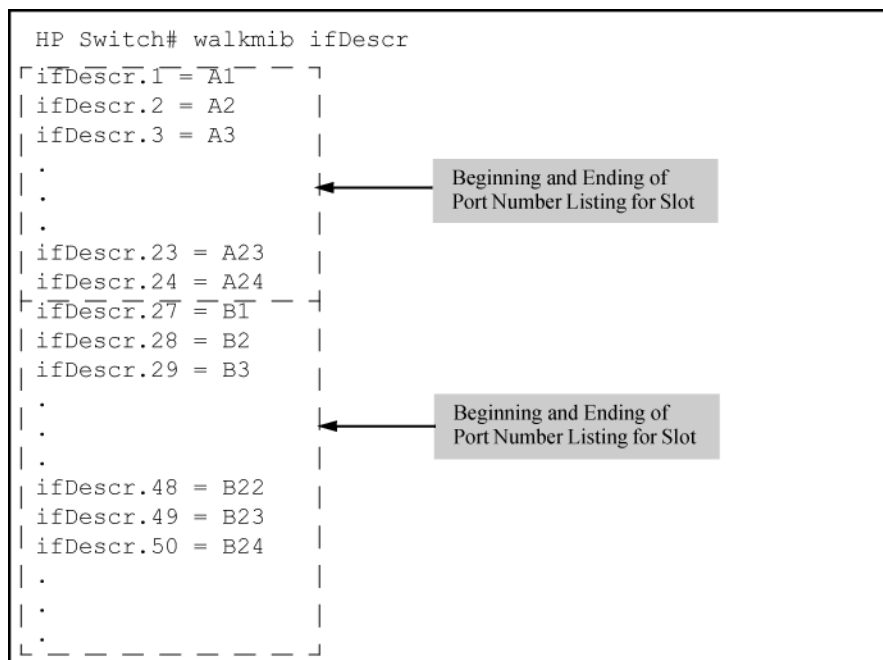
```
HP Switch(config)# setmib hpicmpratelimitportalarmflag.1 -i 1
```

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 157). 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 67 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 │
│ . │
│ . │
└──────────────────┘
```

Configuring a broadcast limit on the switch

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

```
HP Switch(config) # int B1
HP Switch(int B1) # broadcast-limit 1
```

Syntax

```
broadcast-limit [0-99]
```

Enables or disables broadcast limiting for outbound broadcasts on a selected port on the switch.

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

NOTE: You must switch to port context level before issuing the `broadcast-limit` command. This feature is not appropriate for networks requiring high levels of IPX or RIP broadcast traffic.

Syntax

```
show config
```

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

Syntax

```
show running-config
```

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

Example

The following command enables broadcast limiting of 1% of the traffic rate on the selected port on the switch:

```
HP Switch(int B1) # broadcast-limit 1
```

For a 1-Gbps port, this results in a broadcast traffic rate of 10 Mbps.

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 bcast in percent 10
```

or

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

Syntax

```
rate-limit [ bcast | mcast ] in percent 0-100
[no]
rate-limit [ bcast | [mcast ] ] in
```

Enables rate-limiting and sets limits for the specified inbound broadcast or multicast traffic. Only the amount of traffic specified by the percent is forwarded.

Default: Disabled

Example

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 [Figure 68](#). Only 50% of the inbound broadcast traffic will be forwarded.

Figure 68 Inbound broadcast rate-limiting of 50% on port 3

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

HP Switch(eth-3)# show rate-limit bcast
Broadcast-Traffic Rate Limit Maximum %
```

Port	Inbound Limit	Mode	Radius Override
1	Disabled	Disabled	No-override
2	Disabled	Disabled	No-override
3	50	%	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 [Figure 69](#). Only 20% of the multicast traffic will be forwarded.

Figure 69 Example of 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 [Figure 70](#).

Figure 70 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.

Configuring Guaranteed Minimum Bandwidth (GMB) for outbound traffic

For any port or group of ports you can configure either the default minimum bandwidth settings for each outbound priority queue or a customized bandwidth allocation. For most applications, HP recommends configuring GMB with the same values on all ports on the switch so that the outbound traffic profile is consistent for all outbound traffic. However, there may be instances

where it may be advantageous to configure special profiles on connections to servers or to the network infrastructure (such as links to routers, other switches, or to the network core.)
For detailed information about GMB, see [“Guaranteed minimum bandwidth \(GMB\)” \(page 157\)](#).

Syntax

```
[no] int port-list bandwidth-min output
```

Configures the default minimum bandwidth allocation for the outbound priority queue for each port in the *port-list*. In the eight-queue configuration, the default values per priority queue are:

- Queue 1 (low priority): 2%
- Queue 2 (low priority): 3%
- Queue 3 (normal priority): 30%
- Queue 4 (normal priority): 10%
- Queue 5 (medium priority): 10%
- Queue 6 (medium priority): 10%
- Queue 7 (high priority): 15%
- Queue 8 (high priority): 20%

The `no` form of the command disables GMB for all ports in the *port-list*. In this state, which is the equivalent of setting all outbound queues on a port to **0** (zero), a high level of higher-priority traffic can starve lower-priority queues, which can slow or halt lower-priority traffic in the network.

You can configure bandwidth minimums from either the global configuration level (as shown above) or from the port context level. For information on outbound port queues, see [“Per-port outbound priority queues” \(page 157\)](#).

Syntax

```
[no] int port-list bandwidth-min output [ 0-100 | strict ]
```

```
[0-100]
```

Select a minimum bandwidth.

For ports in *port-list*, specifies the minimum outbound bandwidth as a percent of the total bandwidth for each outbound queue. The queues receive service in descending order of priority.

You must specify a bandwidth percent value for all except the highest priority queue, which may instead be set to "strict" mode. The sum of the bandwidth percentages below the top queue cannot exceed 100%. (**0** is a value for a queue percentage setting.)

Configuring a total of less than 100% across the eight queues results in unallocated bandwidth that remains harmlessly unused unless a given queue becomes oversubscribed. In this case, the unallocated bandwidth is apportioned to oversubscribed queues in descending order of priority. For example, if you configure a minimum of 10% for queues 1 to 7 and 0% for queue 8, the unallocated bandwidth is available to all eight queues in the following prioritized order:

Queue 8 (high priority)
Queue 7 (high priority)
Queue 6 (medium priority)
Queue 5 (medium priority)
Queue 4 (normal priority)
Queue 3 (normal priority)
Queue 2 (low priority)
Queue 1 (low priority)

A setting of **0** (zero percent) on a queue means that no bandwidth minimum is specifically reserved for that queue for each of the ports in the *port-list*.

Also, there is no benefit to setting the high-priority queue (queue 8) to **0** (zero) unless you want the medium queue (queue 4) to be able to support traffic bursts above its guaranteed minimum.

[strict]

Provides the ability to configure the highest priority queue as strict. Per-queue values must be specified in priority order, with queue 1 having the lowest priority and queue 8 (or 4, or 2) having the highest priority (the highest queue is determined by how many queues are configured on the switch. Two, four, and eight queues are permitted. (See the `qos queue-config` command.)

The strict queue is provided all the bandwidth it needs. Any remaining bandwidth is shared among the non-strict queues based on need and configured bandwidth profiles (the profiles are applied to the leftover bandwidth in this case.) The total sum of percentages for non-strict queues must not exceed 100.

NOTE: Configuring 0% for a queue can result in that queue being starved if any higher queue becomes over-subscribed and is then given all unused bandwidth.

The switch applies the bandwidth calculation to the link speed the port is currently using. For example, if a 10/100 Mbs port negotiates to 10 Mbps on the link, it bases its GMB calculations on 10 Mbps, not 100 Mbps.

Use `show bandwidth output port-list` to display the current GMB configuration. (The `show config` and `show running` commands do not include GMB configuration data.)

Example

For example, suppose you want to configure the following outbound minimum bandwidth availability for ports A1 and A2:

Priority of outbound port queue	Minimum bandwidth %	Effect on outbound bandwidth allocation
8	20%	Queue 8 has the first priority use of all outbound bandwidth not specifically allocated to queues 1 to 7. If, for example, bandwidth allocated to queue 5 is not being used and queues 7 and 8 become oversubscribed, queue 8 has first-priority use of the unused bandwidth allocated to queue 5.
7	15%	Queue 7 has a GMB of 15% available for outbound traffic. If queue 7 becomes oversubscribed and queue 8 is not already using all of the unallocated bandwidth, queue 7 can use the unallocated bandwidth. Also, any unused bandwidth allocated to queues 6 to queue 1 is available to queue 7 if queue 8 has not already claimed it.
6	10%	Queue 6 has a GMB of 10% and, if oversubscribed, is subordinate to queues 8 and 7 in priority for any unused outbound bandwidth available on the port.
5	10%	Queue 5 has a GMB of 10% and, if oversubscribed, is subordinate to queues 8, 7, and 6 for any unused outbound bandwidth available on the port.
4	10%	Queue 4 has a GMB of 10% and, if oversubscribed, is subordinate to queues, 8, 7, 6, and 5 for any unused outbound bandwidth available on the port.
3	30%	Queue 3 has a GMB of 30% and, if oversubscribed, is subordinate to queues, 8, 7, 6, 5, and 4 for any unused outbound bandwidth available on the port.

Priority of outbound port queue	Minimum bandwidth %	Effect on outbound bandwidth allocation
2	3%	Queue 2 has a GMB of 3% and, if oversubscribed, is subordinate to queues, 8, 7, 6, 5, 4, and 3 for any unused outbound bandwidth available on the port.
1	2%	Queue 1 has a GMB of 2% and, if oversubscribed, is subordinate to all the other queues for any unused outbound bandwidth available on the port.

Either of the following commands configures ports A1 through A5 with bandwidth settings:

```
HP Switch(config) # int a1-a5 bandwidth-min output 2 3 30 10 10 10 15 strict
HP Switch(eth-A1-A5) # bandwidth-min output 2 3 30 10 10 10 15 strict
```

Viewing the current GMB configuration

This command displays the per-port GMB configuration in the `running-config` file.

Syntax

```
show bandwidth output [port-list]
```

Without `port-list`, this command lists the GMB configuration for all ports on the switch.

With `port-list`, this command lists the GMB configuration for the specified ports.

This command operates the same way in any CLI context. If the command lists `Disabled` for a port, there are no bandwidth minimums configured for any queue on the port. (See the description of the `no` form of the `bandwidth-min output` command on page 13-24.)

Example

To display the GMB configuration resulting from either of the above commands:

Figure 71 Listing the GMB configuration

```
HP Switch(config)# show bandwidth output a1-a5
```

```
Outbound Guaranteed Minimum Bandwidth %

Port   Q1  Q2   Q3  Q4   Q5  Q6  Q7  Q8
-----
A1     2   3    30  10   10  10  15  strict
A2     2   3    30  10   10  10  15  strict
A3     2   3    30  10   10  10  15  strict
A4     2   3    30  10   10  10  15  strict
A5     2   3    30  10   10  10  15  strict
```

This is how the preceding listing of the GMB configuration would appear in the `startup-config` file.

Figure 72 GMB settings listed in the `show config output`

```
HP Switch(config)# show config status
Running configuration is same as the startup configuration
HP Switch(config)# show config

Startup configuration:
; J9091A configuration Editor; Created on release #K.15.05.0000x

hostname "HP Switch"
module 1 type J8697A
snmp-server community "public" Unrestricted
vlan 1
    name "DEFAULT_VLAN"
    untagged A1-A24
    ip address dhcp-bootp
    exit
interface A1
    bandwidth-min output 2 3 30 10 10 10 15 strict
    exit
interface A2
    bandwidth-min output 2 3 30 10 10 10 15 strict
    exit
interface A3
    bandwidth-min output 2 3 30 10 10 10 15 strict
    exit
interface A4
    bandwidth-min output 2 3 30 10 10 10 15 strict
    exit
interface A5
    bandwidth-min output 2 3 30 10 10 10 15 strict
    exit
```

Configuring jumbo frame operation

For detailed information about jumbo frames, see [“Jumbo frames” \(page 159\)](#).

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” \(page 150\)](#).) See [Figure 73](#).

Figure 73 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 74.)

Figure 74 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 75.)

Figure 75 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.

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)

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” \(page 159\)](#).

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.

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 or outbound 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.

⚠ CAUTION: Rate-limiting is intended for use on edge ports in a network. HP does not recommend it 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.

NOTE: Rate-limiting also can be applied by a RADIUS server during an authentication client session. (See the *Access Security Guide*.)

The switches also support ICMP rate-limiting to mitigate the effects of certain ICMP-based attacks. For more information, see [“ICMP rate-limiting” \(page 153\)](#).

The mode using bits per second (bps) in releases before K.12.XX has been replaced by the kilobits per second (kbps) mode. Switches that have configurations with bps values are automatically converted when you update your software to the new version. However, you must manually update to kbps values an older config file that uses bps values or it will not load successfully onto a switch running later versions of the software (K.12.XX or greater.)

- The `rate-limit icmp` command specifies a rate limit on inbound ICMP traffic only (See [“ICMP Rate-Limiting”](#) on page 13-9)
- 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 100 Kbps, the value would be 1 Mbps.
 - A 1 to 100 Kbps rate-limit is implemented as a limit of 100 Kbps
 - A limit of 100 to 199 Kbps is also implemented as a limit of 100 Kbps.
 - A limit of 200 to 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%, the traffic flow through that port is limited to no more than 50 Mbps. Similarly, if the same port negotiates a 10 Mbps link, 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 Switch recommends that you use 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.

Example

Either of the following commands configures an inbound rate limit of 60% on ports A3 to A5:

```
HP Switch (config) # int a3-a5 rate-limit all in percent 60
HP Switch (eth-A3-A5)# rate-limit all in percent 60
```

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 for inbound and outbound traffic are separate features. The rate limits for each direction of traffic flow on the same port are configured separately—even the specified limits can be different.
- Rate-limiting and hardware: The hardware will round the actual Kbps rate down to the nearest multiple of 64 Kbps.
- Rate-limiting is visible as an outbound forwarding rate. Because inbound rate-limiting is performed on packets during packet-processing, it is not shown via the inbound drop counters. Instead, this limit is verifiable as the ratio of outbound traffic from an inbound rate-limited port versus the inbound rate. For outbound rate-limiting, the rate is visible as the percentage of available outbound bandwidth (assuming that the amount of requested traffic to be forwarded is larger than the rate-limit.)
- 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.)

In another type of situation, an outbound port can become oversubscribed by traffic received from multiple rate-limited ports. In this case, the actual rate for traffic on the rate-limited ports may be lower than configured because the total traffic load requested to the outbound port exceeds the port's bandwidth, and thus some requested traffic may be held off on inbound.
- 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

As of software version K.15.02.0004, ICMP rate-limiting and classifier-based-rate-limiting operates on the entire packet length instead of just the IP payload part of the packet. As a result, the effective metering rate is now the same as the configured rate. The rate-limiting applies to these modules:

HP device	Product number	Minimum supported software version
HP Switch 24-port 10/100/1000 PoE+ v2 zl Module	J9534A	K.15.02.0004
HP Switch 20-port 10/100/1000 PoE+ / 4-port SFP v2 zl Module	J9535A	K.15.02.0004
HP Switch 20-port 10/100/1000 PoE+ / 2-port 10-GbE SFP+ v2 zl Module	J9536A	K.15.02.0004
HP Switch 24-port SFP v2 zl Module	J9537A	K.15.02.0004
HP Switch 8-port 10-GbE SFP+ v2 zl Module	J9538A	K.15.02.0004
HP 24-port 10/100 PoE+ v2 zl Module	J9547A	K.15.02.0004
HP 20-port Gig-T / 2-port 10-GbE SFP+ v2 zl Module	J9548A	K.15.02.0004
HP 20-port Gig-T / 4-port SFP v2 zl Module	J9549A	K.15.02.0004

HP device	Product number	Minimum supported software version
HP 24-port Gig-T v2 zl Module	J9550A	K.15.02.0004
HP 12-port Gig-T / 12-port SFP v2 zl Module	J9637A	K.15.02.0004

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: This feature should not be used to remove all ICMP traffic from a network. 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 (1 Gbps or 10 Gbps speeds) to be used for ICMP traffic.

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.

Beginning with software release K.12.xx or later, the all-traffic rate-limiting command (`rate-limit all`) and the ICMP rate-limiting command (`rate-limit icmp`) operate differently:

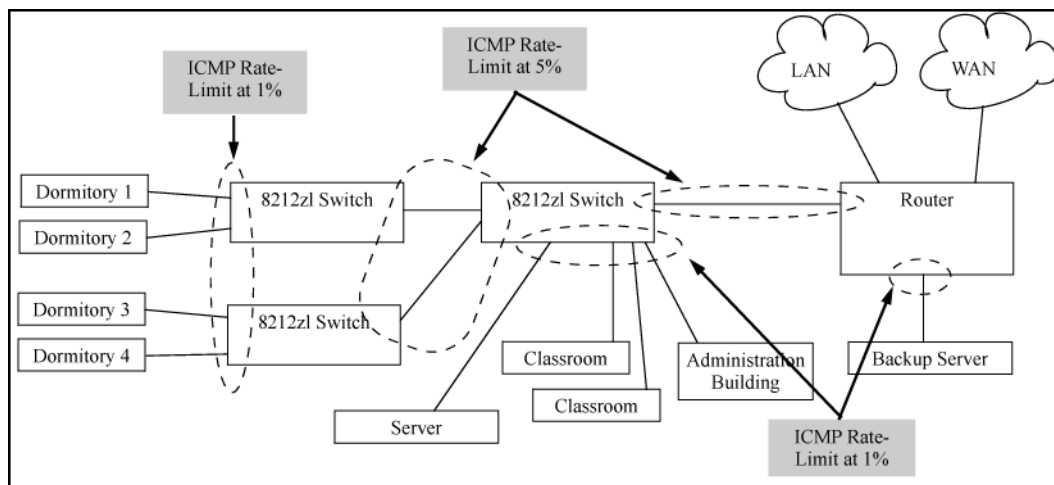
- All-traffic rate-limiting applies to both inbound and outbound traffic and can be specified either in terms of a percentage of total bandwidth or in terms of bits per second;
- ICMP rate-limiting applies only to inbound traffic and can be specified as only a percentage of total bandwidth.

ICMP rate-limiting is not supported on meshed ports. (Rate-limiting can reduce the efficiency of paths through a mesh domain.)

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 76 \(page 155\)](#) 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 76 Example of ICMP rate-limiting



NOTE: When using kbps-mode ICMP rate-limiting, the rate-limiting operates on only 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.

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 port'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.

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.

Guaranteed minimum bandwidth (GMB)

GMB 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. You can configure GMB per-port.

GMB operation

The switch services per-port outbound traffic in a descending order of priority; that is, from the highest priority to the lowest priority. By default, each port offers eight prioritized, outbound traffic queues. Tagged VLAN traffic is prioritized according to the 802.1p priority the traffic carries. Untagged VLAN traffic is assigned a priority of **0** (normal.)

Table 15 Per-port outbound priority queues

802.1p Priority settings in tagged VLAN packets ¹	Outbound priority queue for a given port
1 (low)	1
2 (low)	2
0 (normal)	3
3 (normal)	4

Table 15 Per-port outbound priority queues (continued)

802.1p Priority settings in tagged VLAN packets ¹	Outbound priority queue for a given port
4 (medium)	5
5 (medium)	6
6 (high)	7
7 (high)	8

¹ The switch processes outbound traffic from an untagged port at the "0" (normal) priority level.

You can use GMB to reserve a specific percentage of each port's available outbound bandwidth for each of the eight priority queues. This means that regardless of the amount of high-priority outbound traffic on a port, you can ensure that there will always be bandwidth reserved for lower-priority traffic.

Since the switch services outbound traffic according to priority (highest to lowest), the highest-priority outbound traffic on a given port automatically receives the first priority in servicing. Thus, in most applications, it is necessary only to specify the minimum bandwidth you want to allocate to the lower priority queues. In this case, the high-priority traffic automatically receives all unassigned bandwidth without starving the lower-priority queues.

Conversely, configuring a bandwidth minimum on only the high-priority outbound queue of a port (and not providing a bandwidth minimum for the lower-priority queues) is not recommended, because it may "starve" the lower-priority queues. (See the [158](#).)

NOTE: For a given port, when the demand on one or more outbound queues exceeds the minimum bandwidth configured for those queues, the switch apportions unallocated bandwidth to these queues on a priority basis. As a result, specifying a minimum bandwidth for a high-priority queue but not specifying a minimum for lower-priority queues can starve the lower-priority queues during periods of high demand on the high priority queue. For example, if a port configured to allocate a minimum bandwidth of 80% for outbound high-priority traffic experiences a demand above this minimum, this burst starves lower-priority queues that *do not have a minimum configured*. Normally, this will not altogether halt lower priority traffic on the network, but will likely cause delays in the delivery of the lower-priority traffic.

The sum of the GMB settings for all outbound queues on a given port cannot exceed 100%.

Impacts of QoS queue configuration on GMB operation

The section on "[Configuring Guaranteed Minimum Bandwidth \(GMB\) for outbound traffic](#)" (page 144) assumes the ports on the switch offer eight prioritized, outbound traffic queues. This may not always be the case, however, because the switch supports a QoS queue configuration feature that allows you to reduce the number of outbound queues from eight (the default) to four queues, or two.

Changing the number of queues affects the GMB commands (`interface bandwidth-min` and `show bandwidth output`) such that they operate only on the number of queues currently configured. If the queues are reconfigured, the guaranteed minimum bandwidth per queue is automatically re-allocated according to the following percentages:

Table 16 Default GMB percentage allocations per QoS queue configuration

802.1p priority	8 queues (default)	4 queues	2 queues
1 (lowest)	2%	10%	90%
2	3%		
0 (normal)	30%	70%	
3	10%		
4	10%	10%	10%
5	10%		
6	15%	10%	
7 (highest)	20%		

NOTE: For more information on queue configuration and the associated default minimum bandwidth settings, (see the *Advanced Traffic Management Guide*.)

Impact of QoS queue configuration on GMB commands.

Changing the number of queues causes the GMB commands (`interface bandwidth-min` and `show bandwidth output`) to operate only on the number of queues currently configured. In addition, when the `qos queue-config` command is executed, any previously configured `bandwidth-min output` settings are removed from the startup configuration. For the default GMB percentage allocations per number of queues, see “Default GMB percentage allocations per QoS queue configuration” (page 159).

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 10 Mbps on the HP 3500 switches and 1 Gbps on the other switches covered in this guide 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 10 Mbps on the HP 3500 switches and 1 Gbps on the other switches covered in this guide allow inbound jumbo frames of up to 9220 bytes.

Switch model	Minimum speed for jumbo traffic
3500	10 Mbps
All others in this guide	1 Gbps

Operating rules

- **Required port speed:** This feature allows inbound and outbound jumbo frames on ports operating at a minimum of 10 Mbps on the HP 3500 switches and 1 Gbps on the other switches.
- **Switch meshing:** If you enable jumbo traffic on a VLAN, all meshed ports on the switch are enabled to support jumbo traffic. (On a given meshed switch, every meshed port operating at 1 Gbps or higher becomes a member of every VLAN configured on the switch.)
- **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” \(page 150\)](#).

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.

Operating notes for jumbo traffic-handling

- HP Switch does not recommend configuring a voice VLAN to accept jumbo frames. Voice VLAN frames are typically small, and allowing a voice VLAN to accept jumbo frame traffic can degrade the voice transmission performance.
- You can configure the default, primary, and/or (if configured) the management VLAN to accept jumbo frames on all ports belonging to the VLAN.
- When the switch applies the default MTU (1522-bytes including 4 bytes for the VLAN tag) to a VLAN, all ports in the VLAN can receive incoming frames of up to 1522 bytes. When the switch applies the jumbo MTU (9220 bytes including 4 bytes for the VLAN tag) to a VLAN, all ports in that VLAN can receive incoming frames of up to 9220 bytes. A port receiving frames exceeding the applicable MTU drops such frames, causing the switch to generate an Event Log message and increment the "Giant Rx" counter (displayed by `show interfaces port-list` .)
- The switch allows flow control and jumbo frame capability to co-exist on a port.
- The default MTU is 1522 bytes (including 4 bytes for the VLAN tag.) The jumbo MTU is 9220 bytes (including 4 bytes for the VLAN tag.)
- When a port is not a member of any jumbo-enabled VLAN, it drops all jumbo traffic. If the port is receiving "excessive" inbound jumbo traffic, the port generates an Event Log message to notify you of this condition. This same condition also increments the switch's "Giant Rx" counter.
- If you do not want all ports in a given VLAN to accept jumbo frames, you can consider creating one or more jumbo VLANs with a membership comprising only the ports you want to receive

jumbo traffic. Because a port belonging to one jumbo-enabled VLAN can receive jumbo frames through any VLAN to which it belongs, this method enables you to include both jumbo-enabled and non-jumbo ports within the same VLAN.

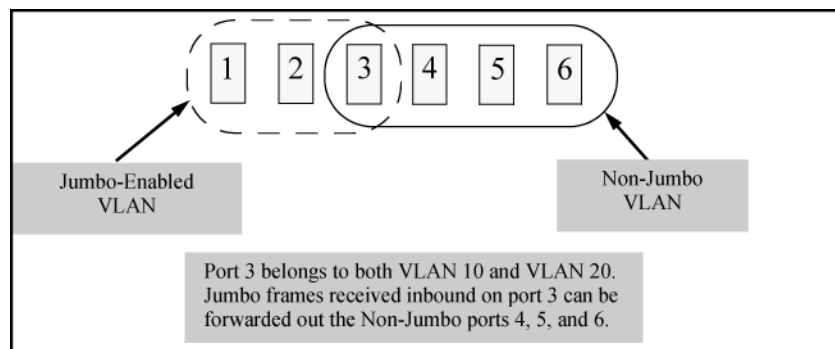
For example, suppose you want to allow inbound jumbo frames only on ports 6, 7, 12, and 13. However, these ports are spread across VLAN 100 and VLAN 200 and also share these VLANs with other ports you want excluded from jumbo traffic. A solution is to create a third VLAN with the sole purpose of enabling jumbo traffic on the desired ports, while leaving the other ports on the switch disabled for jumbo traffic. That is:

	VLAN 100	VLAN 200	VLAN 300
Ports	6-10	11-15	6, 7, 12, and 13
Jumbo-enabled?	No	No	Yes

If there are security concerns with grouping the ports as shown for VLAN 300, you can either use source-port filtering to block unwanted traffic paths or create separate jumbo VLANs, one for ports 6 and 7, and another for ports 12 and 13.

- Outbound jumbo traffic.** Any port operating at 1 Gbps or higher can transmit outbound jumbo frames through any VLAN, regardless of the jumbo configuration. The VLAN is not required to be jumbo-enabled, and the port is not required to belong to any other, jumbo-enabled VLANs. This can occur in situations where a non-jumbo VLAN includes some ports that do not belong to another, jumbo-enabled VLAN and some ports that do belong to another, jumbo-enabled VLAN. In this case, ports capable of receiving jumbo frames can forward them to the ports in the VLAN that do not have jumbo capability, as shown in [Figure 77](#).

Figure 77 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.

- Jumbo traffic in a switch mesh domain.** If a switch belongs to a meshed domain, but does not have any VLANs configured to support jumbo traffic, the meshed ports on that switch drop any jumbo frames they receive from other devices. In this regard, if a mesh domain includes any HP 1600M/2400M/2424M/4000M/8000M switches, along with the switches covered in this guide configured to support jumbo traffic, only the switches covered in this guide receive jumbo frames. The other switch models in the mesh will drop such frames. For more information on switch meshing, see the *Advanced Traffic Management Guide*.

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 10 Mbps on the HP 3500 switches or 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 10 Mbps for HP 3500 switches or 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.

6 Configuring for Network Management Applications

Command Summary

For overview information on using SNMPv3, see “About using SNMP tools to manage the switch” (page 210).

For information on enabling SNMPv3, see “SNMPv3 users” (page 212).

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.

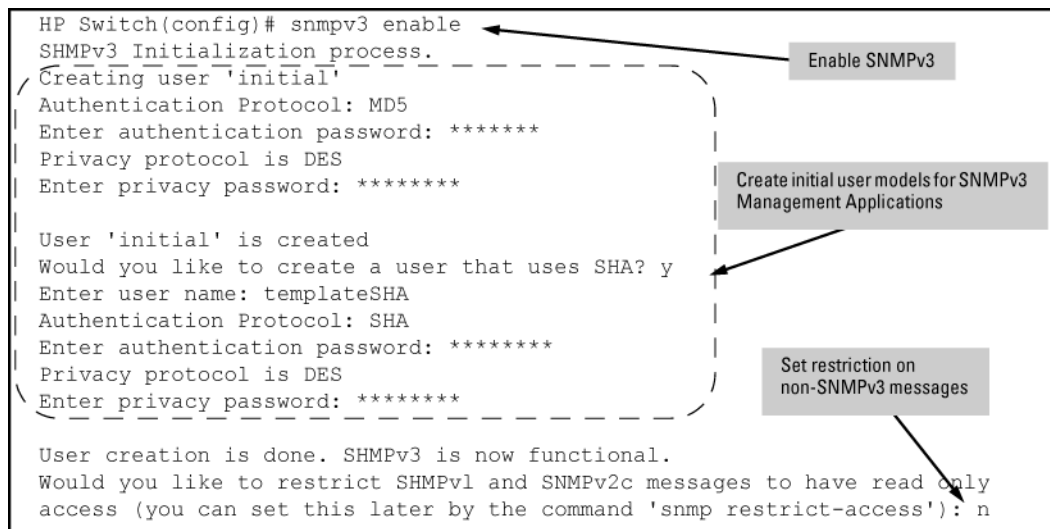
Example

Example 42 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
```



Configuring users (CLI)

For information on users, see “SNMPv3 users” (page 212)

Syntax

```
[no] snmpv3 user user_name
```

Adds or deletes a user entry for SNMPv3. Authorization and privacy are optional, but to use privacy, you must use authorization. When you delete a user, only the *user_name* is required.

```
[ auth md5 | sha auth_pass ]
```

With authorization, you can set either MD5 or SHA authentication. The authentication password *auth_pass* must be 6 to 32 characters and is mandatory when you configure authentication.

(Default: None)

```
[ priv des | aes priv_pass ]
```

With privacy, the switch supports DES (56-bit) and AES (128-bit) encryption. The privacy password `priv_pass` must be 6 to 32 characters and is mandatory when you configure privacy.

(Default: DES)

NOTE: Only AES 128-bit and DES 56-bit encryption are supported as privacy protocols. Other non-standard encryption algorithms, such as AES-172, AES-256, and 3-DES are not supported.

NOTE: For the 5400zl, 3800, and 8200zl switches, when the switch is in enhanced secure mode, commands that take a password as a parameter have the echo of the password typing replaced with asterisks. The input for the password is prompted for interactively. Additionally, the DES option is not available. For more information, see the *Access Security Guide*.

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
```

Viewing management stations configured to access the switch and view authentication and privacy protocols of each station (CLI)

Syntax

```
show snmpv3 user
```

Example

Example 43 Displays information about the management stations configured on VLAN 1 to access the switch

```
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 78 \(page 165\)](#). For more details on the MIBs access for a given group, see ["Group access levels" \(page 212\)](#).

Figure 78 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

Annotations in the image:

- "Add NetworkAdmin to operator noauth group" points to the first `snmpv3 group` command.
- "Add NetworkMgr to managerpriv group" points to the second `snmpv3 group` command.
- "Pre-assigned groups for access by Version 2c and version 1 management applications" points to the first four rows of the output table.

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 212).
<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.

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 213).

Syntax

```
[no] snmpv3 community
```

Maps or removes a mapping of a community name to a group access level. To remove a mapping you need to specify only the `index_name` parameter.

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

Example

Figure 79 (page 166) shows the assigning of the Operator community on MgrStation1 to the CommunityOperatorReadWrite group. Any other Operator has an access level of CommunityOperatorReadOnly.

Figure 79 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
    
```

Two Operator Access Levels

Index	Name	Community Name	Security Name
1		public	CommunityManagerReadWrite
2		Operator	CommunityOperatorReadOnly
3		Manager	CommunityManagerReadWrite
30		Operator	CommunityManagerReadWrite

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 214).)

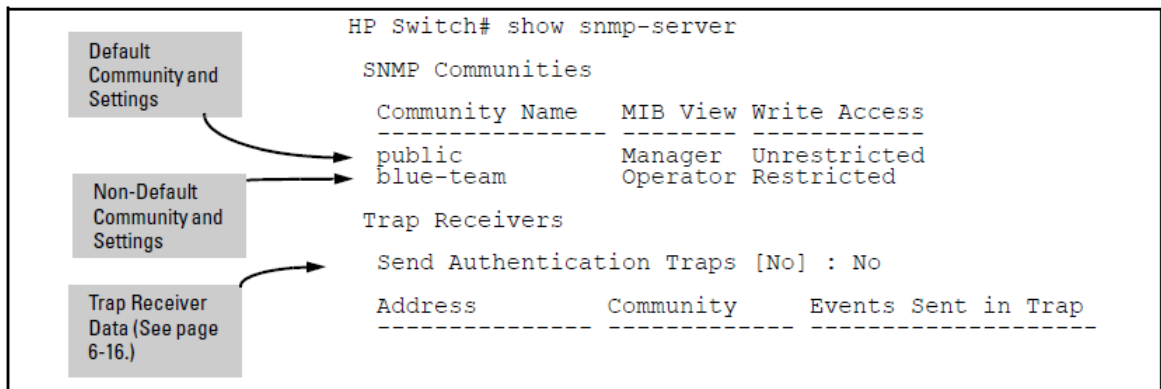
Syntax

```
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 80 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
```

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

- From the Main Menu, select:
 - Switch Configuration...**
 - SNMP Community Names**

Figure 81 The SNMP Communities screen (default values)

Note: This screen gives an overview of the SNMP communities that are currently configured. All fields in this screen are read-only.

```

----- CONSOLE - MANAGER MODE -----
Switch Configuration - SNMP Communities

Community Name  MIB View  Write Access
-----
public          Manager   Unrestricted

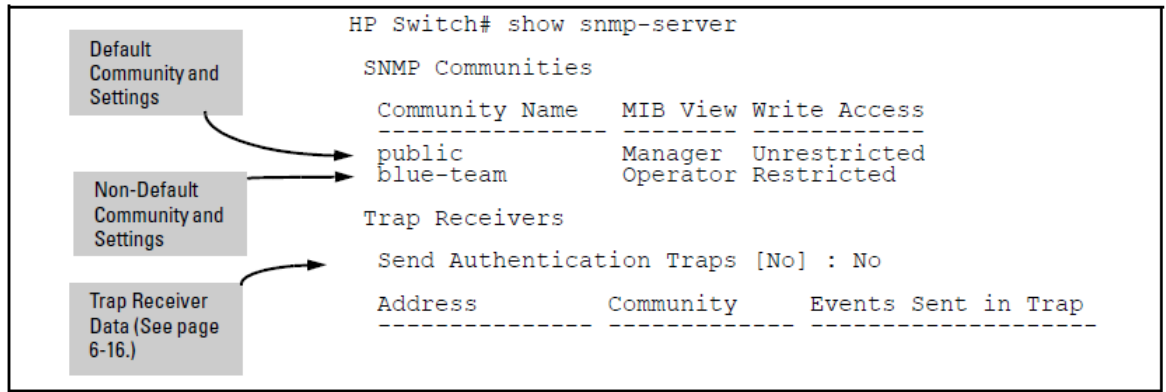
Actions->  Back  Add  Edit  Delete  Help
Return to previous screen.
Use up/down arrow keys to change record selection, left/right arrow ke:

```

Add and Edit options are used to modify the SNMP options. See Figure 8-2.

- Press **[A]** (for **Add**) to display the following screen:

Figure 82 The SNMP add or edit screen



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**.)

Configuring an SNMP trap receiver (CLI)

For information about configuring SNMP trap receivers, see “SNMP trap receivers” (page 215).

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 169).)</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 170).</p>

Table 17 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.

Table 17 Security levels for Event Log messages sent as traps (continued)

Security Level	Action
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
```

Enabling SNMPv2c informs (CLI)

For information about enabling SNMPv2c informs, see "SNMPv2c informs" (page 217).

Syntax

```
[no] snmp-server host [ ipv4-addr | ipv6-addr ]  
community name inform [ retries count ] [ timeout interval ]
```

Enables (or disables) the `inform` option for SNMPv2c on the switch and allows you to configure options for sending SNMP inform requests.

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

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 44 \(page 171\)](#) (note indication of `inform Notify Type` in bold below):

Example 44 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 ???).

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" \(page 212\)](#).) 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 \(CLI\)" \(page 165\)](#).)
 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 <i>notify_name</i></code>	Specifies the name of an SNMPv3 notification configuration.
<code>tagvalue <i>tag_name</i></code>	Specifies the name of a tag value used in other SNMPv3 commands, such as <code>snmpv3 targetaddress params taglist <i>tag_name</i></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 parms_name user user_name
```

Applies the configuration parameters and IP address of an SNMPv3 management station (from the `params params_name` value configured with the `snmpv3 targetaddress` command in Step 5) to a specified SNMPv3 user (from the `user user_name` value configured with the `snmpv3 user` command in Step 2.)

If you enter the `snmpv3 params user` command, you must also configure a security model (`sec-model`) and message processing algorithm (`msg-processing`.)

<pre>[sec-model [ver1 ver2c ver3]]</pre>	<p>Configures the security model used for SNMPv3 notification messages sent to the management station configured with the <code>snmpv3 targetaddress</code> command in Step 5.</p> <p>If you configure the security model as <code>ver3</code>, you must also configure the message processing value as <code>ver3</code>.</p>
<pre>[msg-processing ver1 ver2c ver3 [noaut auth priv]]</pre>	<p>Configures the algorithm used to process messages sent to the SNMPv3 target address.</p> <p>If you configure the message processing value as <code>ver3</code> and the security model as <code>ver3</code>, you must also configure a security services level (<code>noauth</code>, <code>auth</code>, or <code>priv</code>.)</p>

Example

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

Figure 83 Example of an SNMPv3 notification configuration

The diagram shows a sequence of configuration commands on a switch:

```
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.

Enabling or disabling notification/traps for network security failures and other security events (CLI)

For more information, see “Network security notifications” (page 217).

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:

<pre>arp-protect</pre>	<p>If ARP packets are received with an invalid source or destination MAC address, an invalid IP address, or an invalid IP-to-MAC binding.</p>
<pre>auth-server-fail</pre>	<p>If the connection with a RADIUS or TACACS+ authentication server fails.</p>

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 <i>port-list</i>	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.(Default: Disabled.) See “Enabling SNMP traps on Startup Configuration changes” (page 175)

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” (page 400).

Viewing the current configuration for network security notifications (CLI)

Enter the `show snmp-server traps` command, as shown in Figure 84 (page 174). Note that command output is a subset of the information displayed with the `show snmp-server` command in Figure 93 (page 180).

Figure 84 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-changeport-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.

Enabling SNMP traps on running configuration changes (CLI)

For more information, see [“Enabling SNMP traps on running configuration changes \(CLI\)” \(page 175\)](#).

Syntax

```
[no] snmp-server enable traps running-config-change [ transmission-interval  
0-4294967295 ]
```

Enables SNMP traps being sent when changes to the running configuration file are made.

(Default: Disabled)

`transmission-interval 0-2147483647` controls the egress rate for generating SNMP traps for the running configuration file. The value configured specifies the time interval in seconds that is allowed between the transmission of two consecutive traps.

None of the running configuration change events that occur within the specified interval generate SNMP traps, although they are logged in the Configuration Changes History Table.

A value of 0 (zero) means there is no limit; traps can be sent for every running configuration change event.

(Default: Zero)

Enabling SNMP traps on Startup Configuration changes

You can send a specific SNMP trap for any configuration change made in the switch's startup configuration file when the change is written to flash. Changes to the configuration file can occur when executing a CLI write command, executing an SNMP set command directly using SNMP, or when using the WebAgent

NOTE: A log message is always generated when a startup configuration change occurs. An example log entry is:

```
I 07/06/10 18:21:39 02617 mgr: Startup configuration changed by SNMP.  
New seq. number 8
```

The corresponding trap message is sent if the `snmp-server enable traps startupconfig- change` command is configured.

Syntax

```
[no]snmp-server enable traps startup-config-change
```

Enables notification of a change to the startup configuration. The change event is logged. Default: Disabled

An example of configuring the command from the CLI is shown in Figure 6-16. The number that displays when show config is executed is global for the switch and represents the startup configuration sequence number.

Figure 85 Enabling notification of changes to the Startup Configuration file

```
Switch(config)# snmp-server enable traps startup-config-change

Switch(config)# show config

Startup configuration: 16 ← The number "16" is global for the switch and represents the startup
                           configuration sequence number.

; J8697A Configuration Editor; Created on release #K.14.54

hostname "Switch"
module 1 type J8702A
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24, B1-B10
  ip address dhcp-bootp
  exit
snmp-server community "public" unrestricted
```

Figure 6-17 displays an example of the fields in the trap when a change is made via SNMP (station ip=0xAC161251 (172.22.18.81), no username is set, and the new sequence number is 16.)

Figure 86 Fields when the SNMP trap is set

```

+ Internet Protocol, Src: 172.22.18.57 (172.22.18.57), Dst: 172.22.18.81 (172.22.18.81)
+ User Datagram Protocol, Src Port: snmp (161), Dst Port: snmptrap (162)
+ Simple Network Management Protocol
  version: version-1 (0)
  community: public
  data: trap (4)
    trap
      enterprise: 1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1 (SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1)
      agent-addr: 172.22.18.57 (172.22.18.57)
      generic-trap: enterpriseSpecific (6)
      specific-trap: 6
      time-stamp: 65437
      variable-bindings: 6 items
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.9 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.9): 16
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.0.1 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.0.1): 2
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.0.2 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.0.2): 1
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.0.3 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.0.3): AC161251
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.0.4 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.0.4): <MISSING>
        + SNMPv2-SMI::enterprises.11.2.14.11.5.1.7.1.29.1.0.5 (1.3.6.1.4.1.11.2.14.11.5.1.7.1.29.1.0.5): 1

```

Viewing configuration file change information (CLI)

Syntax

```
show running-config [ changes-history [ 1-32 ] ][ detail ]
```

Displays the history up to 32 events for changes made to the running-configuration file, as shown in Figure 87 (page 177) and Figure 88 (page 177). The changes are displayed in descending order, the most recent change at the top of the list. You can specify from 1 to 32 entries for display.

The detail option displays a more detailed amount of information for the configuration changes. Figure 89 (page 177) and Figure 90 (page 177) display detailed information for configuration changes history.

Examples

Figure 87 Example of output for running configuration changes history for all ports

```
HP Switch(config)# show running-config changes-history
Running Config Last Changed   : 02/19/10 16:30:09
Number of Changes Since Reboot : 150086

Event ID   Config
-----
150086    CLI      02/19/10 16:30:09
150085    SNMP     02/03/10 14:50:12
150084    SNMP     02/03/10 14:50:12
150083    SNMP     02/03/10 14:45:59
150082    SNMP     02/03/10 14:57:15
150081    SNMP     02/03/10 13:11:00
150080    SNMP     02/03/10 13:11:00
150079    CLI      01/18/10 09:09:17
```

Figure 88 Example of output for running configuration changes history

```
HP Switch(config)# show running-config changes-history 6
Running Config Last Changed   : 08/04/10 16:35:31
Number of Changes since Reboot : 120

Event ID   Config
-----
120        CLI      08/04/10 16:35:31
119        CLI      08/04/10 16:34:01
118        SNMP     08/04/10 15:32:22
117        WEBUI    08/03/10 12:55:21
116        MENU     07/01/10 01:45:26
115        CLI      06/23/10 11:34:23
```

Figure 89 Example of detailed output for running configuration changes history

```
HP Switch(config)# show running-config changes-history 3 detail

Event ID   : 120
User       : switch_admin
Remote IP Address : 10.11.12.4
Config Method : CLI
Date       : 08/04/10
Time       : 16:35:31

Event ID   : 119
User       : switch_admin
Remote IP Address : 10.11.12.4
Config Method : CLI
Date       : 08/04/10
Time       : 16:34:01

Event ID   : 118
User       : switch_admin
Remote IP Address : 10.11.12.4
Config Method : SNMP
Date       : 08/04/10
Time       : 15:32:22
```

Figure 90 Example of output for running config changes history with detail

```
HP Switch(config)# show running-config changes-history detail
Running Config Last Changed: 01/01/90 00:35:44
Number of changes since last boot : 6

Event ID   : 6
User       :
Remote IP Address :
Config Method : CLI
Date       : 01/01/90
Time       : 00:35:44

Event ID   : 5
User       :
Remote IP Address :
Config Method : CLI
Date       : 01/01/90
Time       : 00:35:39

Event ID   : 4
User       :
Remote IP Address :
Config Method : CLI
Date       : 01/01/90
Time       : 00:35:33

Event ID   : 3
User       :
Remote IP Address :
Config Method : CLI
Date       : 01/01/90
Time       : 00:35:27
```

Figure 91 (page 178) displays the current status (enabled/disabled) of the SNMP trap type for running-configuration changes.

Figure 91 Example of SNMP trap configuration status information

```

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
Running Configuration Changes : Enabled ← SNMP trap status for running-config changes
is enabled.

Address      Community      Events  Type  Retry  Timeout
-----
173.33.25.201  public          None   trap  3      15

Excluded MIBs
    
```

Configuring the source IP address for SNMP notifications (CLI)

For more information, see “Source IP address for SNMP notifications” (page 218).

Syntax

```
[no] snmp-server response-source [ dst-ip-of-request [ ipv4-addr |
ipv6-addr ] | loopback0-7 ]
```

Specifies the source IP address of the SNMP response PDU. The default SNMP response PDU uses the IP address of the active interface from which the SNMP response was sent as the source IP address.

The `no` form of the command resets the switch to the default behavior (compliant with rfc-1517.) (Default: Interface IP address)

<code>dst-ip-of-request</code>	Destination IP address of the SNMP request PDU that is used as the source IP address in an SNMP response PDU.
<code>[ipv4-addr ipv6-addr]</code>	User-defined interface IP address that is used as the source IP address in an SNMP response PDU. Both IPv4 and IPv6 addresses are supported.
<code>loopback 0-7</code>	IP address configured for the specified loopback interface that is used as the source IP address in an SNMP response PDU. If multiple loopback IP addresses are configured, the lowest alphanumeric address is used.

Example

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 | loopback0-7 ]
```

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>	IP 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 Figure 92 (page 179).

Figure 92 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

Trap Pdu Source-IP Information
Selection Policy : Configured IP
Ip Address       : 10.10.10.10
```

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 93 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 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 of information displayed for show mac-count-notify traps command

```
HP Switch (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 of 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 of 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

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					

Configuring the MAC address table change option

When enabled, this feature allows the generation of SNMP traps for each MAC address table change. Notifications can be generated for each device that connects to a port and for devices that are connected through another device (daisy-chained.)

The snmp-server enable traps mac-notify command globally enables the generation of SNMP trap notifications upon MAC address table changes.

Syntax

```
[no]snmp-server enable traps mac-notify [mac-move |  
trap-interval <0- 120>]
```

Globally enables or disables generation of SNMP trap notifications.

trap-interval	The time interval (in seconds) that trap notifications are sent. A value of zero disables the interval and traps are sent as events occur. If the switch is busy, notifications can be sent prior to the configured interval. Notifications may be dropped in extreme instances and a system warning is logged. The range is 0-120 seconds. Default: 30 seconds.
mac-move	Configures the switch to capture data for MAC addresses that are moved from one port to another port. The <code>snmp-server enable traps mac-notify</code> command must have been enabled in order for this information to be sent as an SNMP notification.

Example of trap-interval option

```
HP Switch (config)# snmp-server enable traps mac-notify  
trap-interval 60
```

Example of mac-move option

```
HP Switch (config)# snmp-server enable traps mac-notify mac-move
```

Additional mac-notify options for per-port Mac changes

Use the following command to configure SNMP traps for learned or removed MAC addresses on a per-port basis.

NOTE: The switch will capture learned or removed events on the selected ports, but will not send an SNMP trap unless `mac-notify` has been enabled with the `snmp-server enable traps mac-notify` command.

Syntax

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

When this command is executed without the **learned** or **removed** option, it enables or disables the capture of both learned and removed MAC address table changes for the selected ports in `<port-list>`

<port-list>	Configures MAC address table changes capture on the specified ports. Use all to capture changes for all ports on the switch.
learned	Enables the capture of learned MAC address table changes on the selected ports.
removed	Enables the capture of removed MAC address table changes table on the selected ports.

Example of configuring traps on a per-port basis for learned MAC addresses

```
HP Switch(config)# mac-notify traps 5-6 learned  
HP Switch(config)# show mac-notify traps 5-6  
Mac Notify Trap Information  
Mac-notify Enabled : Yes  
Mac-move Enabled : Yes  
Trap-interval : 60  
Port    MAC Addresses      trap learned/removed  
-----  
5       Learned  
6       Learned
```

Example of configuring traps on a port-bases for removed MAC addresses

```
HP Switch(config)# mac-notify traps 3-4 removed
Switch(config)# show mac-notify traps
Mac Notify Trap Information
Mac-notify Enabled : Yes
Mac-move Enabled : Yes
Trap-interval : 60
Port      MAC Addresses      trap learned/removed
-----
1         None
2         None
3         Removed
4         Removed
```

Configuring the mac-notify option at the interface context level

You can also execute the `mac-notify traps` command from the interface context.

Example of the interface context for MAC-notify traps command

```
HP Switch(config)# int 11
HP Switch(int-11)# mac-notify traps learned
```

Viewing mac-notify traps configuration information

Use the `show mac-notify traps` command to display information about SNMP trap configuration for MAC Address Table changes.

Syntax

```
show mac-notify traps [port-list]
```

Example of information for SNMP trap configuration

Displays SNMP trap information for all ports, or each port in the port-list.

```
HP Switch(config)# show mac-notify traps
Mac Notify Trap Information
Mac-notify Enabled : Yes
Mac-move Enabled : Yes
Trap-interval : 60
Port      MAC Addresses      trap learned/removed
-----
1         None
2         None
3         Removed
4         Removed
5         Learned
6         Learned
```

The configured `mac-notify` commands are displayed in the `show running-configuration` output.

Example of running config file with mac-notify parameters configured

```
HP Switch(config)# show running-config
Running configuration:
; J9087A Configuration Editor; Created on release #R.11.XX
hostname "Switch"
snmp-server community "public" Unrestricted
snmp-server host 15.255.133.236 "public"
snmp-server host 15.255.133.222 "public"
snmp-server host 15.255.133.70 "public"
snmp-server host 15.255.134.235 "public"
```

```

vlan 1
  name "DEFAULT_VLAN"
  untagged 1-28
  ip address dhcp-bootp
  exit
snmp-server enable traps mac-notify mac-move
snmp-server enable traps mac-notify trap-interval 60
snmp-server enable traps mac-notify
mac-notify traps 5-6 learned
mac-notify traps 3-4 removed

```

Configuring listening mode (CLI)

For more information, See [“Configuring listening mode \(CLI\)” \(page 184\)](#).

Syntax

```
snmp-server [ listen [ oobm | data | both ] ]
```

Enables or disables inbound SNMP access on a switch.

Use the `no` version of the command to disable inbound SNMP access.

The `listen` parameter is available only on switches that have a separate out-of-band management port. Values for this parameter are:

oobm	Inbound SNMP access is enabled only on the out-of-band management port.
data	Inbound SNMP access is enabled only on the data ports.
both	Inbound SNMP access is enabled on both the out-of-band management port and on the data ports. This is the default value.

For more information on out-of-band management, see [“Network Out-of-Band Management \(OOBM\) for the 6600 Switch” \(page 484\)](#).

The `listen` parameter is not available on switches that do not have a separate out-of-band management port.

Configuring sFlow (CLI)

The following sFlow commands allow you to configure sFlow instances via the CLI. For more information, see [“Advanced management: RMON” \(page 219\)](#).

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.

Syntax

```
[no] sflow receiver-instance destination <ipv4 | ipv6> [udp-port-num] oobm
```

This command provides a configurable option for sending sFlow packets to a destination through the OOBM port on the switch. The sFlow collector collects sample packets through the OOBM port, allowing the monitoring of network traffic. Both IPv4 and IPv6 addresses are supported.

The command 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>**

oobm: Use the OOBM port to reach the specified sFlow receiver.

Example 45 Example

Figure 94 sFlow Destination is OOBM port

```
HP Switch(config)# sflow 1 destination 192.168.2.3 6000 oobm
```

Figure 95 Output showing OOBM Support Enabled

```
HP Switch# show sflow 1 destination
Destination Instance      : 1
sflow                    : Enabled
Datagrams Sent           : 0
Destination Address       : 192.168.2.3
Receiver Port             : 6043
Owner                     : Administrator, CLI-Owned, Instance 1
Timeout (seconds)        : 2147479533
Max Datagram Size        : 1400
Datagram Version Support : 5
OOBM Support              : Enabled
```

Figure 96 Output of the running-config File showing the sFlow Destination is the OOBM Port

```
HP Switch# show running-config
Running configuration:
; J9091A Configuration Editor; Created on release #K.15.06.0006
; Ver #01:0d:0c

hostname "HP Switch"
module 1 type J8702A
module 7 type J8708A
module 12 type J8702A
sflow 1 destination 192.168.2.3 oobm
vlan 1
  name "DEFAULT VLAN"
  untagged A1-A24,G1-G4,L1-L24
  ip address dhcp-bootp
  exit
snmp-server community "public" unrestricted
```

Viewing sFlow Configuration and Status (CLI)

The following sFlow commands allow you to display sFlow configuration and status via the CLI. Figure 98 (page 187) 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.

Figure 97 Example of 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 [Figure 98](#) (page 187).

Figure 98 Example of 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)       9999530
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 99](#) (page 187). You can specify a list or range of ports for which to view sampling information.

Figure 99 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 | Enabled | Rate | Header | Dropped | Samples | Polling | Enabled | Interval
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
A1 | Yes (2) | Yes | 128 | 128 | 1234567890 | 0 | Yes (1) | Yes | 60
A2 | --- | --- | --- | --- | 0 | 0 | No | No | 30
A3 | No (1) | No | 100 | 100 | 898703 | 0 | No | No | 0
A4 | Yes (3) | Yes | 50 | 128 | 0 | 0 | No (3) | No | 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.

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)” on page 6-74.

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 191).

Example of viewing the general LLDP configuration

show lldp config produces the following display when the switch is in the default 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)

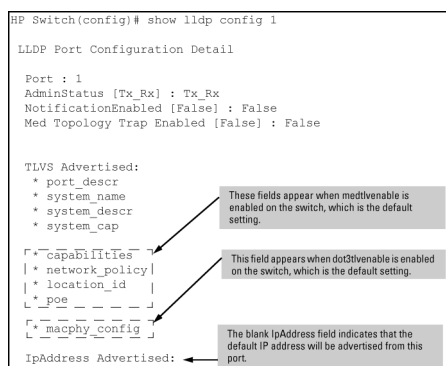
Syntax

```
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 “SNMP notification support” (page 226). For information on the other configurable settings displayed by this command, see “Configuring per-port transmit and receive modes (CLI)” (page 191).

Figure 100 Per-port configuration display



Enabling or disabling LLDP operation on the switch (CLI)

For more information, see [“LLDP operation on the switch” \(page 226\)](#).

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

To disable LLDP on the switch:

```
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.

Changing the time-to-live for transmitted advertisements (CLI)

For more information, see [“Time-to-Live for transmitted advertisements” \(page 226\)](#).

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-interval` to 2, which would result in a Time-to-Live of 30 seconds.

```
HP Switch(config)# lldp holdtime-multiplier 2
```

Changing the delay interval between advertisements generated by value or status changes to the LLDP MIB (CLI)

For more information, see “Delay interval between advertisements generated by value or status changes to the LLDP MIB” (page 226).

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.

NOTE: For the 5400zl, 3800, and 8200zl switches, when the switch is in enhanced secure mode, the following prompt appears before the sensitive information for the `setmib` command is displayed:

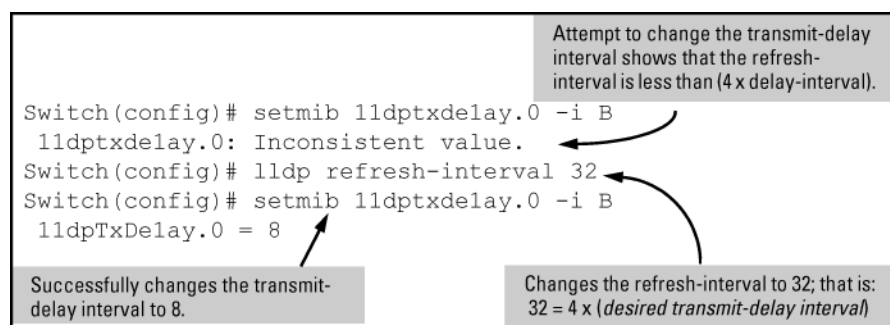
The `setmib` command should not be used in enhanced secure mode.

For more information, see the *Access Security Guide*.

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 101 (page 190).)

Figure 101 Changing the transmit-delay interval



Changing the reinitialization delay interval (CLI)

For more information, see “Reinitialization delay interval” (page 226).

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
```

Enabling LLDP data change notification for SNMP trap receivers (CLI)

Syntax

```
[no] lldp enable-notification port-list
```

Enables or disables each port in *port-list* for sending notification to configured SNMP trap receivers if an LLDP data change is detected in an advertisement received on the port from an LLDP neighbor.

(Default: Disabled)

For information on configuring trap receivers in the switch, see [“SNMP notifications”](#) (page 214).

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

For more information, see [“Changing the minimum interval for successive data change notifications for the same neighbor”](#) (page 226).

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.

txonly	Configures the specified ports to transmit LLDP packets, but block inbound LLDP packets from neighbor devices.
rxonly	Configures the specified ports to receive LLDP packets from neighbors, but block outbound packets to neighbors.
tx_rx	Configures the specified ports to both transmit and receive LLDP packets. (This is the default setting.)
disable	Disables LLDP packet transmit and receive on the specified ports.

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” \(page 226\)](#).

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
```

Configuring support for port speed and duplex advertisements (CLI)

For more information, see [“Support for port speed and duplex advertisements” \(page 227\)](#).

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.

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” \(page 227\)](#).

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 102 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 [Figure 104 \(page 194\)](#) through [Figure 105 \(page 195\)](#).

Figure 103 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 ← The VLAN ID TLV is being advertised.

IpAddress Advertised:
:
:
```

Figure 104 Example of local device LLDP information

```
HP Switch(config)# show lldp info local-device a1

LLDP Local Port Information Detail

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

Port VLAN ID : 1 ← The information that LLDP used in its advertisement.
```

Figure 105 Example of 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
```

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 \(media-endpoint-discovery\)” \(page 228\)](#).

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 [“About determining the switch port number included in topology change notification traps” \(page 234\)](#).)
- The LLDP-MED class of the device detected on the port ([“LLDP-MED endpoint device classes” \(page 229\)](#).)

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” \(page 214\)](#). 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.

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” \(page 230\)](#).

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” \(page 197\)](#).)
- 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.

<code>capabilities</code>	<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>
<code>network-policy</code>	<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 capability TLV is already enabled.</p> <p>For more information, see “Network policy advertisements” (page 230).</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 197).</p> <p>(Default: Enabled)</p> <p>NOTE: When disabled, this TLV cannot be enabled unless the capability TLV is already enabled.</p>
poe	<p>This TLV enables the switch port to advertise its current PoE state and to read the PoE requirements advertised by the LLDP-MED endpoint device connected to the port.</p> <p>(Default: Enabled)</p> <p>NOTE: When disabled, this TLV cannot be enabled unless the capability TLV is already enabled.</p> <p>For more on this topic, see “PoE advertisements” (page 231)</p>

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 [“PoE advertisements” \(page 231\)](#).

To display the current PoE configuration on the switch, use the following commands:

```
show power brief port-list
show power port-list
```

For more information on PoE configuration and operation, see [“Power Over Ethernet \(PoE/PoE+\) Operation” \(page 80\)](#).

Configuring location data for LLDP-MED devices

For more information, see [“Location data for LLDP-MED devices” \(page 231\)](#).

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> command. (For a complete list of country codes, visit www.iso.org .)
WHAT	A single-digit number specifying the type of device to which the location data applies: 0 : Location of DHCP server 1 : Location of switch 2 : Location of LLDP-MED endpoint (recommended application) This field is required in a <code>civic-addr</code> command.

<p>Type/Value Pairs [CA-TYPE/CA-VALUE]</p>	<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 (CA-TYPE), and the second value in a pair is expected to be the corresponding civic address data (CA-VALUE.)</p> <p>For example, if the CA-TYPE 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 CA-TYPE.</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>CA-TYPE: This is the first entry in a type/value pair and is a number defining the type of data contained in the second entry in the type/value pair (CA-VALUE.) Some examples of CA-TYPE specifiers include:</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 CA-TYPE specifiers, see Table 6-5 (page 231).</p> <p>CA-VALUE: This is the second entry in a type/value pair and is an alphanumeric string containing the location information corresponding to the immediately preceding CA-TYPE entry.</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 CA-TYPE 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>
<p><code>elin-addr emergency-number</code></p>	<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>

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.

Example

In the default configuration, the switch information currently available for outbound LLDP advertisements appears similar to the display in [Figure 106 \(page 200\)](#).

Figure 106 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 00
System Name  : HP Switch1
System Description : HP J9091A Switch 3500y1, revision K.15.06...
System Capabilities Supported:bridge
System Capabilities Enabled:bridge

Management Address :
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
```

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). For more on this topic, refer to "Remote Management Address" on page 6-52.

Figure 107 Example of 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
```

Viewing the current port speed and duplex configuration on a switch port

For more information, see [“About displaying the current port speed and duplex configuration on a switch port” \(page 233\)](#).

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 6-4 \(page 223\)](#).

Examples

Example 46 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 3500yl
2	00 11 85 cf 66 60	8	8	HP Switch 3500yl

Figure 108 An LLLDP-MED listing of an advertisement received from an LLDP-MED (VoIP telephone) source

```
HP Switch(config)# show lldp info remote-device 1

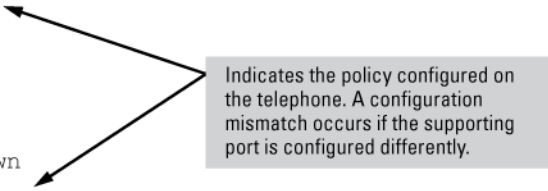
LLDP Remote Device Information Detail

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 K.15.06.0000x
PortDescr       : LAN Port

System Capabilities Supported : bridge, telephone
System Capabilities Enabled   : bridge, telephone

Remote Management Address

MED Information Detail
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
```



Indicates the policy configured on the telephone. A configuration mismatch occurs if the supporting port is configured differently.

Viewing LLDP statistics

For more information, see [“Displaying LLDP statistics” \(page 233\)](#).

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	<p>The number of neighbor deletions from the MIB for AgeOut Count and forced drops for all ports.</p> <p>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.</p> <p>The device receiving the shutdown packet deletes all information about the neighbor received on the applicable inbound port and increments the counter.</p>
Neighbor Entries Dropped Count	<p>The number of valid LLDP neighbors the switch detected, but could not add.</p> <p>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 233).</p>
Neighbor Entries AgeOut Count	The number of LLDP neighbors dropped on all ports because of Time-to-Live expiring.

Per-Port LLDP Counters:

NumFramesRecv	<p>The total number of valid, inbound LLDP advertisements received from any neighbors on <code>port-list</code> .</p> <p>Where multiple neighbors are connected to a port through a hub, this value is the total number of LLDP advertisements received from all sources.</p>
NumFramesSent	The total number of LLDP advertisements sent from <code>port-list</code> .
NumFramesDiscarded	<p>The total number of inbound LLDP advertisements discarded by <code>port-list</code> .</p> <p>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 233). This can also be an indication of advertisement formatting problems in the neighbor device.</p>
Frames Invalid	<p>The total number of invalid LLDP advertisements received on the port.</p> <p>An invalid advertisement can be caused by header formatting problems in the neighbor device.</p>
TLVs Unrecognized	<p>The total number of LLDP TLVs received on a port with a type value in the reserved range.</p> <p>This can be caused by a basic management TLV from a later LLDP version than the one currently running on the switch.</p>

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 47 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	NumFramesRecvd	NumFramesSent	NumFramesDiscarded
A1	97317	97843	0
A2	21	12	0
A3	0	0	0
A4	446	252	0
A5	0	0	0
A6	0	0	0
A7	0	0	0
A8	0	0	0

Example 48 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
```

Configuring CDP mode

To set the CDP mode to pass-through or receive only, enter this command.

Syntax

```
[no]cdp moden [pass-through | rxonly]
```

Sets the selected mode of CDP processing.

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

Example 49 Show CDP with the default CDP configuration

This example shows the 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-num]]	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

Example 50 CDP neighbors table listing

This example displays the CDP devices that the switch has detected by receiving their CDP packets.

```
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 LLDP-MED)
- 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.

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

tx_rx

Enable transmit and receive mode.

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

Example

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

Example

Show CDP output when CDP Run is disabled.

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

Example

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
```

Example

show cdp output when cdp run and cdp mode rxonly are enabled. When CDP mode is not pre-standard voice, the admin-status column is not displayed.

```
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

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.

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
```

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-CCCCCC (CDP)
- 0180c2- 00000e (LLDP)
- 0180c2-000003 (EAPOL)

Example

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

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
```

About 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 the *Advanced Traffic Management Guide*.

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*. (The latest version of this guide is available on the HP Networking website.) For information on the Management VLAN feature, see the *Advanced Traffic Management Guide*.

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 213))
- 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/support

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 213).)
2. Configure the appropriate trap receivers. (See “SNMP notifications” (page 214).)

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*.)

- △ 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. If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address.

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 163).)
2. Configure the appropriate SNMP users (see “SNMPv3 users” (page 212).)
3. Configure the appropriate SNMP communities (see “SNMPv3 communities” (page 213).)
4. Configure the appropriate trap receivers (see “SNMP notifications” (page 214).)

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*.)

SNMP version 3 (SNMPv3) adds some new commands to the CLI for configuring SNMPv3 functions. To enable SNMPv3 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.

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 “About adding users” (page 212).)
2. Assign users to Security Groups based on their security model with the `snmpv3 group` command (see “Assigning users to groups (CLI)” (page 165).)

- △ **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.

About 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 109 (page 212).

Figure 109 Adding SNMPv3 users and displaying SNMPv3 configuration

```
HP Switch(config)# snmpv3 user NetworkAdmin
HP Switch(config)# snmpv3 user NetworkMgr auth md5 authpass priv privpass
HP Switch(config)# show snmpv3 user
```

Status and Counters - SNMP v3 Global Configuration Information

User Name	Auth. Protocol	Privacy Protocol
initial	MD5	CFB AES-128
NetworkAdmin	MD5	CBC-DES

Group access levels

The switch supports eight predefined group access levels, shown in Table 6-3 (page 213). 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 18 Predefined group access levels

Group name	Group access type	Group read view	Group write view
managerpriv	Ver3 Must have Authentication and Privacy	ManagerReadView	ManagerWriteView
managerauth	Ver3 Must have Authentication	ManagerReadView	ManagerWriteView
operatorauth	Ver3 Must have Authentication	OperatorReadView	DiscoveryView
operatornoauth	Ver3 No Authentication	OperatorReadView	DiscoveryView
commanagerrw	Ver2c or Ver1	ManagerReadView	ManagerWriteView
commanagerr	Ver2c or Ver1	ManagerReadView	DiscoveryView
comoperatorrw	Ver2c or Ver1	OperatorReadView	OperatorReadView
comoperatorr	Ver2c or Ver1	OperatorReadView	DiscoveryView

Each view allows you to view or modify a different set of MIBs:

- **Manager Read View** – access to all managed objects
- **Manager Write View** – access to all managed objects except the following:
 - vacmContextTable
 - vacmAccessTable
 - vacmViewTreeFamilyTable
- **OperatorReadView** – no access to the following:
 - icfSecurityMIB
 - hpSwitchIpTftpMode
 - vacmContextTable
 - vacmAccessTable
 - vacmViewTreeFamilyTable
 - usmUserTable
 - snmpCommunityTable
- **Discovery View** – Access limited to samplingProbe MIB.

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 166).)

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."

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 215)
- "Configuring an SNMP trap receiver (CLI)" (page 169)
- "Enabling SNMPv2c informs (CLI)" (page 170)

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 171).

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 175).

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 \(CLI\)” \(page 178\)](#)
 - [“Viewing SNMP notification configuration \(CLI\)” \(page 179\)](#)

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.

SNMP trap when MAC address table changes

An SNMP trap is generated when a laptop/PC is removed from the back of an IP phone and the laptop/PC MAC address ages out of the MAC table for the HP Switch 2920 and HP 5400 series switch.

The `mac-notify` trap feature globally enables the generation of SNMP trap notifications on MAC address table changes (learns/moves/removes/ages.)

The following command enables trap for aged MAC addresses:

Syntax

```
HP Switch(config)# [no] mac-notify traps [port-list] aged
```

Example

For port 1 the command is:

Syntax

```
HP Switch(config)# mac-notify traps 1 aged
```

Show command

Use the following show command to display the different mac-notify traps configured on an interface:

Syntax

```
HP Switch # show mac-notify traps
```

Displays the following information:

```
Mac Notify Trap Information
Mac-notify Enabled : No
Mac-move Enabled : No
Trap-interval : 30
Port    MAC Addresses trap learned/removed/aged
-----
1      Learned, Removed & Aged
2      Removed & Aged
3      Learned & Aged
4      Learned & Removed
5      Aged
6      Learned
7      Removed
```

Example

For port 1 the command would be as follows

```
HP Switch # show mac-notify traps 1
```

Displays the following information:

```
1 Aged
```

SNMP trap when power supply is inserted or removed

SNMP traps generate while inserting or removing a powered up Power Supply Unit (PSU) without pulling out the power cable and also when removing a powered down PSU from the HP Switch 5406 Series. RMON log events are used to generate SNMP traps for PSU insertion and removal in both powered up and powered down states.

Log Event

```
Chassis: Power Supply 1 inserted
Chassis: Power Supply 1 removed while powered
Chassis: Power Supply 2 removed while not powered
```

Example

Power Supply inserted while powered off:

```
W 09/13/13 09:10:18 03834 chassis: AM1: Power Supply 1 inserted
W 09/13/13 09:10:19 00071 chassis: AM1: Power Supply failure: Supply: 1, Failures: 4
```

Example

Power Supply inserted while powered on:


```
W 09/13/13 09:06:20 03834 chassis: AM1: Power Supply 1 inserted
W 09/13/13 09:06:21 00071 chassis: AM1: Power Supply OK: Supply: 1, Failures: 2
```

Example

Power Supply removed while powered off:

```
W 09/13/13 09:08:57 03835 chassis: AM1: Power Supply 1 removed while not powered
W 09/13/13 09:08:57 00071 chassis: AM1: Power Supply failure: Supply: 1, Failures: 3
```

Example

Power Supply inserted while powered on:

```
W 09/13/13 09:03:36 03835 chassis: AM1: Power Supply 1 removed while powered
W 09/13/13 09:03:36 00071 chassis: AM1: Power Supply failure: Supply: 1, Failures: 2
```

About 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.

SNMPv2c informs

On a switch enabled for SNMPv2c, you can use the `snmp-server host inform` command ([“Enabling SNMPv2c informs \(CLI\)” \(page 170\)](#)) 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.

Network security notifications

By default, a switch is enabled to send the SNMP notifications listed in [“Supported Notifications” \(page 214\)](#) 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 \(CLI\)” \(page 169\)](#)
- [“Configuring SNMPv3 notifications \(CLI\)” \(page 171\)](#)

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

SNMP traps on running configuration changes

You can send a specific SNMP trap for any configuration change made in the switch's running configuration file. The trap will be generated for changes made from any of these interfaces:

- CLI
- Menu
- SNMP (remote SNMP set requests.)

The SNMP trap contains the following information.

Information	Description
Event ID	An assigned number that identifies a specific running configuration change event.
Method	Method by which the change was made—CLI, Menu, or remote SNMP. For configuration changes triggered by internal events, the term "Internal-Event" is used as the source of the change.
IP Address Type	Indicates the source address type of the network agent that made a change. This is set to an address type of "unknown" when not applicable.
IP address	IP address of the remote system from which a user accessed the switch. If not applicable, this is an empty string and nothing is displayed, for example, if access is through a management console port.
User Name	User name of the person who made the change. Null if not applicable.
Date and Time	Date and time the change was made.

The SNMP trap alerts any interested parties that someone has changed the switch's configuration and provides information about the source for that change. It does not specify what has been changed.

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 178) and `snmp-server trap-source` (page 178) commands.

Listening mode

For switches that have a separate out-of-band management port, you can specify whether a configured SNMP server listens for SNMP queries over the OOBM interface, the data interface, or both. By default, the switch listens over both interfaces.

This option is not available for switches that do not have a separate OOBM port. For more information on network OOBM, see “[Network Out-of-Band Management \(OOBM\) for the 6600 Switch](#)” (page 484).

The listening mode is set with parameters to the `snmp-server` command.

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.

The CLI supports the configuration of RMON alarm threshold settings. The settings can be saved in the configuration file.

Syntax

```
[no] rmon alarm entry number alarm-variable sampling-interval absolute  
| delta rising-threshold threshold-value1 falling-threshold2  
threshold-value2 owner string
```

This command configures RMON sampling periods and threshold parameters. The **no** option removes the alarm entry.

entry number <1-65535>: An alarm number that uniquely identifies the alarm threshold entry.

alarm-variable <object-string>: Object identifier of the particular variable to be sampled. Variables must be of type Integer in order to be sampled.

sampling-interval <5-65535>: Time interval in seconds over which data is sampled and compared with the rising-threshold and the falling-threshold.

absolute: The value of the selected variable is compared directly with the thresholds at the end of the sampling interval.

NOTE: If the **absolute** option is used for alarm variables of counter-type, an RMON trap is generated only once, when the threshold limit is reached. The RMON trap is never generated again.

It is recommended that you use the **delta** option instead when using a counter-type alarm variable.

delta: The value of the selected variable at the last sample is subtracted from the current value, and the difference is compared with the thresholds.

rising-threshold <threshold-value 1>: An integer value for the upper threshold for the sampled statistic. A single event is generated when the current sampled value of the specified statistic

becomes greater than or equal to this threshold, and if the value at the last sampling intervals was less than this threshold.

NOTE: The value of the rising-threshold must be greater than the value of the falling-threshold.

falling-threshold <threshold-value2>: An integer value for the lower threshold for the sampled statistic. A single event is generated when the current sampled value of the specified statistic becomes less than or equal to this threshold, and if the value at the last sampling interval was greater than this threshold.

owner <string>: The name of the owner of this alarm.

Examples

Figure 110 Configuring the RMON Alarm Parameters in the CLI

```
HP Switch(config)# rmon alarm 1 1.3.6.1.2.1.16.1.1.1.4.1 100 absolute rising-  
threshold 500 falling-threshold 300 owner Joe
```

Figure 111 Removing an RMON Alarm

```
HP Switch(config)# no rmon alarm 1
```

Figure 112 The show Command Output for a Specific Alarm

```
HP Switch(config)# show rmon alarm 1  
  
Alarm table 1 owned by Joe is VALID.  
Sample Type          : absolute  
Object Variable      : 1.3.6.1.2.1.16.1.1.1.4.1 <etherStatsOctets.1>  
  
Sampling Interval    : 100(sec)  
Rising Threshold     : 500  
Falling Threshold    : 300  
Startup Alarm        : risingOrFallingAlarm  
Last Sampled Value   : 2550  
Last Threshold Time  : Fri Jun 19 10:03:31 2012
```

Figure 113 The display Command Output for a Specific Alarm

```
HP Switch(config)# display rmon alarm 1  
  
Alarm table 1 owned by Joe is VALID.  
Samples type         : absolute  
Variable formula     : 1.3.6.1.2.1.16.1.1.1.4.1 <etherStatsOctets.1>  
  
Sampling interval    : 100(sec)  
Rising threshold     : 500  
Falling threshold    : 300  
Startup alarm        : risingOrFallingAlarm  
Latest value         : 2550  
Last threshold time  : Fri Jun 19 10:03:31 2012
```

Figure 114 Output of the running-config File Displaying the Configured RMON Alarm Parameters

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

Running configuration:

; J9470A Configuration Editor; Created on release #K.15.13.0000x
; Ver #04:2f.ff.3f.ef:04
hostname "HP-3500-24"
module 1 type j94dda
snmp-server community "public" unrestricted
snmp-server host 15.255.133.156 community "public"
snmp-server host 15.255.133.146 community "public"
vlan 1
  name "DEFAULT_VLAN"
  untagged 1-24
  ip address dhcp-bootp
  exit
vlan 2
  name "VLAN2"
  ip address 10.10.10.100 255.255.255.0
  exit
spanning-tree
rmon alarm 1 "etherStatsOctets.1" 100 absolute rising-threshold 500
  falling-threshold 300 owner "Joe"
```

CLI-configured sFlow with multiple instances

In earlier software releases, sFlow was configured on the switch via SNMP using a single sFlow instance. Beginning with software release K.11.34, 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.

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 228).

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 \(media-endpoint-discovery\)” \(page 228\)](#).

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 \(media-endpoint-discovery\)” \(page 228\)](#).

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 \(page 189\)](#).)

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 \(media-endpoint-discovery\)” \(page 228\)](#).

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 \(page 189\)](#).)

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 \(page 189\)](#).)

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 191).) 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 (SNMP notification support (page 226).)

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 192).)

Table 19 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 "MAC Address Management" (page 463)
Remote Management Address			
Type ^{5, 3}	N/A	Always Enabled	Shows the network address type.

Table 19 Data available for basic LLDP advertisements *(continued)*

Data type	Configuration options	Default	Description
Address ⁵	Default or Configured		Uses a default address selection method unless an optional address is configured. See “Remote management address” (page 224) .
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.
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. (See [“Changing the time-to-live for transmitted advertisements \(CLI\)” \(page 189\)](#).)

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” \(page 225\)](#).

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 [“Troubleshooting” \(page 379\)](#). 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 \(CLI\)” \(page 187\)](#).)
- 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 228).)

LLDP operating rules

For additional information specific to LLDP-MED operation, see “[LLDP-MED \(media-endpoint-discovery\)](#)” (page 228).

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 ([Section \(page 192\)](#).) (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.

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.

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`.

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.

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.

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.

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.

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)

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” \(page 201\)](#).

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.

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)

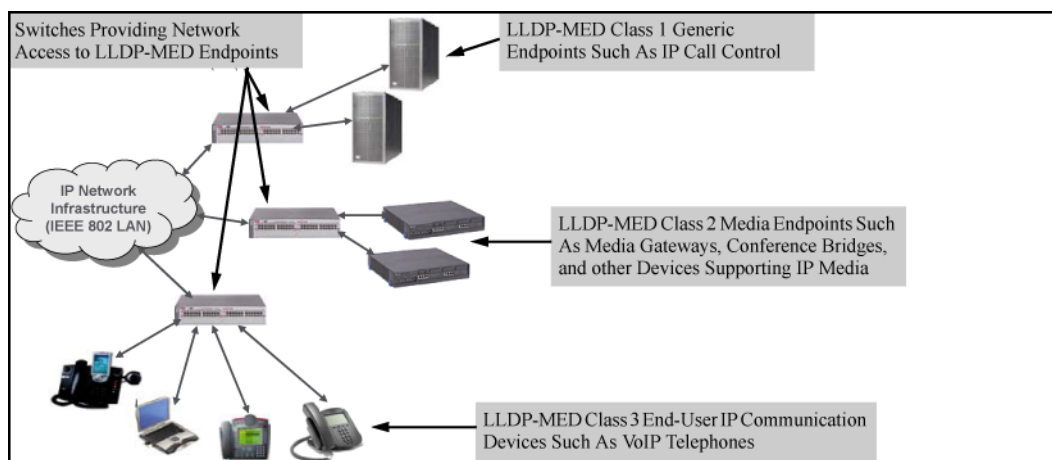
LLDP-MED (ANSI/TIA-1057/D6) extends the LLDP (IEEE 802.1AB) industry standard to support advanced features on the network edge for Voice Over IP (VoIP) endpoint devices with specialized capabilities and LLDP-MED standards-based functionality. LLDP-MED in the switches uses the standard LLDP commands described 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 115 Example of LLDP-MED network elements



LLDP-MED endpoint support

LLDP-MED interoperates with directly connected IP telephony (endpoint) clients having these features and services:

- Autonegotiate speed and duplex configuration with the switch
- Use the following network policy elements configured on the client port
- Voice VLAN ID

- 802.1p (Layer 2) QoS
- Diffserv codepoint (DSCP) (Layer 3) QoS
- Discover and advertise device location data learned from the switch
- Support ECS (such as E911, 999, and 112)
- Advertise device information for the device data inventory collected by the switch, including:

<ul style="list-style-type: none"> • Hardware revision • Firmware revision • Software revision 	<ul style="list-style-type: none"> • Serial number • Manufacturer name • Model name 	<ul style="list-style-type: none"> • Asset ID
---------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------	--------------------------------------------------------------

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

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; page 14-41), 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.

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 \(page 197\)](#))

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 \(CLI\)](#)” [\(page 193\)](#).)

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 *Advanced Traffic Management Guide*.

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.

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.)

About 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 the *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 20 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

Table 20 Some location codes used in CA-TYPE fields *(continued)*

Location element	Code	Location element	Code
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:

Description	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 51 \(page 233\)](#) shows the commands for configuring and displaying the above data.

Example 51 Example of 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
```

About 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` (Example 17 (page 48)) and `show lldp info remote-device [port-list]` (Figure 34 (page 87)) 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 201).

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.

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 \(page 192\)](#)) 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 \(CLI\)” \(page 189\)](#).

Mandatory TLVs

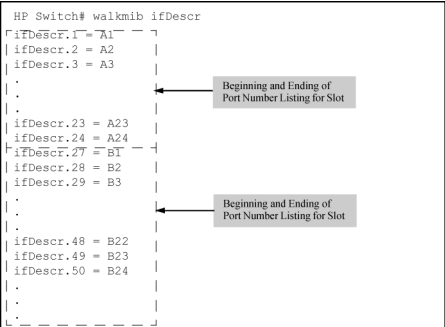
All mandatory TLVs required for LLDP operation are also mandatory for LLDP-MED operation.

About 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 116 \(page 234\)](#).

Figure 116 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
.
.
.
L-----L
```



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.

7 Chassis Redundancy (8200zl Switches)

Command Summary

For an overview of Chassis Redundancy, see [“Overview of chassis redundancy \(8200zl switches\)” \(page 256\)](#).

For information about how to use redundant management, see [“About using redundant management” \(page 257\)](#).

Viewing management module redundancy status (CLI)

You can display the status of both the management and fabric redundant modules using this command:

Syntax

```
show redundancy
```

Displays the status of the management and fabric modules.

Example

The output for the `show redundancy` command is seen in [Figure 117 \(page 237\)](#).

Figure 117 `show redundancy` command for management and fabric modules

```
HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Nonstop switching enabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 0
Last Failover  :

Slot Module Description                               Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200zl  Active   K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200zl  Standby  K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200zl      Enabled
FM2  HP Switch J9093A Fabric Module 8200zl      Enabled
```

Enabling or disabling redundant management (CLI)

For more information, see [“How the management modules interact” \(page 256\)](#).

There are two modes for management module redundancy—warm standby mode (the default) and Nonstop switching mode. In warm-standby mode, 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.

In Nonstop switching mode, 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.

Syntax

```
[no] redundancy management-module [nonstop-switching]
```

Allows enabling or disabling of redundant management. The current active module continues to be the active module on boot unless you use the `redundancy active-management` command to enable redundant behavior.

(Default: Warm-standby redundancy mode)

The `nonstop-switching` parameter sets the redundancy mode to Nonstop switching.

You are prompted with "All configuration files and software images on the off-line management module will be overwritten with the data from the current active management module. During initial syncing from active to standby management module configuration changes are disallowed. Do you want to continue [y/n]?"

When the `nonstop-switching` option is *not* selected, the switch enters warm-standby redundancy mode.

You are prompted with "All configuration files and software images on the off-line management module will be overwritten with the data from the current active management module. Do you want to continue [y/n]?"

The `no` version of the command disables redundant management. You are prompted with this message: "The other management module may reboot and it will no longer be used for system redundancy, except in the case of a hardware failure of the active management module. Do you want to continue [y/n]?"

Example

The `redundancy management-module` command in [Figure 118 \(page 238\)](#) shows **warm-standby redundant management** being **enabled**. The `show redundancy` command displays "**Mgmt Redundancy**" as **warm-standby redundancy enabled**. Management Module 1 (MM1) is the active management module and Management Module 2 (MM2) is the standby management module.

Figure 118 Enabling warm-standby redundancy

```
HP Switch(config)# redundancy management-module
All configuration files and software images on the off-line management
module will be overwritten with the data from the current active
management module. Do you want to continue [y/n]? y

HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Warm-standby redundancy enabled ← Redundancy enabled
Rapid Switchover Stale Timer : 1

Statistics
-----
Failovers      : 0
Last Failover  :

Slot Module Description                Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200z1 Active   K.15.01.000x Secondary
MM2  HP Switch J9092A Management Module 8200z1 Standby  K.15.01.000x Secondary

FM1  HP Switch J9093A,Fabric Module 8200z1  Enabled
FM2  HP Switch J9093A,Fabric Module 8200z1  Enabled
```

The `redundancy management-module` command in [Figure 119 \(page 239\)](#) shows Non-stop switching redundant management being **enabled**. The `show redundancy` command displays "**Mgmt Redundancy**" as **Nonstop switching enabled**. Management Module 1 (MM1) is the standby management module and Management Module 2 (MM2) is the active management module.

Figure 119 Enabling nonstop-switching redundancy

```
HP Switch(config)# redundancy management-module nonstop-switching
All configuration files and software images on the off-line management module
will be overwritten with the data from the current active management module.
During initial syncing from active to standby management module configuration
changes are disallowed. Do you want to continue [y/n]? y

HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Nonstop switching enabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 0
Last Failover  :

Slot Module Description                Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200z1 Standby  K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200z1 Active   K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200z1  Enabled
FM2  HP Switch J9093A Fabric Module 8200z1  Enabled
```

← Redundancy enabled

The no version of the redundancy management-module command is used to disable management module redundancy on the switch, as seen in Figure Figure 120 (page 239). The show redundancy command displays "**Mgmt Redundancy**" as **Nonstop switching disabled**. The standby management module in slot MM1 is now offline. The management module in slot MM2 remains the active management module.

NOTE: HP recommends that you leave management module redundancy enabled. If the active management module has a hardware failure, the standby module may take over and may have an old configuration since file synchronization has not occurred when management module redundancy was disabled.

The no redundancy management-module command allows you to shut down a management module that is not functioning correctly without physically removing the module. If you want to remove the module, first perform the shutdown procedure as explained in "Hotswapping out the active management module" (page 245) and then remove the module.

Figure 120 Disabling redundancy

```
HP Switch(config)# no redundancy management-module
The other management module may reboot and it will no longer be used for system
redundancy except in the case of a hardware failure of the active management
module. Do you want to continue[y/n]? y

HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Nonstop switching disabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 1
Last Failover  : Tue Mar 19 12:42:31 2009

Slot Module Description                Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200z1 Offline  K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200z1 Active   K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200z1  Enabled
FM2  HP Switch J9093A Fabric Module 8200z1  Enabled
```

← Nonstop switching disabled

The `redundancy management-module` command shows Nonstop switching redundant management being enabled. The `show redundancy` command displays “Mgmt Redundancy” as Nonstop switching enabled. Management Module 1 (MM1) is the standby management module and Management Module 2 (MM2) is the active management module.

Example

Enabling non-stop switching redundancy.

```
HP Switch(config)# redundancy management-module nonstop-switching
All configuration files and software images on the off-line management module
will be overwritten with the data from the current active management module.
During initial syncing from active to standby management module configuration
changes are disallowed. Do you want to continue [y/n]? y
HP Switch(config)# show redundancy
Settings
-----
Mgmt Redundancy : Nonstop switching enabled
Rapid Switchover Stale Timer : 0
Statistics
-----
Failovers : 0
Last Failover :
Slot Module      Description                               Status   SW Version   Boot Image
-----
MM1  HP J9092A Management Module 8200z1  Standby  K.15.01.000x Primary
MM2  HP J9092A Management Module 8200z1  Active   K.15.01.000x Primary
FM1  HP J9093A Fabric Module 8200z1      Enabled
FM2  HP J9093A Fabric Module 8200z1      Enabled
```

The `no version of the redundancy management-module` command is used to disable management module redundancy on the switch, as seen in Figure 7-4. The `show redundancy` command displays “Mgmt Redundancy” as Nonstop switching disabled. The standby management module in slot MM1 is now offline. The management module in slot MM2 remains the active management module.

NOTE: HP recommends that you leave management module redundancy enabled. If the active management module has a hardware failure, the standby module may take over and may have an old configuration since file synchronization has not occurred when management module redundancy was disabled.

The `no redundancy management-module` command allows you to shut down a management module that is not functioning correctly without physically removing the module. If you want to remove the module, first perform the shutdown procedure as explained in “Hotswapping Out the Active Management Module” on page 7-25, and then remove the module.

Example

Disabling redundancy:

```
HP Switch(config)# no redundancy management-module
The other management module may reboot and it will no longer be used for system
redundancy except in the case of a hardware failure of the active management
module. Do you want to continue[y/n]? y
HP Switch(config)# show redundancy
Settings
-----
Mgmt Redundancy : Nonstop switching disabled
Rapid Switchover Stale Timer : 0
Statistics
-----
Failovers : 1
Last Failover : Tue Mar 19 12:42:31 2009
Slot Module      Description                               Status   SW Version   Boot Image
-----
MM1  HP J9092A Management Module 8200z1  Offline  K.15.01.000x Primary
MM2  HP J9092A Management Module 8200z1  Active   K.15.01.000x Primary
FM1  HP J9093A Fabric Module 8200z1      Enabled
FM2  HP J9093A Fabric Module 8200z1      Enabled
```


Transitioning from no redundancy to nonstop switching

While the switch is transitioning from no redundancy mode to Nonstop switching mode, no configuration changes are allowed. The management modules are syncing information during the transition period.

Setting the Rapid Switchover Stale Timer

Use the Rapid Switchover Stale Timer to set the amount of time that you want route and neighbor table entries to be re-added to the Forwarding Information Base on the active management module after a failover has occurred.

Layer 3 applications and protocols rely on existing routing information in the FIB. They restart and operate as if the switch performed a quick reset.

When a failover occurs, the interface modules and the fabric modules continue forwarding Layer 3 traffic based on the information in the FIB. The transitioning standby management module marks all routes in the FIB as “stale”. The routing protocols restart, reestablish their neighbors and reconverge. As the routes are added in again, the route’s stale designation is removed. After the Rapid Switchover Stale Timer expires, the remaining stale route entries are removed. Multicast flows are also removed; the multicast application re-adds the flows after failover completes.

Syntax

```
redundancy rapid-switchover <0-2147483647>
```

Allows configuration of a timer (in seconds) for Layer 3 forwarding of packets when Nonstop switching is configured for redundancy. 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 (entries not re-added) are removed.

A setting of zero indicates that no Layer 3 Nonstop switching behavior is wanted.

When the switch fails over, the FIB entries and corresponding hardware entries are removed. Default: 90 seconds

To display information about stale FIB routes, enter the `show tech route stale` command. The VLAN ID and IP route are shown, as well as other information used only for technical support.

Directing the standby module to become active

To make the standby management module become the active management module, use the `redundancy switchover` command. The switch will switchover after all files have finished synchronizing.

In Nonstop switching mode:

- The switchover occurs quickly and seamlessly. No reboot is needed.
- There is no interruption in switching operations.

In warm-standby mode:

- The switchover may take a couple of minutes if there have been recent configuration file changes or if you have downloaded a new operating system.
- The standby module finishes booting and becomes the active module.

The formerly active module becomes the standby module if it passes selftest.

Syntax

```
redundancy switchover
```

Causes a switchover to the standby module.

For Nonstop switching, the warning displays: A nonstop switching failover will occur; L2 operations will not be interrupted. This management module will now reboot and will become the standby module! You will need to use the other management module's console interface. Do you want to continue [y/n]?

In warm-standby mode the warning displays: A warm failover will occur; all networking operations will be interrupted. This management module will now reboot and will become the standby module! You will need to use the other management module's console interface. Do you want to continue [y/n]?

If management module redundancy has been disabled, or there is no standby module, or the standby module is not in standby mode, this message displays: The other management module does not exist or is not in standby mode. An example of the redundancy switchover command when the switch is in Nonstop switching mode is shown in the example below.

Example

Redundancy switchover command when in nonstop switching mode.

```
HP Switch(config)# redundancy switchover
A nonstop switching failover will occur; L2 operations will not be interrupted.
This management module will now reboot and will become the standby
module! You will need to use the other management module's console interface.
Do you want to continue [y/n]? y
This management module will now boot from the primary image and will
become the standby module! You will need to use the other management module's
console interface. Do you want to continue [y/n]? y
ROM information:
Build directory: /sw/rom/build/bmrom(t2g)
Build date: Oct 15 2009
Build time: 08:24:27
Build version: K.15.01
Build number: 13040
Select profile (primary):
Booting Primary Software Image...
...
Standby Console>
```

Setting the rapid switchover stale timer (CLI)

For more information about transitioning from no redundancy to nonstop switching and setting the rapid switchover stale timer, see [“Transition from no redundancy to nonstop switching” \(page 257\)](#) and [“About setting the rapid switchover stale timer” \(page 257\)](#).

Syntax

```
redundancy rapid-switchover 0-2147483647
```

Allows configuration of a timer (in seconds) for Layer 3 forwarding of packets when nonstop switching is configured for redundancy. 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 (entries not re-added) are removed.

A setting of zero indicates that no Layer 3 Nonstop switching behavior is wanted. When the switch fails over, the FIB entries and corresponding hardware entries are removed.

(Default: 45 seconds)

To display information about stale FIB routes, enter the `show tech route stale` command. The VLAN ID and IP route are shown, as well as other information used only for technical support.

Directing the standby module to become active (CLI)

For more information, see “About directing the standby module to become active” (page 257).

Syntax

```
redundancy switchover
```

Causes a switchover to the standby module.

For nonstop switching, the warning displays: "A nonstop switching failover will occur; L2 operations will not be interrupted. This management module will now reboot and will become the standby module! You will need to use the other management module's console interface. Do you want to continue [y/n]?"

In warm-standby mode the warning displays: "A warm failover will occur; all networking operations will be interrupted. This management module will now reboot and will become the standby module! You will need to use the other management module's console interface. Do you want to continue [y/n]?"

If management module redundancy has been disabled, or if there is no standby module, or if the standby module is not in standby mode, this message displays:

The other management module does not exist or is not in standby mode

Example

Figure 121 (page 243) shows an example of the `redundancy switchover` command when the switch is in **nonstop switching** mode.

Figure 121 The `redundancy switchover` command when in nonstop switching mode

```
HP Switch(config)# redundancy switchover
A nonstop switching failover will occur; L2 operations will not be interrupted. This management module will now reboot and will become the standby module! You will need to use the other management module's console interface. Do you want to continue [y/n]? y
This management module will now boot from the primary image and will become the standby module! You will need to use the other management module's console interface. Do you want to continue [y/n]? y

ROM information:
  Build directory: /sw/rom/build/bmrom(t2g)
  Build date:      Oct 15 2009
  Build time:      08:24:27
  Build version:   K.15.01
  Build number:    13040
Select profile (primary):

Booting Primary Software Image...
.
.
.

Standby Console>
```

Setting the active management module for next boot (CLI)

Syntax

```
redundancy active-management [ management-module1 | management-module2
| standby ]
```

The specified module becomes the active management module at the next system boot. This message displays: On the next system boot, the module specified will become active.

This command does not take effect if the standby management module has failed selftest.

management-module1	Configures management-module 1 as the active management module for the next system boot.
management-module2	Configures management-module 2 as the active management module for the next system boot.
standby	Configures the current standby module as the active management module for the next system boot if management module redundancy is enabled. If redundancy is disabled, it becomes enabled as a standby module at the next boot or failover event.

If the specified management module is not there or is in failed mode, this message displays:

The specified module is not present or is in failed state.

Example

Figure 122 (page 244) shows an example of setting **management module 2** to be the **active management module**.

Figure 122 Setting a management module to be active on the next boot

```

HP Switch(config)# redundancy active-management management-module2
On the next system boot, the management-module2 will become active.

HP Switch(config)# boot system
/boot occurs...

HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Nonstop Switching enabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 0
Last Failover :

Slot Module Description                Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200z1 Standby  K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200z1 Active   K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200z1   Enabled
FM2  HP Switch J9093A Fabric Module 8200z1   Enabled

```

If management module redundancy has been disabled and you specify the standby module with the active-management command, upon rebooting, the offline module becomes the standby module. The state of redundancy (enabled or disabled) is based on the value in the configuration file in the offline (now standby) module. The configuration files have not been synchronized if management module redundancy has been disabled. An example of making the offline management module become the standby management module when **redundancy** is disabled is shown in Figure 123 (page 245).

Figure 123 Showing the results of switching to standby module when redundancy is disabled

```

HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Nonstop switching disabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 0
Last Failover :

Slot Module Description                               Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200z1  Active   K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200z1  Offline  K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200z1    Enabled
FM2  HP Switch J9093A Fabric Module 8200z1    Enabled

HP Switch(config)# redundancy active-management standby
On the next system boot, the standby will become active.
Redundancy and Synchronization have been disabled, so it will
not have current configurations.

HP Switch(config)# boot
The other management module is not in standby mode and this command will
not cause a switchover. System will reboot from primary image.
Do you want to continue [y/n]? y

(After system reboots...)

HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Nonstop switching disabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 0
Last Failover :

Slot Module Description                               Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200z1  Standby  K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200z1  Active   K.15.01.000x Primary
    
```

Nonstop switching disabled


When redundancy is disabled and the redundancy active-management standby command is executed, the offline MM becomes the active MM.

Hotswapping out the active management module

For information about management module switchover and hotswapping management modules, see “Management module switchover” (page 263) and “About hotswapping out the active management module” (page 264).

1. On the management module to be hotswapped out, press the **MM Shutdown** button. It is located between the **Module Operation** and **Component Status** LEDs. (See Figure 124 (page 245).)

Figure 124 The MM Shutdown button



2. The **Dwn** LED to the right of the **MM Shutdown** button begins flashing green. File synchronization will complete before shutdown occurs.

3. The standby module takes control and the switchover occurs. It is now the active management module.
4. The **Dwn** LED on the management module being hotswapped out turns green and all other LEDs go out when it is OK to remove the module.
5. The module being hotswapped out goes into offline mode. In the offline mode, the module cannot take over when the active module fails over.

NOTE: If you remove the active management module without pressing the **MM Shutdown** button, any files that may have been in the process of synchronizing will not finish synchronizing to the standby module and all file transfer is aborted.

For information about hotswapping in a management mode and software version mismatch between active and hotswapped modules, see [“Hotswapping in a management module”](#) (page 264) and [“Software version mismatch between active and hotswapped module”](#) (page 265).


Resetting the management module

The **MM Reset** button, shown in [Figure 125](#) (page 246), found on each management module reboots its management module. If the management module is active and management module redundancy is enabled, switchover occurs. The standby management module is notified immediately. It then takes over and becomes the active management module. If the **MM Reset** button is pressed on the standby management module, that module reboots but no other switch operations are affected. The active management module remains in control.

If management module redundancy is disabled, the active management module reboots and remains in control, as long as it passes selftest.

CAUTION: HP does not recommend using the **MM Reset** button to trigger a switchover. Files being copied over at the time of the reset will be aborted.

Figure 125 The **MM Reset** button on the 8200zl management module



For information about downloading a new software version, see [“About downloading a new software version”](#) (page 266).

For information about synchronizing files after downloading, see [“File synchronization after downloading”](#) (page 266).

For information about software version mismatches after downloading, see [“Potential software version mismatches after downloading”](#) (page 266).

Viewing management information

For information about active management module commands, see [“Disable management module redundancy with only one module present”](#) (page 269). The sections that follow provide information about the `show` commands that display various aspects of management information.

Syntax

```
show modules [details]
```

Displays information about the installed modules, including:

- The slot in which the module is installed
- The module description
- The serial number
- The status
- Core dump
- Model Version

Additionally, the part number (J number) and serial number of the chassis is displayed.

Example

```
show modules
Chassis: 8212z1 J9091A          Serial Number:  LP713BX004
Allow V1 Modules:  Yes

Core Mod
Slot Module Description          Serial Number  Status  Dump  Ver
-----
MM1 HP J9092A Management Module 8200z1  sg844bp012   Active  NO    1
SSM HP J9095A System Support Module          SG911BZ00N
FM1 HP J9093A Fabric Module 8200z1          SG911BQ015   Enabled -    1
FM2 HP J9093A Fabric Module 8200z1          SG911BQ04T   Enabled -    1
A   HP J9536A 20p GT PoE+/2p SFP+ v2 z1... SG0607T124   Up      YES   2
B   HP          Enh Svs v2 z1 Module          Up      YES   2
C   HP J8702A 24p Gig-T z1 Module          Up      NO    1
D   HP J9840A Adv Svs v2 z1 Module          ID3ZG6N008   Up      YES   2
E   HP J8705A Gig-T/SFP z1 Module          Up      NO    1
F   HP J9857A Adv Svs v2 z1 Module          SG2ZFNX166   Up      YES   2
G   HP J8708A 4p 10G CX4 z1 Module          Up      NO    1
H   HP J9154A Services z1 Module          SG811GG01N   Up      NO    1
I   HP J9051A Wireless Edge Services z1... SG660ZB095   Up      NO    1
J   HP J9545A ONE Adv Svs z1 Module          SG9604P933   Up      NO    1
K   HP J9051A Wireless Edge Services z1... 1111        Up      NO    1
L   HP J9154A Services z1 Module          SG811GG01M   Up      NO    1
```

Viewing information about the management and fabric modules

The `show redundancy` command displays information about the management and fabric modules. It displays the flash image last booted from, even if the `boot set-default` command has been set to change the flash booted from on the next boot.

Example

Figure 126 show redundancy command

```
HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Nonstop switching enabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 0
Last Failover  :

Slot Module Description                Status  SW Version  Boot Image
-----
MM1 HP J9092A Management Module 8200z1  Standby  K.15.01.000x Primary
MM2 HP J9092A Management Module 8200z1  Active   K.15.01.000x Secondary

FM1 HP J9093A Fabric Module 8200z1    Enabled
FM2 HP J9093A Fabric Module 8200z1    Enabled
```

The active management module was last booted from secondary flash. The standby management module was last booted from primary flash.

Viewing information about the redundancy role of each management module

The `show redundancy` command with the `detail` option displays information about the redundancy role of each management module, as well as statistical information such as how long the module has been up.

Example

Figure 127 show redundancy detail command

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

Redundancy Information:

Slot Role      Card Up Since      Role Since          Redundancy State
-----
1  Active      11/11/09 23:40:22      11/04/09 23:33:15      Active
2  Standby     11/11/09 23:40:24      11/04/09 23:33:15      Nonstop switching

Fail-Over Log:

Slot Role      Time              Reason
-----
2  Standby     11/01/09 10:16:04      Standby Reset
2  Active      11/02/09 17:46:03      Hot Swap
1  Standby     11/03/09 15:39:06      Standby Reset
1  Active      11/04/09 09:25:39      Switchover
```

Viewing which software version is in each flash image

The `show flash` command displays which software version is in each flash image. The **Default Boot** field displays which flash image will be used for the next boot.

Example

Figure 128 show flash command

```
HP Switch(config)# show flash
Image          Size(Bytes)  Date      Version
-----
Primary Image  : 7463821  09/05/09  K.15.00.0001
Secondary Image: 7463821  09/05/09  K.15.00.0001

Boot Rom Version: K.15.07
Default Boot    : Primary
```

Will boot from primary flash on the next boot.

Viewing system software image information for both management modules

The `show version` command displays system software image information for both management modules, as well as which module is the active management module and which is the standby management module. The `Boot Image` field displays which flash image last booted from, even if the `boot set-default` command has been set to change the flash booted from on the next boot. The output of the `show version` command when redundancy is enabled is shown in [Figure 129 \(page 249\)](#).

Example

Figure 129 show version command when redundancy is enabled

```
HP Switch(config)# show version
Management Module 1: Standby
Image stamp:      /sw/code/build/btm(t2g)
                  Mar  5 2009 13:20:59
                  K.15.01.0001
                  351
Boot Image:      Primary

Management Module 2: Active
Image stamp:      /sw/code/build/btm(t2g)
                  Mar  5 2009 13:20:59
                  K.15.01.0001
                  351
Boot Image:      Primary
```

Both management modules were booted from primary flash.

When redundancy is disabled, the output of the `show version` command changes, as shown in [Figure 130 \(page 249\)](#).

Example

Figure 130 show version command when redundancy is disabled

```
HP Switch(config)# show version
Management Module 1: Redundancy and Synchronization has been disabled;
                    enable with the 'redundancy' command.

Management Module 2: Active
Image stamp:      /sw/code/build/btm(t2g)
                  Mar  5 2009 13:20:59
                  K.15.01.0001
                  351
Boot Image:      Primary
```

Viewing the status of the switch and its management modules

The `show logging` command displays the status of the switch and its management modules. See [“Displaying module events” \(page 253\)](#). To show log messages in reverse chronological order (most recent messages displayed first), enter `show log -r`.

Example

Figure 131 show log command output

```
HP Switch(config)# show logging
Keys:   W=Warning   I=Information
        M=Major     D=Debug   E=Error
----- Event Log listing: Events Since Boot -----
I 10/28/09 21:45:42 00061 system: AM1: -----
I 10/28/09 21:45:42 00062 system: AM1: Mgmt Module 1 went down without saving cr
ash information
M 10/28/09 21:45:42 03002 system: AM1: System reboot due to Reset Switch
I 10/28/09 21:45:42 02759 chassis: AM1: Savepower LED timer is OFF.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot A configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot B configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot C configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot D configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot E configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot F configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot G configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot H configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot I configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot J configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot K configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot L configured ON.
I 10/28/09 21:45:42 00937 chassis: AM1: Fabric 1 inserted
I 10/28/09 21:45:42 00937 chassis: AM1: Fabric 2 inserted
I 10/28/09 21:45:43 00092 dhcp: AM1: Enabling Auto Image Config Download via
DHCP and turning off auto-tftp if enabled
I 10/28/09 21:45:43 00690 udpf: AM1: DHCP relay agent feature enabled
I 10/28/09 21:45:43 02637 srcip: AM1: TACACS admin policy is 'outgoing interface'
I 10/28/09 21:45:43 02638 srcip: AM1: TACACS oper policy is 'outgoing interface'
```

AM1 = Active management module in slot 1
AM2 = Active management module in slot 2
SM1 = Standby management module in slot 1
SM2 = Standby management module in slot 2

Standby management module commands

The standby management module, by design, has very little console capability. You can use three commands—`show flash`, `show version`, and `show redundancy`. The `show redundancy` command displays when a management module is in standby mode.

Viewing redundancy status on the standby module (CLI)

Use the `show redundancy` command to display redundancy status on the standby module, as shown in Figure 132 (page 250). This command displays the flash image last booted from, even if the `boot set-default` command has been set to change the flash booted from on the next boot.

Example

Figure 132 show redundancy command for standby module

```
Standby Console> show redundancy

Settings
-----
Mgmt Redundancy : Nonstop Switching Enabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 1
Last Failover  : Mon Sep 26 09:50:40 2009

Slot Module Description                               Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200z1  Active  K.15.01.0001 Secondary
MM2  HP Switch J9092A Management Module 8200z1  Standby K.15.01.0001 Primary

FM1  HP Switch J9093A Fabric Module 8200z1      Enabled
FM2  HP Switch J9093A Fabric Module 8200z1      Enabled
```

The active management module was last booted from secondary flash. The standby management module was last booted from primary flash.

Viewing the flash information on the standby module (CLI)

Use the `show flash` command to display the flash information on the standby module, as shown in [Figure 133 \(page 251\)](#). The **Default Boot** field displays which **flash image** will be used for the next boot.

Example

Figure 133 `show flash` command for standby module

```
Standby Console> show flash
Image           Size(Bytes)   Date   Version
-----
Primary Image   : 7493854   09/21/09 K.15.00.0001
Secondary Image : 7463821   09/05/09 K.15.00.0001

Boot Rom Version: K.15.07
Default Boot    : Primary
```

← Will boot from primary flash on the next boot.

Viewing the version information on the standby module (CLI)

Use the `show version` command to display the version information on the standby module, as shown in [Figure 134 \(page 251\)](#). The **Boot Image** field displays which flash image was last booted from, even if the `boot set-default` command has been set to change the flash booted from on the next boot. Unlike executing the `show version` command on an active management module, this command shows only the running version of software on the standby management module.

Example

Figure 134 `show version` command for standby module

```
Standby Console> show version
Image stamp:   /sw/code/build/btm(t2g)
               Mar 21 2009 15:03:31
               K.15.01.0001
               1617
Boot Image:    Primary ← Was booted from primary flash.
```

Setting the default flash for boot (CLI)

For more information about the boot command, see [“CLI commands affected by redundant management” \(page 269\)](#).

You can set which flash image to boot from as the default image on boot by using this command:

Syntax

```
boot set-default flash[ primary | secondary ]
```

Sets the flash image to boot from on the next boot.

primary	Boots the primary flash image.
secondary	Boots the secondary flash image.

Example

[Figure 135 \(page 252\)](#) shows an example of the output when the command is used to set the boot default to secondary flash.

Figure 135 `boot set-default` command defaulting to secondary flash

```
HP Switch(config)# show flash
Image                Size(Bytes)    Date    Version
-----
Primary Image       : 7463821    11/05/09 K.15.01.0001
Secondary Image     : 7463821    11/05/09 K.15.01.0001

Boot Rom Version: K.15.07
Default Boot      : Primary

HP Switch(config)# boot set-default flash secondary
This command changes the location of the default boot. This
command will change the default flash image to boot from
secondary. Hereafter, 'reload' and 'boot' commands will boot
from secondary. Do you want to continue [y/n]? y

HP Switch(config)# show flash
Image                Size(Bytes)    Date    Version
-----
Primary Image       : 7463821    03/05/09 K.15.01.0001
Secondary Image     : 7463821    03/05/09 K.15.01.0001

Boot Rom Version: K.15.07
Default Boot      : Secondary
```

Booting the active management module from the current default flash (CLI)

Use the `reload` command to boot the active management module from the current default flash (You can change the default flash with the `boot set-default` command. See [“Setting the default flash for boot \(CLI\)”](#) (page 251).) Switchover occurs if redundancy is enabled and the standby management module is in standby mode. If redundancy is disabled or the standby management module is not present, the `reload` command boots the system.

NOTE: The `reload` command is a "warm" reboot; it skips the Power on Self Test routine.

Syntax

```
reload <cr>
```

Boots (warm reboot) the active management module. Switchover to the standby management module occurs if management module redundancy is enabled. If redundancy is disabled or if there is no standby management module, the `reload` command boots the system.

NOTE: If the running config file is different from the stored config file, you are prompted to save the config file. The `reload at/after` versions of this command do not display a prompt to save configuration file changes: the changes are lost on the scheduled reload.

Example

Figure 136 reload command with redundancy enabled

```
HP Switch(config)# reload
This command will cause a switchover to the other management module
which may not be running the same software image and configurations.
Do you want to continue [y/n]? y

(Boots...)

HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Nonstop Switching Enabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 1
Last Failover : Mon April 30 09:10:11 2009

Slot Module Description                               Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200z1 Active   K.15.01.0001 Primary
MM2  HP Switch J9092A Management Module 8200z1 Standby K.15.01.0001 Primary
```

For information on other commands that are affected by redundant management, see [“Additional commands affected by redundant management”](#) (page 272).

Displaying module events

Viewing log events (CLI)

The log file displays messages about the activities and status of the management modules. Enter this command to display the messages:

Syntax

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

Displays log events.

The event messages are tagged with the management module state and the management module slot (AM1 or AM2, SM1 or SM2.) Synchronization is maintained by syncing the standby management module log events with the active management module. In this way, events are available for both management modules. Only the active management module events are shown unless you select the `-s` option. This option works like the `-a` option, except that the events for both the active management module and standby management module are displayed.

For more information on command options available with the `show logging` command, see [“Troubleshooting”](#) (page 379).

Example

Figure 137 Log file listing

```
HP Switch(config)# show logging
Keys:  W=Warning  I=Information
      M=Major    D=Debug  E=Error
----- Event Log listing: Events Since Boot -----
I 10/28/09 21:45:42 00061 system: AM1: -----
I 10/28/09 21:45:42 00062 system: AM1: Mgmt Module 1 went down without saving
      crash information
M 10/28/09 21:45:42 03002 system: AM1: System reboot due to Reset Switch
I 10/28/09 21:45:42 02759 chassis: AM1: Savepower LED timer is OFF.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot A configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot B configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot C configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot D configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot E configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot F configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot G configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot H configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot I configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot J configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot K configured ON.
I 10/28/09 21:45:42 02751 chassis: AM1: LEDs for module in slot L configured ON.
I 10/28/09 21:45:42 00937 chassis: AM1: Fabric 1 inserted
I 10/28/09 21:45:42 00937 chassis: AM1: Fabric 2 inserted
I 10/28/09 21:45:43 00092 dhcp: AM1: Enabling Auto Image Config Download via
      DHCP and turning off auto-tftp if enabled
I 10/28/09 21:45:43 00690 udpf: AM1: DHCP relay agent feature enabled
I 10/28/09 21:45:43 02637 srcip: AM1: TACACS admin policy is 'outgoing interface'
I 10/28/09 21:45:43 02638 srcip: AM1: TACACS oper policy is 'outgoing interface'
```

AM1 = Active management module in slot 1
AM2 = Active management module in slot 2
SM1 = Standby management module in slot 1
SM2 = Standby management module in slot 2

Copying crash file information to another file (CLI)

Crash logs for all modules are always available on the active management module. You can use the `copy crash-log` and `copy crash-data` commands to copy the information to a file of your choice.

Syntax

```
copy crash-log [ slot-id | mm ] tftp ip-address filename
```

Copies the crash logs of both the active and standby management modules to a user-specified file. If no parameter is specified, files from all modules (management and interface) are concatenated.

slot-id	Retrieves the crash log from the module in the specified slot.
mm	Retrieves the crash logs from both management modules and concatenates them.

Syntax

```
copy crash-data [ slot-id | mm ] tftp ip-address filename
```

Copies the crash data of both the active and standby management modules to a user-specified file. If no parameter is specified, files from all modules (management and interface) are concatenated.

slot-id	Retrieves the crash data from the module in the specified slot.
mm	Retrieves the crash data from both management modules and concatenates them.

Viewing saved crash information

Syntax

```
show boot-history
```

Displays the system boot log.

Example

Figure 138 The system boot log file

```
HP Switch Switch 8200zl$ show boot-history

Mgmt Module 1 -- Saved Crash Information (most recent first):
=====
Mgmt Module 1 in Active Mode went down: 11/07/09 14:48:36
Operator warm reload from CONSOLE session.

Mgmt Module 1 in Active Mode went down: 11/07/09 11:43:10
Operator cold reboot from CONSOLE session.

Mgmt Module 2 -- Saved Crash Information (most recent first):
=====
No Saved Crash Information
```

Enabling and disabling fabric modules (CLI)

The fabric modules can be enabled or disabled even if they are not present in the switch. You cannot disable both fabric modules at the same time; one must be enabled.

Use this command to enable or disable the redundant fabric modules. Disabling one fabric module reduces the overall switching capacity of the 8200zl series switches. On some networks where network utilization is less than 50%, you may not notice any degradation of performance.

Syntax

```
redundancy fabric-module [ 1 | 2 ] [ enable | disable ]
```

Allows enabling or disabling of fabric modules. (You cannot have both fabric modules disabled at the same time.)

Default: Both fabric modules are enabled.

NOTE: The redundant fabric modules do not support nonstop switching.

Example

Figure 139 Disabling a fabric module

```
HP Switch(config)# redundancy fabric-module 2 disable
HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Nonstop switching enabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 0
Last Failover :

Slot Module Description                Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200z1 Active   K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200z1 Standby  K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200z1 Enabled
FM2  HP Switch J9093A Fabric Module 8200z1 Disabled
```

Overview of chassis redundancy (8200zl switches)

The HP 8200zl switches provide high availability through the use of hot-swappable, redundant management modules. In the event of a failure on the active management module, management module redundancy allows a quick and unattended transition from the active management module to the standby management module. The standby management module now becomes the active management module. Management module redundancy keeps the switch operating and reduces network downtime.

The advantages of redundant management are:

- Maintaining switch operation if a hardware failure occurs on the active management module
- Minimizing restart time caused by the failure of a management module
- Hotswapping a failed management module with no downtime

Nonstop switching with redundant management modules

Beginning with software version K.15.01, you can use either nonstop switching or warm-standby redundant management.

The advantages of nonstop switching are:

- Quick, seamless transition to the standby management module; no reboot is necessary
- Switching of packets continues without interruption

How the management modules interact

When the switch boots up, the management modules run selftest to decide which is the active module and which is the standby module (see [“Notes on how the active module is determined” \(page 274\).](#)) The module that becomes active finishes booting and then brings up the interface modules and ports.

If you are using nonstop switching mode, 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. If the active management module fails or if there is a manual switchover, switching continues without interruption.

If you are using warm-standby mode, the standby module boots to a certain point, syncs basic files such as the config and security files, and finishes booting only if the active management module fails or you choose to change which module is the active module.

The two management modules communicate by sending heartbeats back and forth.

About using redundant management

The CLI commands for redundant management are shown at the beginning of the chapter. Additionally, some other commands are affected by redundant management (See “[CLI commands affected by redundant management](#)” (page 269).)

Transition from no redundancy to nonstop switching

While the switch is transitioning from no redundancy mode to nonstop switching mode, no configuration changes are allowed. The management modules are syncing information during the transition period.

About setting the rapid switchover stale timer

After a failover has occurred, use the rapid switchover stale timer to set the amount of time that you want route and neighbor table entries to be re-added to the FIB on the active management module.

Layer 3 applications and protocols rely on existing routing information in the FIB. They restart and operate as if the switch performed a quick reset.

When a failover occurs, the interface modules and the fabric modules continue forwarding Layer 3 traffic based on the information in the FIB. The transitioning standby management module marks all routes in the FIB as "stale". The routing protocols restart, reestablish their neighbors and reconverge. As a route is added in again, the route's stale designation is removed. After the rapid switchover stale timer expires, the remaining stale route entries are removed. Multicast flows are also removed; the multicast application re-adds the flows after failover completes.

About directing the standby module to become active

To make the standby management module become the active management module, use the `redundancy switchover` command. The switch will switchover after all files have finished synchronizing.

In nonstop switching mode:

- The switchover occurs quickly and seamlessly; no reboot is needed.
- There is no interruption in switching operations.

In warm-standby mode:

- The switchover may take several minutes if there have been recent configuration file changes or if you have downloaded a new operating system.
- The standby module finishes booting and becomes the active module.

The formerly active module becomes the standby module if it passes selftest.

Nonstop switching with VRRP

When Nonstop VRRP is enabled, VRRP continues to operate in its current state when a failover from the AMM to the SMM occurs. This provides an additional layer of redundancy in a switched network. VRRP state information is maintained between MMs so that VRRP operations resume immediately after failover from the AMM to SMM. Because of this quick resumption of operations there is no failover to the backup VRRP router in the network. The Master VRRP router continues to be active and operate as is.

The command for enabling Nonstop mode for VRRP must be executed in VRRP context. For more information on enabling and configuring VRRP, see the *Multicast and Routing Guide*.

Syntax

```
(vrrp#) [no]  
nonstop
```

Enabling Nonstop VRRP allows the VRRP router to retain control of IP addresses when the AMM fails over. The VRRP Backup router does not take control of the virtual IP addresses on the network.

The no version of the command disables Nonstop VRRP.

When Nonstop behavior is disabled, failure of the AMM on the VRRP Master results in the VRRP Backup router taking control of the virtual IP addresses on the network.

The commands must be executed in VRRP context.

NOTE: Before this command is executed, the command `redundancy management nonstop-switching` should be configured. Any prerequisites required for VRRP configuration commands, such as IP routing being enabled, remain as required prerequisites.

Default: Disabled

Example

Example 52 Example of enabling nonstop switching for VRRP and then displaying the output

This example shows nonstop VRRP being enabled. The `show vrrp config` command output displays the enabled status (see bold line below.)

```
HP Switch(vlan-10-vrid-1)# nonstop
HP Switch(vlan-10-vrid-1)# show vrrp config

VRRP Global Configuration Information

VRRP Enabled      : Yes
Traps Enabled     : Yes
Virtual Routers Respond to Ping Requests : Yes
VRRP Nonstop Enabled: Yes

VRRP Virtual Router Configuration Information

Vlan ID : 10
Virtual Router ID : 1

Administrative Status [Disabled] : Enabled
Mode [Uninitialized] : Backup
Priority [100] : 150
Advertisement Interval [1] : 1
Preempt Mode [True] : True
Preempt delay time : 0
Respond to Virtual IP Ping Requests [Yes] : Yes
Primary IP Address : Lowest

IP Address      Subnet Mask
-----
10.0.202.87     255.255.0.0
```

Example nonstop routing configuration

Example 53 Example of configuring the owner routing switch

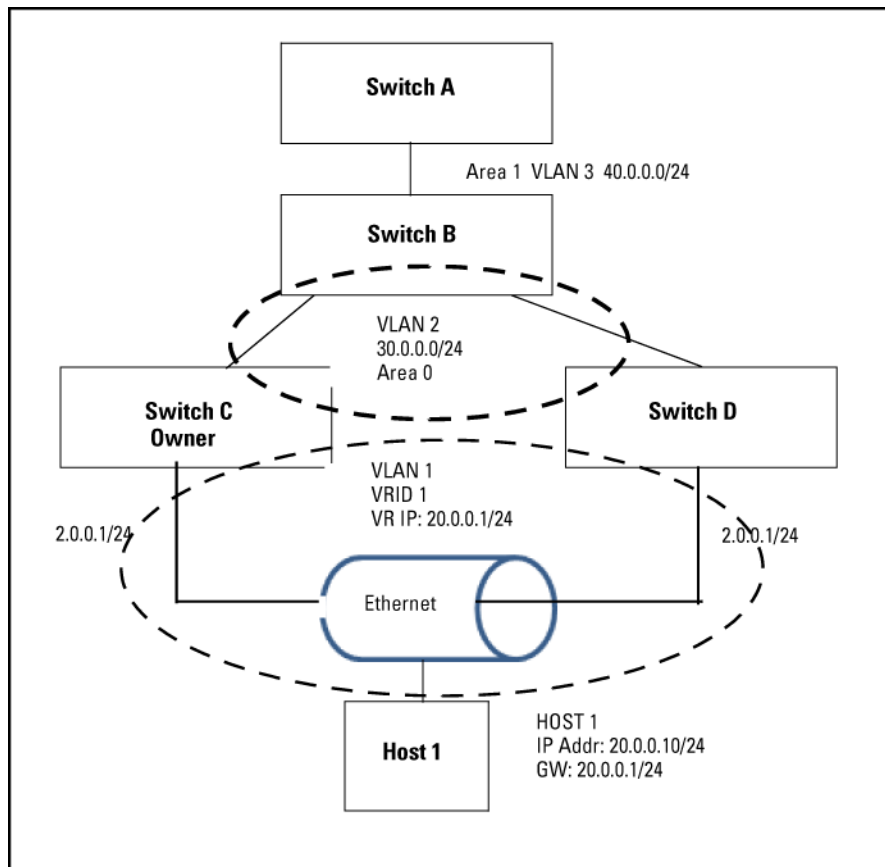
```
HP Switch C(config)# ip routing
HP Switch C(config)# router vrrp
HP Switch C(vrrp)# enable
HP Switch C(vrrp)# vlan 201
HP Switch C(vlan-201)# untag a1-a10
HP Switch C(vlan-201)# ip address 20.0.0.1/24
HP Switch C(vlan-201)# vrrp vrid 1
HP Switch C(vlan-201-vrid-1)# owner
HP Switch C(vlan-201-vrid-1)# virtual-ip-address 20.0.0.1/24
HP Switch C(vlan-201-vrid-1)# enable
```

Example 54 Example of configuring the backup routing switch

```
HP Switch D(config)# ip routing
HP Switch D(config)# router vrrp
HP Switch D(vrrp)# enable
HP Switch D(vrrp)# vlan 201
HP Switch D(vlan-201)# untag a1-a10
HP Switch D(vlan-201)# ip address 20.0.0.2/24
HP Switch D(vlan-201)# vrrp vrid 1
HP Switch D(vlan-201-vrid-1)# backup
HP Switch D(vlan-201-vrid-1)# virtual-ip-address 2.1.1.1/24
HP Switch D(vlan-201-vrid-1)# enable
```

The configuration is shown graphically in [Figure 140 \(page 260\)](#).

Figure 140 Example of nonstop routing configuration



Nonstop forwarding with RIP

On a Nonstop RIP router, the traffic does not get re-routed when the MM fails over. A request packet is sent on failover that asks for the router's peers to send routing updates to the requesting router. There is no loss of routed traffic.

Nonstop forwarding with OSPFv2 and OSPFv3

On a Nonstop OSPFv2 router, failover of a MM does not result in the OSPF v2 router being removed from the OSPFv2 domain. A restart request is sent by the Nonstop OSPFv2 router to the neighboring OSPFv2 routers, after which the graceful restart process begins. This behavior applies to OSPFv3 as well.

A graceful restart allows an OSPF routing switch to stay on the forwarding path while being restarted. The routing switch sends "grace LSAs" that notify its neighbors that it intends to perform a graceful restart. During the configurable grace period, the restarting switch's neighbors continue to announce the routing switch in their LSAs as long as the network topology remains unchanged. The neighbors run in "helper mode" while the routing switch restarts.

Graceful restart will fail under these conditions:

- There is a topology change during the graceful restart period. The helper switches exit helper mode and adjacencies are lost until the restarting switch rebuilds the adjacencies.
- The neighbor switches do not support helper mode.

For more information on OSPFv2 and OSPFv3 graceful restart, see RFC 3623 and RFC 5187.

Enabling nonstop forwarding for OSPFv2

The routing switch must be in `ospf` context when enabling Nonstop forwarding for OSPFv2. To enable Nonstop forwarding, enter this command.

Syntax

```
(ospf) # [no]
nonstop
```

Enables nonstop forwarding for OSPFv2.

The `no` version of the command disables nonstop forwarding.

The commands must be executed in `ospf` context.

Default: Disabled

Example 55 Example of enabling nonstop forwarding for OSPFv2

```
HP Switch(ospf) # nonstop
```

Configuring restart parameters for OSPFv2

Syntax

```
(ospf) # [no]
restart interval 1-1800 [strict-lsa-checking]
```

Specify the graceful restart timeout interval in seconds.

The `no` version of the command sets the restart parameters to the default values.

Default: Disabled

`interval 1-1800` The graceful restart timeout interval (grace period) in seconds.
Default: 120 seconds

`strict-lsa-checking` Used in OSPFv2 context to enable or disable strict LSA operation in a network segment for a neighboring router that is attempting a graceful restart. When enabled, this operation halts Helper mode support if a change in LSAs (topology change) is detected during the neighbor's restart period.

The `no` form of this command disables strict LSA operation.

Default: Strict LSA operation enabled

Viewing OSPFv2 nonstop forwarding information

To display the status of Nonstop forwarding information, enter the `show ip ospf general` command.

Example 56 Example of output showing status of nonstop forwarding for OSPFv2

```
HP Switch(config)# show ip ospf general

OSPF General Status

OSPF protocol      :enabled
Router ID          :10.10.10.80
.
.
.
Nonstop forwarding : Enabled
Graceful Restart Interval : 500
Graceful Restart Helper Mode : Enabled
.
.
.
```

Enabling nonstop forwarding for OSPFv3

The routing switch must be in `ospf3` context when enabling Nonstop forwarding for OSPFv3. To enable nonstop forwarding, enter this command.

Syntax

```
(ospf3)# [no]
nonstop
```

Enables nonstop forwarding for OSPFv3.

The `no` version of the command disables nonstop forwarding.

The commands must be executed in `ospf3` context.

Default: Disabled

Example 57 Example of enabling nonstop forwarding for OSPFv3

```
HP Switch(ospf3)# nonstop
```

Configuring restart parameters for OSPFv3

Syntax

```
(ospf3)# [no]
restart interval 1-1800 [strict-lsa-checking]
```

Specify the graceful restart timeout interval in seconds.

The `no` version of the command sets the restart parameters to the default values. Default: Disabled

`interval 1-1800` The graceful restart timeout interval (grace period) in seconds.
Default: 120 seconds

`strict-lsa-checking` Used in OSPFv3 context to enable or disable strict LSA operation in a network segment for a neighboring router that is attempting a graceful restart. When enabled, this operation halts Helper mode support if a change in LSAs (topology change) is detected during the neighbor's restart period.

The `no` form of this command disables strict LSA operation.

Default: Strict LSA operation enabled

Viewing OSPFv3 nonstop forwarding information

To display the status of Nonstop forwarding information, enter the `show ipv6 ospf3 general` command.

Example 58 Example of output showing status of nonstop forwarding for OSPFv3

```
HP Switch(config)# show ipv6 ospf3 general
```

```
OSPFv3 General Status
```

```
OSPFv3 protocol :enabled
Router ID       :10.10.10.80
.
.
.
Nonstop forwarding : Enabled
Graceful Restart Interval : 500
Graceful Restart Helper Mode : Enabled
.
.
.
```

Hotswapping management modules

Management module switchover

Events that cause a switchover

There are a number of events that can cause the active management module to switchover to the standby management module when management module redundancy is enabled:

- The active management module crashes
- The standby management module does not receive a heartbeat from the active management module
- The `redundancy switchover` command is executed
- The **MM Reset** button on the active management module is pressed
- The **MM Shutdown** button on the active management module is pressed
- The `boot` or `boot active` command is executed
- The `reload` command is executed
- There is a hardware failure on the active management module

In all of these cases, the standby management module takes control and performs the actual switchover. The reason for the switchover is entered in log messages on the newly active management module and to any configured Syslog servers.

What happens when switchover occurs

When a switchover occurs, the features that support nonstop switching continue to operate in an uninterrupted manner. See [“Nonstop switching features” \(page 276\)](#) for a list of the supported features.

The features that do not support nonstop switching perform as if the switch had just finished booting; however, no actual boot time occurs.

NOTE:

When meshing configuration changes are made on a redundant management system, you must execute `write mem` and then the `boot system` command to boot *both* management modules for the changes to be activated.

Meshing is not supported by nonstop switching.

NOTE:

If the switch is a querier and a failover occurs, the querier continues to be the same on the standby management module; no new querier election process occurs on the standby management module.

When switchover will not occur

There are some events for which a switchover is not triggered:

- When a `boot system` command is executed
- When the **Clear** button on the System Support module is pressed
- When management module redundancy is disabled, unless there is a hardware failure and the system is rebooted.

When a management module crashes while the other management module is rebooting

If the uncommon situation occurs where the active management module (MM1) is trying to reboot and the standby management module (MM2) also crashes, the switch attempts to recover from the crash and eventually the standby management module becomes the active management module if it passes self-test. However, traffic can be disrupted for as long as five minutes before the newly active management module (MM2) has finished rebooting.

About hotswapping out the active management module

You can hotswap out the active management module and have switch operations taken over by the standby management module by following the correct shutdown procedure on the active module using the **MM Shutdown** button. When the **MM Shutdown** button is pressed, any file synchronization in progress completes before the shutdown begins, and then a graceful shutdown of that management module occurs.

For information on using the **MM Shutdown** button, see “Hotswapping out the active management module” (page 245).

When the standby module is not available

If you have disabled management module redundancy with the `no redundancy management-module` command, or the standby module failed selftest, the **Dwn** LED does not turn green to indicate it is OK to hotswap out the active management module.

NOTE: If you remove the active management module without pressing the **MM Shutdown** button, any files that may have been in the process of synchronizing will not finish synchronizing to the standby module and all file transfer is aborted.

Hotswapping in a management module

If another management module is hotswapped in while there is an active management module booted up, the newly hotswapped management module becomes the standby module.

No negotiating is needed as to which module becomes the active management module, because there is already a functioning active management module. However, the following conditions must be met to determine if the hotswapped module can become a standby management module:

- The hotswapped module must pass selftest
- Management module redundancy is not administratively disabled (using the `no redundancy management-module` command.) If the active management module's config file has redundancy administratively disabled, the hotswapped management module goes into "offline" mode.

In nonstop switching mode—The active management module's files and features are synced with the standby management module. Heartbeats are sent back and forth, and the standby management module is ready to quickly take over in the event of a switchover or a failure on the active management module.

In warm-standby mode—The standby management module partially boots up and heartbeats are sent back and forth with the active management module.

Software version mismatch between active and hotswapped module

If the software version in the hotswapped module does not match the software version in the active module, the following occurs:

1. The active module sends the primary and secondary images in flash to the hotswapped module.
2. The module that was hotswapped in then reboots if necessary to primary or secondary flash, whichever matches (if it does not already match.)
3. After the hotswapped management module finishes booting, it is sent the config and other critical files from the active management module.
4. The hotswapped management module goes into standby mode and is ready to take over in case of a switchover.

NOTE: After the `boot standby` command is executed, if the software versions on the active management module and the standby management module are not compatible, the standby module does not sync with the active management module. The standby module then enters warm-standby redundancy mode.

Other software version mismatch conditions

The following steps describe the behavior that may when a new software image is installed in secondary flash of the AMM and a `redundancy switchover` command is executed.

1. A new software image, K.15.04.0002 containing ROM upgrade K.15.12 is installed in secondary flash of the AMM/MM1.
2. The AMM/MM1 automatically syncs the images to the secondary flash in the SMM/MM2. Now both AMM/MM1 and SMM/MM2 have identical software and ROM in secondary flash.
3. The SMM/MM2 is booted from secondary. It boots into the new K.15.04.0002 software version. The new ROM is applied and the SMM/MM2 reboots.
4. After the SMM/MM2 finishes rebooting, it reconnects to the AMM/MM1 and prepares to take the standby role by rebooting.
5. However, the AMM/MM1 is running software version K.15.03.0008 in its primary flash, and the SMM/MM2 is running software version K.15.04.0002 in its secondary flash, so the SMM/MM2 pauses its reboot because of the software mismatch.
6. If a `redundancy switchover` command is executed, the AMM/MM1 will give control to the SMM/MM2, which can then finish booting and become the new AMM/MM2. This is the warm-start behavior.

7. The SMM/MM1 (former AMM/MM1) reboots, but unless the reboot is executed from secondary flash, it reboots into primary flash, which contains the older software version K.15.03.0008 with no ROM upgrade.
8. If the SMM/MM1 is forced to boot from secondary before executing the `redundancy switchover` command, it will boot into the new K.15.04.0002 software and upgrade the ROM. After the reboot that occurs with the ROM upgrade, the SMM/MM1 connects to the new AMM/MM2 and takes the standby role.

About downloading a new software version

File synchronization after downloading

After downloading a new software version to either the primary or secondary flash of the active management module, the software version is immediately copied to the corresponding flash (primary or secondary) of the standby module, unless the standby module failed selftest or redundancy was disabled with the `no redundancy management-module` command.

The configuration files, including which configuration file to use for that flash image, are synchronized. For example, if the active management module is using `config1`, the standby module is also synchronized to use `config1`.

Table 21 Example of upgrading software version K.15.01.0003 to version K.15.01.0004

	Newer code to secondary flash		New code to primary flash	
	Active MM	Standby MM	Active MM	Standby MM
Software version downloaded to Primary flash image	K.15.01.0003	K.15.01.0003	K.15.01.0004	K.15.01.0004
Software version downloaded to Secondary flash image	K.15.01.0004	K.15.01.0004	K.15.01.0003	K.15.01.0003

NOTE: For information about testing new software versions, see [“Setting the default flash for boot \(CLI\)” \(page 251\)](#).

After installing the new software to the active management module, wait a few minutes, and then verify that the standby management module has been synchronized with the new software as well (use the `show flash` command.) If the default flash for boot is set correctly, you can start the standby management module on the new software by executing the `boot standby` command. This does not interrupt current switch operations yet. After the standby management module has rebooted and is ready for takeover in standby mode (you can verify this using the `show redundancy` command—see [syntax \(page 237\)](#)), you can now switch over to the management module running the newer software with this command:

```
HP Switch# redundancy switchover
```

This causes a switchover to the management module that received the new software version, which becomes the active management module. This method incurs the least amount of network downtime for booting. If downtime is not an issue, use the `boot system` command. Both management modules are then running the new software version.

Potential software version mismatches after downloading

When a new software version is downloaded to the active management module, it is immediately copied to the corresponding flash (primary or secondary) in the standby management module, unless redundancy has been disabled. If the standby management module is rebooted, it will be running a different software version than the active management module. You can direct the standby

module to boot from the non-corresponding flash image that has a different software version during the actual reboot process of the standby module when the prompt to select the **Boot Profile** appears, as shown in [Figure 141 \(page 267\)](#).

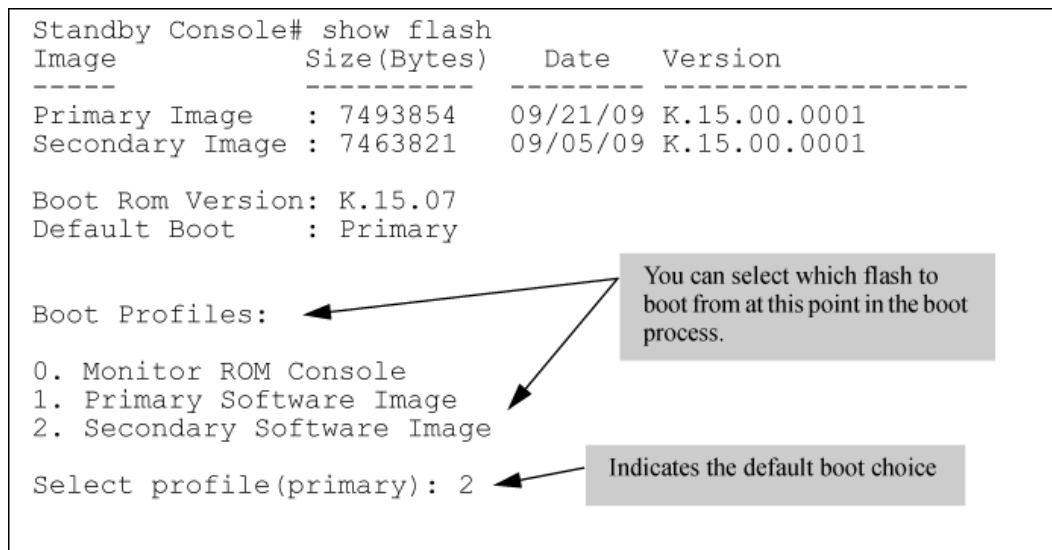
Figure 141 Booting the standby management module to secondary flash

```
Standby Console# show flash
Image          Size(Bytes)   Date   Version
-----
Primary Image  : 7493854   09/21/09 K.15.00.0001
Secondary Image: 7463821   09/05/09 K.15.00.0001

Boot Rom Version: K.15.07
Default Boot    : Primary

Boot Profiles:
0. Monitor ROM Console
1. Primary Software Image
2. Secondary Software Image

Select profile(primary): 2
```



CAUTION: If you have booted one module out of primary flash and one module out of secondary flash, and the secondary flash is running a prior software version because the latest version was never copied over from the primary flash, you will have a software version mismatch. The configuration file may not work with that software version. For more information, see [“Software version mismatch between active and hotswapped module” \(page 265\)](#).

The standby module enters warm-standby redundancy mode and boots to a certain point, syncs basic files such as the config and security files, and finishes booting only if the active management module fails or you choose to change which module is the active module..

Additionally, if a switchover occurs, or if you reboot to make the standby module become the active module, any configuration file changes made may not work on the active module if it has a different software version from the standby module.

When you enter the `show redundancy` command and a software version mismatch exists, a warning message is displayed, as shown at the bottom of [Figure 142 \(page 268\)](#).

Figure 142 Example of a software version mismatch between the active and standby modules

```
HP Switch(config)# show version
Management Module 1: Active
Image stamp:      /sw/code/build/btm(t2g)
                  Mar 15 2007 12:28:32
                  K.15.01.0001
                  64
Boot Image:      Primary

Management Module 2: Standby
Image stamp:      /sw/code/build/btm(t2g)
                  Mar 21 2007 14:24:38
                  K.15.01.0002
                  789
Boot Image:      Secondary

HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Warm-standby redundancy enabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 0
Last Failover :

Slot Module Description                               Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200z1  Active   K.15.01.0001 Primary
MM2  HP Switch J9092A Management Module 8200z1  Standby  K.15.01.0002 Secondary

FM1  HP Switch J9093A Fabric Module 8200z1      Enabled
FM2  HP Switch J9093A Fabric Module 8200z1      Enabled

Warning: Standby module is running a different software version and may be using
a different configuration file. Configuration changes on active management
module may not take effect on a failover.
```

Downloading a software version serially if the management module is corrupted

If the software version on a management module becomes corrupted, you may need to do a serial download to restore the affected module. The non-corrupted management module becomes the active module. You can then use the serial port on the corrupted management module to download a new software version. When the corrupted module is rebooted, the software version in the corrupted module is immediately overwritten by the software version in the active management module. Both management modules should now operate on the same software version.

About turning off redundant management

Disable management module redundancy with two modules present

To troubleshoot a suspect management module, you may want to operate the switch with redundant management disabled by entering this command:

```
HP Switch(config)# no redundancy management-module
```

After executing this command, the second management module will not boot into standby mode—it is offline and no longer receives configuration file changes from the active module. The active management module updates its config file with the information that redundancy is disabled.

NOTE: Even if redundancy has been disabled, the specified management module becomes the active management module at the next system boot if you use the `redundancy active-management` command. You are warned that you may not be using current configurations. See “[Setting the active management module for next boot \(CLI\)](#)” (page 243).

The second management module is enabled as the active management module in the event of a hardware failure of the first management module.

Figure 143 (page 269) shows that redundant management was disabled.

Figure 143 Results of disabling redundancy

```
HP Switch(config)# no redundancy management-module
The other management module may reboot and it will no longer be used for system
redundancy except in the case of a hardware failure of the active
management module. Do you want to continue [y/n]? y

HP Switch(config)# show redundancy

Settings
-----
Mgmt Redundancy : Nonstop switching disabled
Rapid Switchover Stale Timer : 0

Statistics
-----
Failovers      : 0
Last Failover :

Slot Module Description                               Status  SW Version  Boot Image
-----
MM1  HP Switch J9092A Management Module 8200z1  Offline   K.15.01.000x Primary
MM2  HP Switch J9092A Management Module 8200z1  Active    K.15.01.000x Primary

FM1  HP Switch J9093A Fabric Module 8200z1   Enabled
FM2  HP Switch J9093A Fabric Module 8200z1   Enabled
s
```

Disable management module redundancy with only one module present

If you disable redundancy when there is only one management module in the switch, and then you insert a second management module, the second module never goes into standby mode. You must re-enable redundant management using this command:

```
HP Switch(config)# redundancy management-module
```

The currently active module remains active on boot (assuming no selftest failure) unless you make the newly inserted management module active using this command:

```
HP Switch(config)# redundancy active-management standby
```

The standby management module becomes the active management module.

Active management module commands

Show modules

The `show modules` command displays information about all the modules in the switch, as well as additional component information for the following:

- System Support Modules (SSM)—identification, including serial number
- Mini-GBICS—a list of installed mini-GBICs displaying the type, "J" number, and serial number (when available)

CLI commands affected by redundant management

Several existing commands have changes related to redundant management.

boot command

In redundant management systems, the `boot` or `boot active` command causes a switchover to the standby management module as long as the standby module is in standby mode. This message displays:

```
This management module will now reboot and will become the
standby module! You will need to use the other management
module's console interface. Do you want to continue [y/n]?
```

If you select **y**, switchover is initiated by the standby management module, which becomes the active management module after boot completes.

If the standby module is not in standby mode (for example, it is in failed mode or offline mode), switchover to the standby module does not occur. The system is rebooted and this message displays:

```
The other management module is not in standby mode and this
command will not cause a switchover, but will reboot the
system, do you want to continue [y/n]?
```

If the other management module is not present in the switch, the system simply reboots.

The `boot` command has these options.

Command	Action
<code>boot cr</code>	Reboots the active management module from the flash image that is specified for the default boot. This can be changed with the <code>boot set-default flash</code> command. You can select which image to boot from during the boot process itself. (See Figure 144 (page 271) .) The switch will switchover to the standby management module. NOTE: This is changed from always booting from primary flash. You are prompted with a message, which indicates the flash being booted from.
<code>boot active</code>	Boots the active management module. The switch starts to boot from the default flash image. You can select which image to boot from during the boot process itself. (See Figure 144 (page 271) .) The switch will switchover to the standby management module. If a second management module is not present in the switch, the system is rebooted.
<code>boot standby</code>	Boots the standby management module. The switch does not switchover. If the standby module is not present, this message displays: "The other management module is not present."
<code>boot system [flash [primary secondary]]</code>	Boots both the active and standby management modules. You can specify the flash image to boot from.
<code>boot set-default flash primary secondary</code>	Sets the default flash for the next boot to primary or secondary. You see this message: "This command changes the location of the default boot. This command will change the default flash image to boot from flash chosen>. Hereafter, 'reload' and 'boot' commands will boot from flash chosen>. Do you want to continue [y/n]?"

You can select a **boot profile** during the reboot process, as shown in [Figure 144 \(page 271\)](#). If you make no selection, the boot defaults to the image displayed as the default choice (shown in parentheses.)

Figure 144 The management module rebooting, showing boot profiles to select

```
Boot Profiles:
0. Monitor ROM Console
1. Primary Software Image
2. Secondary Software Image

Select profile(primary): 2

Booting Secondary Software Image...
```

An example of the boot command with the **default flash** set to **secondary** is shown in [Figure 145](#) (page 271).

Figure 145 Showing boot command with default flash set to secondary

```
HP Switch(config)# boot set-default flash secondary
This command changes the location of the default boot. This command will
change the default flash image to boot from secondary image. Hereafter,
'reload' and 'boot' commands will boot from secondary image. Do you want
to continue [y/n]? y

HP Switch(config)# show flash
Image          Size(Bytes)  Date   Version
-----
Primary Image   : 7476770   11/01/09 K.15.01.0001
Secondary Image : 7476770   11/01/09 K.15.01.0001

Boot Rom Version: K.15.07
Default Boot    : Secondary

HP Switch(config)# boot
This management module will now reboot from secondary and will become
the standby module! You will need to use the other management module's
console interface. Do you want to continue [y/n]?
```

CAUTION: For a given reboot, the switch automatically reboots from the `startup-config` file assigned to the flash (primary or secondary) being used for the current reboot. The `startup-default` command can be used to set a boot configuration policy. This means that both the flash image and one of the three configuration files can be specified as the default boot policy.

Boot and reload commands with OSPFv2 or OSPFv3 enabled

It is now possible to gracefully shut down OSPFv2 or OSPFv3 routing on HP switches without losing packets that are in transit. OSPF neighbors are informed that the router should not be used for forwarding traffic, which allows for maintenance on the switch without interrupting traffic in the network. There is no effect on the saved switch configuration.

Prior to a switch shutdown, the CLI/SNMP `reload` command or the CLI `boot` command is executed to initiate the sending of OSPF "empty Hello list" messages on the interfaces that are part of the OSPF routing configuration. After a small delay (approximately 2 seconds) that allows the messages to be transmitted on all applicable interfaces the `boot` or `reload` command continues.

Modules operating in nonstop mode

When a switch is in standalone mode and OSPF routing is enabled, the "empty Hello list" is transmitted whenever the `boot` or `reload` command is executed.

When the switch is operating in nonstop switching mode (redundant), and a single module is being reloaded or booted, the standby module notifies neighboring switches of the management module failover. If the failover fails, the "empty Hello list" is transmitted before the switch is rebooted.

When a switch is operating with multiple management modules in warm standby mode, the "empty Hello list" is sent when a `reload` or `boot` command is executed. The standby management module sends out OSPF Hello packets after becoming the active management module.

Additional commands affected by redundant management

The other existing commands operate with redundant management as shown below.

Command	Action
<code>auto-tftp</code>	If a new image is downloaded using <code>auto-tftp</code> , the active management module downloads the new software version to both the active and standby modules. Rebooting after the <code>auto-tftp</code> completes reboots the entire system.
<code>banner</code>	The banner will not be seen on the standby module, only the active module.
<code>chassislocate</code>	If the management module performs a switchover, the LED does not remain lit.
<code>clear</code>	The <code>clear crypto</code> command causes public keys to be deleted from both modules when the second module is in standby mode.
<code>console</code>	Console settings, such as mode, flow-control, and baud-rate, are the same on both management modules. There cannot be individual settings for each management module.
<code>copy</code>	Files are automatically sync'd from the active management module to the standby management module. When no parameter is specified with the <code>copy crash-data</code> or <code>copy crash-log</code> command, files from all modules (management and interface) are concatenated. See "Copying crash file information to another file (CLI)" (page 254) . NOTE: If redundancy is disabled or the standby module failed selftest, the <code>copy</code> command affects only the active management module.
<code>copy core-dump [mm standby flash xmodem usb filename]</code>	The <code>copy core-dump standby flash</code> command copies the standby management module's coredump to the active management module's flash. The destination file is fixed as <code>dumpM1.cor</code> or <code>dumpM2.cor</code> , depending on which module is the standby management module. The <code>copy core-dump [mm standby flash xmodem usb filename]</code> command copies the core file of the active management module or the standby management module to a USB flash drive or to an xmodem host.
<code>core-dump management-module</code>	Enables or disables a core dump on a management module.
<code>crypto</code>	Authentication files for ssh or the https server are copied to the standby management module. The <code>clear crypto</code> command deletes the public keys from both modules when the second module is in standby mode.
<code>erase flash</code>	Erases the software version on the active and standby modules. If redundancy has been disabled, or if the standby module has not passed selftest, the flash is not erased on the standby module.
<code>erase config</code>	Erases the config file on the active and standby modules. If redundancy has been disabled, or if the standby module has not passed selftest, the config file is not erased on the standby module.
<code>erase startup-config</code>	Affects both modules if the second module is in standby mode. If redundancy has been disabled, or if the standby module has not passed selftest, the <code>startup-config</code> file is not erased on the standby module.

Command	Action
fastboot	When fastboot is enabled, this information is saved to the standby management module when the config files are sync'd. The fastboot value is used during the next boot on both modules.
front-panel-security factory-reset password-clear password-recovery	This command and its options affect only the active management module. See the section "Front-Panel Button Functions" in the <i>Access Security Guide</i> for more information about resetting the switch.
kill	Does not affect the console on the standby module.
log	Log messages from a formerly active management module are available on the current active management module after a switchover.
password (set or clear)	Affects only the active management module until a switchover occurs, at which time it affects the new active module.
startup-default	Affects both modules. The config file is immediately sent to the standby module and also becomes the default on that module when the next boot occurs.
update	Affects only the active module. The standby may become the active module when the updated active module is booted.
write	A write memory updates the config file in flash on the active module. The file is then sync'd to the standby module.

Using the WebAgent for redundant management

The WebAgent can be used to display information about the active and standby management modules.

Online Help is available for the WebAgent, which you can open by clicking on the question mark (?) in the upper right corner of any of the WebAgent screens. An example redundancy screen is shown in [Figure 146 \(page 273\)](#).

To access the redundancy information in the WebAgent:

1. In the WebAgent navigation panel, click System.
2. Click Redundancy. The following screen displays.

Figure 146 Example of redundancy screen in the WebAgent

The screenshot shows the WebAgent interface for System > Redundancy. At the top right, there is a [Reboot] button and a question mark icon. The main content is divided into three sections:

- Management Module Status:** A table with columns Slot, Status, Description, Software Version, and Boot Image.

Slot	Status	Description	Software Version	Boot Image
MM1	Active	ProCurve J9092A Management Module 8200zl	K.15.01.0000x	Secondary
MM2	Offline	ProCurve J9092A Management Module 8200zl	K.15.01.0000x	Secondary
- Redundancy:** A section with the following values:
 - Redundancy: Enabled
 - Failovers: 0
 - Last Failover: 0 Seconds
- Fabric Module Status:** A table with columns Slot, Status, and Description.

Slot	Status	Description
FM1	Enabled	ProCurve J9093A, Fabric Module 8200zl
FM2	Enabled	ProCurve J9093A, Fabric Module 8200zl

Notes on how the active module is determined

Both management modules run selftest routines to determine which module becomes the active management module and which becomes the standby management module. The module that was last active in the chassis is given precedence and becomes the "active" module. This module is the one that is booted going forward. If a module fails selftest and is unable to communicate with the other module, it does not take control as the management module. The other management module takes control and becomes the active module.

If both modules fail selftest, the fault LED flashes and neither module is operational.

NOTE: You are not allowed to switchover to a management module that is not in standby mode. The module must have passed selftest and be in standby mode.

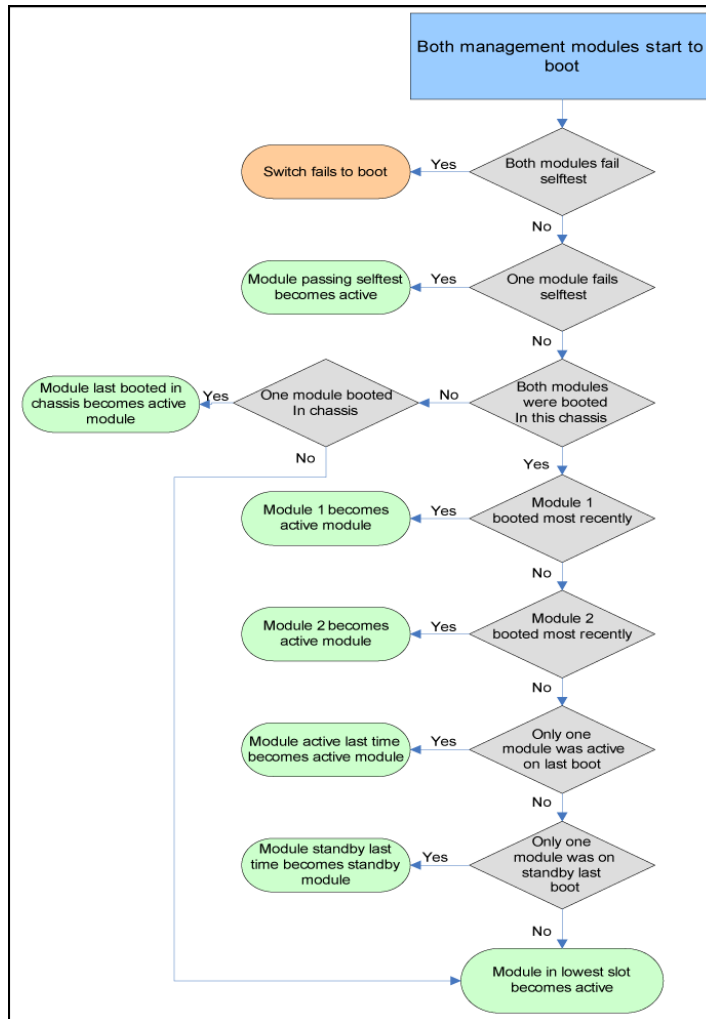
The entire boot decision process works as follows:

1. If there is only one management module, that is the active management module.
2. If one module is already booted and operational, a newly inserted module or the other management module booting always becomes the standby module. The standby module does not become active unless a switchover occurs.
3. If there are two management modules and one fails selftest, the one that passes selftest becomes the active management module.
4. If only one of two modules was ever booted in the chassis, that module is given precedence.
5. The module that was active on the last boot becomes the active management module. This guarantees that the active module has the latest configuration data.
6. If both management modules have previously booted in this chassis and were "active" the last time booted, the module that booted most recently becomes the active management module.
7. If none of the above conditions are applicable, the module in the lowest slot becomes the active management module.

To see a diagram of this process, see [“Active module decision flow chart at boot” \(page 275\)](#).

Diagram of the decision process

Figure 147 Active module decision flow chart at boot



CLI commands and syncing

The following CLI commands can be executed during initial syncing between the active management module and the standby management module, which occurs when the standby module is inserted or after a reboot of the system. All other CLI commands will not be executed until after the initial syncing completes.

During initial syncing, no SNMP set requests are executed, except the SNMP request for ping.

Operator commands		
dir	menu	traceroute6
enable	ping	dbgstack
exit	ping6	wireless-services
link-test	show	services
logout	traceroute	

Manager commands		
boot system	copy running-config	page
boot active	copy startup-config	print
boot standby	copy event-log	redo
configure	copy core-dump	reload
copy command-output	recopy	repeat
copy config tftp	display	task-monitor
copy config xmodem	end	telnet
copy crash-data	getMIB	terminal
copy crash-log	kill	walkMIB
copy flash tftp	licenses	write-terminal
copy flash xmodem	log	redundancy

Management module redundancy features

Nonstop switching features

Nonstop switching features are synced at initialization of the standby management module.

802.1X and Web/MAC authentication	Spanning Tree (MSTP)
MAC Lockout/Lockdown	GVRP
ACLs/Qos Policies	Loop Protection
Power over Ethernet	LACP
Port Security	Syslog
DHCP Snooping	UDLD
Dynamic ARP Protection	Virus Throttling
Dynamic IP Lockdown	LLDP

Unsupported zl modules

ZL modules/controllers that do not support the nonstop switching feature include the following:

- HP ONE Services zl Module (J9289A)
- HP Threat Management Services zl Module (J9155A)
- HP Threat Management Services zl Module with 1-year IDS/IPS subscription (J9156A)
- HP Wireless Edge Services zl Module (J9051A) and Redundant Wireless Services zl Module (J9052A)
- HP MSM765zl Mobility Controller (J9370A)

During a nonstop switching failover, unsupported modules will not failover seamlessly to the standby module. A nonstop switching failover causes a forced reboot on these modules. After rebooting, these modules then sync with the newly active management module and begin operation again. Module traffic is disconnected until the module completes the reboot process.

Hot swapping of management modules

Use the MM Shutdown button on the front of the management module before removal. The Shutdown button ensures that the management module is shut down properly. If nonstop switching is enabled, using the Shutdown button prior to removal ensures failover to the standby module will be successful.

Rapid routing switchover and stale timer

With K.15.01.0031, nonstop switching supports only Layer 2 functions on the switch. During a failover, traffic routed through the switch at Layer 3 will see an interruption. When a failover from active to standby occurs, the routing table is "frozen." All routes that existed at the time of the failover are marked as "stale." While dynamic routing protocols running at the time may act as if the routing protocol has been restarted and rebuilds the table, the switch on which the failover occurred continues to rout traffic using the 'stale routes.'

The "stale timer" begins counting when the switchover occurs. When the "stale timer" expires, any routes that are still marked as stale are purged from the routing table. Because of the nature of rapid routing switchover, if there are multiple simultaneous failures, network loops could occur or traffic could flow through unpredictable paths.

Use caution if setting the "rapid-switchover" timer higher than the default. To disable "rapid routing switchover" and to ensure that all routing is based on the most current routing protocol information, set the "rapid-switchover" timer to 0.

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

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-MAC 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.



9 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/support.

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.

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, Xmodem, and USB. 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 “[About downloading switch software](#)” (page 303).

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 www.procurve.com.)
- 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 flash using TFTP (CLI)

Syntax

```
copy tftp flash ip-address remote-file [ primary | secondary ] [oobm]
```

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.

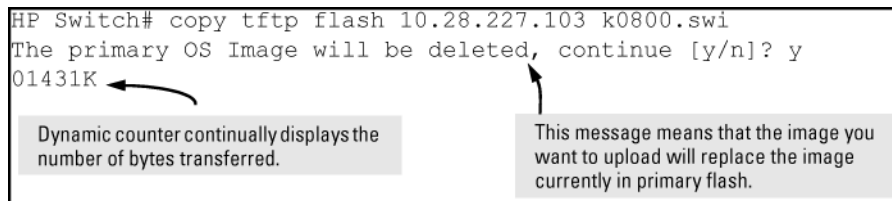
Example

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:

Figure 148 Example of 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
01431K
```



Dynamic counter continually displays the number of bytes transferred.

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

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 the *Basic Operation Guide*.

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 the *Basic Operation Guide*.

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 [ listen oobm | data | both ] ]
```

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.

For switches that have a separate out-of-band management port, the `listen` parameter in a **server** configuration allows you to specify whether transfers take place through the out-of-band management (oobm) interface, the data interface, or both. For more information on OOBM, see [“Network Out-of-Band Management \(OOBM\) for the 6600 Switch” \(page 484\)](#).

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 \(CLI\)" \(page 283\)](#) 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 Configuration Guide](#).

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 \(CLI\)" \(page 282\)](#).)

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.

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 149 \(page 284\)](#). (The term "OS" or "operating system" refers to the switch software):

Figure 149 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 150 \(page 284\)](#) appears:

Figure 150 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 the *Basic Operation Guide*.

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

For troubleshooting information on download failures, see [“Troubleshooting TFTP download failures” \(page 303\)](#).

Enabling SCP and SFTP

For more information about secure copy and SFTP, see [“Using SCP and SFTP” \(page 304\)](#).

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 the *Access Security Guide*. 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” \(page 305\)](#).

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 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 the *Basic Operation Guide*.

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 the *Basic Operation Guide*.

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 switch software using USB (CLI)

This procedure assumes that:

- A software version for the switch has been stored on a USB flash drive. (The latest software file is typically available from the HP Switch Networking website at www.hp.com.)
- The USB device has been plugged into the switch's USB port.

Before you use the procedure:

- Determine the name of the software file stored on the USB flash drive (for example, `k0800.swi`.)
- Decide whether the image will be installed in the primary or secondary flash. For more on primary/secondary flash memory and related boot commands, see the *Basic Operation Guide*.

Syntax

```
copy usb flash filename [ primary | secondary ]
```

This command automatically downloads a switch software file to primary or secondary flash. If you do not specify the flash destination, the USB download defaults to primary flash.

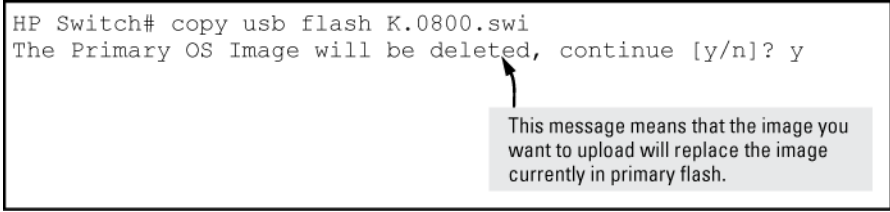
Example

To copy a switch software file named `k0800.swi` from a USB device to primary flash:

1. Execute `copy` as shown below:

Figure 151 The command to copy switch software from USB

```
HP Switch# copy usb flash K.0800.swi
The Primary OS Image will be deleted, continue [y/n]? y
```



When the switch finishes copying the software file from the USB device, it displays this progress message:

Validating and Writing System Software to the Filesystem....

2. When the copy finishes, you must reboot the switch to implement the newly loaded software. 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 the *Basic Operation Guide*.

3. To confirm that the software downloaded correctly, execute `show system` and check the **Firmware revision** line.

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.

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 ] [oobm]
```

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.

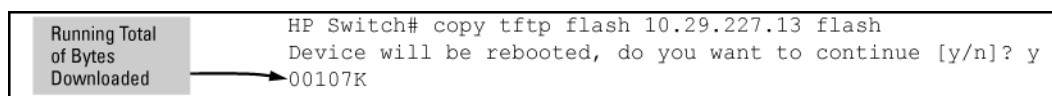
For switches that have a separate OOBM port, the `oobm` parameter specifies that the TFTP traffic must come in through the OOBM interface. If this parameter is not specified, the TFTP traffic comes in through the data interface. The `oobm` parameter is not available on switches that do not have a separate OOBM port. For more information on OOBM, see [“Network Out-of-Band Management \(OOBM\) for the 6600 Switch” \(page 484\)](#) in this guide.

If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

Example

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:

Figure 152 Switch-to-switch, from primary in source to either flash in destination



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 ] [oobm]
```

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.

For switches that have a separate out-of-band management port, the `oobm` parameter specifies that the TFTP traffic must come in through the out-of-band management interface. If this parameter is not specified, the TFTP traffic comes in through the data interface. The `oobm` parameter is not available on switches that do not have a separate out-of-band management port. For more information on out-of-band management, see “[Network Out-of-Band Management \(OOBM\) for the 6600 Switch](#)” (page 484) in this guide.

Example

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 59 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
```

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.

Copying software images

NOTE: For details on how switch memory operates, including primary and secondary flash, see the *Basic Operation Guide*.

TFTP: Copying a software image to a remote host (CLI)

Syntax

```
copy flash tftp ip-addr filename [oobm]
```

Copies the primary flash image to a TFTP server.

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the OOBM interface. If this parameter is not specified, the transfer is through the data interface.

The `oobm` parameter is not available on switches that do not have a separate OOBM port. For more information on OOBM, see [“Network Out-of-Band Management \(OOBM\) for the 6600 Switch” \(page 484\)](#) in this guide.

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.

USB: Copying a software image to a USB device (CLI)

To use this method, a USB flash memory device must be connected to the switch's USB port.

Syntax

```
copy flash usb filename
```

Uses the USB port to copy the primary flash image from the switch to a USB flash memory device.

Example

To copy the primary image to a USB flash drive:

1. Insert a USB device into the switch's USB port.
2. Execute the following command:

```
HP Switch# copy flash usb k0800.swi
```

where `k0800.swi` is the name given to the primary flash image that is copied from the switch to the USB device.

Transferring switch configurations

For more information, see [“About transferring switch configurations” \(page 310\)](#).

TFTP: Copying a configuration file to a remote host (CLI)

Syntax

```
copy startup-config | running-config tftp ip-addr remote-file [ pc  
| unix ] [oobm]
```

```
copy config filename tftp ip-addr remote-file [ pc | unix ] [oobm]
```

This command can copy a designated config file in the switch to a TFTP server. For more information, see the *Basic Operation Guide*.

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the OOBM interface. If this parameter is not specified, the transfer is through the data interface.

The `oobm` parameter is not available on switches that do not have a separate OOBM port. For more information, see [“Network Out-of-Band Management \(OOBM\) for the 6600 Switch” \(page 484\)](#).

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-addr remote-file [ pc  
| unix ] [oobm]
```

```
copy tftp config filename ip-addr remote-file [ pc | unix ] [oobm]
```

This command can copy a configuration from a remote host to a designated config file in the switch. For more information, see the *Basic Operation Guide*.

For more information on flash image use, see the *Basic Operation Guide*.

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the OOBM interface. If this parameter is not specified, the transfer is through the data interface.

The `oobm` parameter is not available on switches that do not have a separate OOBM port. For more information, see [“Network Out-of-Band Management \(OOBM\) for the 6600 Switch” \(page 484\)](#).

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 [oobm]
```

Copies a customized command file to the switch (see [Example 60 \(page 292\)](#).)

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the out-of-band management interface. If this parameter is not specified, the transfer is through the data interface. The `oobm` parameter is not available on switches that do not have a separate OOBM port. For more information, see ["Network Out-of-Band Management \(OOBM\) for the 6600 Switch"](#) (page 484).

Example

Example 60 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.

Example

Example 61 The `show tech custom` command

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.)

```
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 the *Basic Operation Guide*.

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 the *Basic Operation Guide*.

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 the *Basic Operation Guide*.

Syntax

```
reload
```

Reboots from the flash image currently in use.

(For more on these commands, see the *Basic Operation Guide*.)

USB: Copying a configuration file to a USB device (CLI)

To use this method, a USB flash memory device must be connected to the switch's USB port.

Syntax

```
copy startup-config usb filename
copy running-config usb filename
```

Uses the USB port to copy a designated configuration file from the switch to a USB flash memory device. For more information, see the *Basic Operation Guide*.

Example

To copy the startup configuration file to a USB flash drive:

1. Insert a USB device into the switch's USB port.
2. Execute the following command:

```
HP Switch# copy startup-config usb HP Switch-config
```

where `HP Switch-config` is the name given to the configuration file that is copied from the switch to the USB device.

USB: Copying a configuration file from a USB device (CLI)

To use this method, the switch must be connected via the USB port to a USB flash drive on which is stored the configuration file you want to copy. To execute the command, you will need to know the name of the file to copy.

Syntax

```
copy usb startup-config filename
```

Copies a configuration file from a USB device to the startup configuration file on the switch.

Example

To copy a configuration file from a USB device to the switch:

1. Insert a USB device into the switch's USB port.
2. Execute the following command:

```
HP Switch# copy usb startup-config HP Switch-config
```

where `HP Switch-config` is the name of the file to copy.

3. At the prompt, press **[Enter]** to reboot the switch and implement the newly downloaded software.

Transferring ACL command files

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

Syntax

```
copy tftp command-file ip-addr filename.txt unix | pc [oobm]
```

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

<code>ip-addr</code>	The IP address of a TFTP server available to the switch
<code>filename.txt</code>	A text file containing ACL commands and stored in the TFTP directory of the server identified by <code>ip-addr</code>

[unix pc]	The type of workstation used for serial, Telnet, or SSH access to the switch CLI
[oobm]	For switches that have a separate out-of-band management port, specifies that the transfer will be through the out-of-band management interface. (Default is transfer through the data interface.)

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 the *Access Security Guide*.)
- 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 [Figure 153 \(page 295\)](#)), and continues to implement the remaining ACL commands in the file.

Figure 153 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.
 5. exit
Switch(config)# show running
Running configuration:

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

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

Manually executing **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 the *Access Security Guide*.)
- Adds to an existing ACL.

USB: Uploading an ACL command file from a USB device (CLI)

Syntax

```
copy usb command-file filename.txt [ unix | pc ]
```

Copies and executes the named text file from a USB flash drive and executes the ACL commands in the file.

filename.txt	A text file containing ACL commands and stored in the USB flash drive
unix pc	The type of workstation used to create the text file.

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 the *Access Security Guide*.)
- 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 USB flash drive.

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 usb command-file 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 in the tftp example shown in [Figure 153 \(page 295\)](#)), and continues to implement the remaining ACL commands in the file.

Copying diagnostic data to a remote host, USB device, PC, or UNIX workstation

Copying command output to a destination device (CLI)

Syntax

```
copy command-output "cli-command" tftp ip-address
filepath-filename [oobm]
copy command-output "cli-command" usb filename
copy command-output "cli-command" xmodem
```

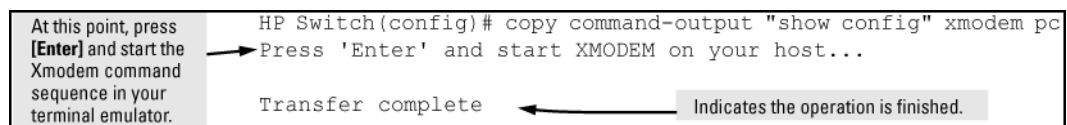
These commands direct the displayed output of a CLI command to a remote host, attached USB device, or to a serially connected PC or UNIX workstation.

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the OOBM interface. If this parameter is not specified, the transfer is through the data interface. The `oobm` parameter is not available on switches that do not have a separate OOBM port. For more information on out-of-band management, see [“Network Out-of-Band Management \(OOBM\) for the 6600 Switch”](#) (page 484).

Example

To use Xmodem to copy the output of `show config` to a serially connected PC:

Figure 154 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 smm [ tftp | usb | xmodem ]
copy event-log tftp ip-address filepath_filename [oobm]
copy event-log usb filename
copy event-log xmodem filename
```

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

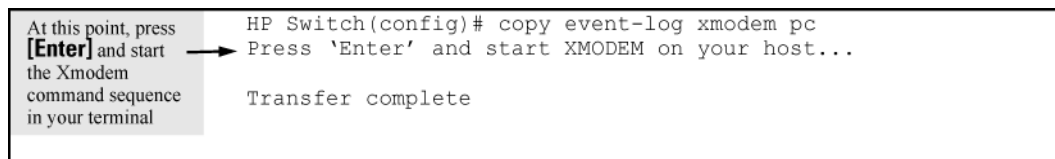
When used with the `smm` option, the entire Event Log, both active management module events and standby management module events, is copied to the selected host, USB device, or serially connected PC or UNIX workstation.

For switches that have a separate OOBM port, the `oobm` parameter specifies that the transfer is through the OOBM interface. If this parameter is not specified, the transfer is through the data interface. The `oobm` parameter is not available on switches that do not have a separate OOBM port. For more information on OOBM, see [“Network Out-of-Band Management \(OOBM\) for the 6600 Switch”](#) (page 484).

Example

To copy the event log to a PC connected to the switch:

Figure 155 Sending event log content to a file on an attached PC



Copying crash data content to a destination device (CLI)

This command uses TFTP, USB, or Xmodem to copy the Crash Data content to a destination device. You can copy individual slot information or the management module's switch information. If you do not specify either, the command defaults to the management function's data.

Syntax

```
copy crash-data [ slot-id | master ]
tftp ip-address filename [oobm]
copy crash-data [ slot-id | mm ]
usb filename
copy crash-data [ slot-id | mm ]
xmodem
```

These commands copy the crash data content to a remote host, attached USB device, or to a serially connected PC or UNIX workstation.

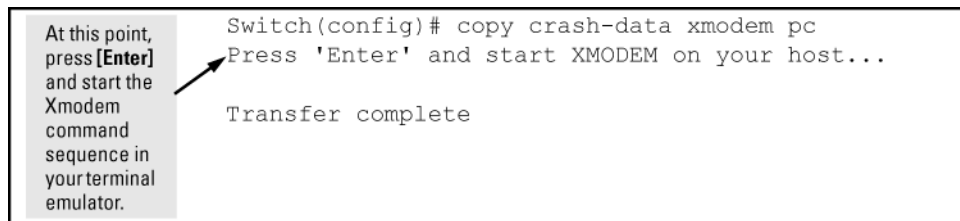
slot-id	a - h—Retrieves the crash log or crash data from the processor on the module in the specified slot
mm	Retrieves crash log or crash data from the switch's chassis processor. When "mm" is specified, crash files from both management modules are copied.
oobm	For switches that have a separate OOBM port, specifies that the transfer is through the OOBM interface. (Default is transfer through the data interface.)

You can copy individual slot information or the management module (mm) switch information. If you do not specify either, the command defaults to the mm data.

Example

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

Figure 156 Copying switch crash data content to a PC



Copying crash data with redundant management (CLI)

When you use redundant management, the `copy crash-data` command operates somewhat differently:

Syntax

```
copy crash-data [ slot-id | mm ] tftp ip-address filename
[oobm]
```

Copies the crash data of both the active and standby management modules to a user-specified file. If no parameter is specified, files from all modules (management and interface) are concatenated.

slot-id	Retrieves the crash data from the module in the specified slot.
mm	Retrieves the crash data from both management modules and concatenates them.
oobm	For switches that have a separate OOBM port, specifies that the transfer is through the OOBM interface. (Default is transfer through the data interface.)

Copying crash log data content to a destination device (CLI)

Syntax

```
copy crash-log [ slot-id | mm ] tftp ip-address filepath and
filename [oobm]
copy crash-log [ slot-id | mm ] usb filename
copy crash-log [ slot-id | mm ] xmodem
```

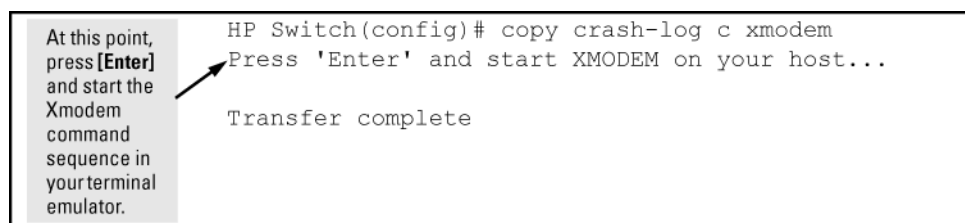
These commands copy the Crash Log content to a remote host, attached USB device, or to a serially connected PC or UNIX workstation. You can copy individual slot information or the management module (mm) switch information.

slot-id	a - h—Retrieves the crash log from the processor on the module in the specified slot
mm	Retrieves the crash log from the switch's chassis processor. When mm is specified, crash files from both management modules are copied.
oobm	For switches that have a separate OOBM port, specifies that the transfer is through the OOBM interface. (Default is transfer through the data interface.) If you do not specify either, the command defaults to the mm data.

Example

To copy the crash log for slot C to a file in a PC connected to the switch:

Figure 157 Sending a crash log for slot C to a file on an attached PC



Copying crash logs with redundant management (8200zl) (CLI)

When you use redundant management, the `copy crash-log` command operates somewhat differently.

Syntax

```
copy crash-log [ slot-id | mm ] tftp ip-address filepath and
filename [oobm]
```

Copies the crash logs of both the active and standby management modules to a user-specified file. If no parameter is specified, files from all modules (management and interface) are concatenated.

slot-id	Retrieves the crash log from the module in the specified slot.
mm	Retrieves the crash logs from both management modules and concatenates them.
oobm	For switches that have a separate OOBM port, specifies that the transfer is through the OOBM interface. (Default is transfer through the data interface.)

Copying coredumps from the standby management module (8200zl switches)

It is important that the coredump files on the standby management module are accessible for diagnostic purposes.

Syntax

```
copy core-dump [ mm usb filename | standby flash | usb filename ]
```

Copies the management module coredump or the standby management module coredump to the active management module flash or to a USB flash drive (see [Figure 158 \(page 300\)](#).)

flash	Copies the core file of the standby management module to the flash of the active management module. The destination file is fixed as dumpM1.cor or dumpM2.cor, depending on which module is the standby management module.
usb filename	Copies the management module's core file or the standby management module's core file to a USB flash drive. The optional filename defaults to dumpM1.cor or dumpM2.cor, depending on which module is the standby management module

While the file is being copied, the number of bytes transferred and the percentage of the total is displayed. Management module core files can be quite large. Use **Cntl-C** to cleanly cancel the transfer.

Example

Figure 158 Copying the standby coredump to flash

```
HP Switch(config)# copy core-dump standby flash
02816K of 26899K (10%)
```

If there is no coredump on the standby management module, the following error message displays:

```
Standby MM coredump does not exist.
```

If there is not enough destination space before or during the transfer to flash or USB, the following error message displays:

```
Insufficient FLASH space to complete the file copy.
```

Flight data recorder

The Flight Data Recorder (FDR) log collects information that is "interesting" when the switch is not performing correctly, but has not crashed. Runtime logs are written to FDR memory while the switch is running, and crashtime logs are collected and stored in the FDR buffer during a switch crash.

Syntax

```
copy fdr-log [[slot slot-list] | [mm-active [[current] | [previous]]]
| [mm-standby] | [all]]
tftp [[hostname] | [ip-addr]]
filename
```

Copies `fdr-log` files to a user-specified file.

<code>all</code>	Copies all the log files from both management modules and all slots.
<code>mn-active</code>	Copies the active management module's log.
<code>mn-standby</code>	Copies the standby management module's log.
<code>slot</code>	Retrieves the crash log from the module in the identified slots.

Enabling or disabling the USB port with the CLI

This feature allows configuration of the USB port with either the CLI or SNMP.

Syntax

```
usb-port
no usb-port
```

Enables the USB port. The `no` form of the command disables the USB port and any access to the device.

Viewing the status of the USB port (CLI)

Syntax

```
show usb-port
```

Displays the status of the USB port. It can be enabled, disabled, or not present. (See [Figure 159 \(page 301\)](#) or [Figure 160 \(page 301\)](#), depending on your version.)

Example

Figure 159 show `usb-port` command output on version K.13.59 and later

```
HP Switch(config)# show usb-port

USB port status: enabled
USB port power status: power on      (USB device detected in port)
USB port reset status: USB reset not required
```

Figure 160 show `usb-port` command output on version K.14.XX

```
HP Switch(config)# show usb-port

USB port status: enabled
USB port power status: power on      (USB device detected in port)
```

One of the following messages indicates the presence or absence of the USB device:

- Not able to sense device in USB port
- USB device detected in port
- No USB device detected in port

The reset status messages can be one of the following (K.13.XX only):

- Undetermined USB reset requirement
- USB reset not required
- USB device reset required for USB autorun

The autorun feature works only when a USB device is inserted and the USB port is enabled.

Using USB autorun

For more information, see [“About using USB autorun” \(page 311\)](#).

The general process for using USB autorun is as follows (*steps 1, 2, and 7 require an upcoming update to PCM+, as described above*):

1. Create an AutoRun file using PCM+.

See the HP Switch Manager documentation for details.

NOTE: Creating the AutoRun file in PCM+ includes the following steps:

- a. Specify the target device or devices.
 - b. Create the CLI script to be executed on the target devices.
 - c. Determine if the file will be signed and/or encrypted.
 - d. Determine if the file will be 'run once' (moved to a 'processed' directory on execution) or 'run many' (kept in the root directory of the flash drive from where it can be executed again.)
-

2. Deploy the AutoRun file to a USB flash drive.
3. (If required) Enable the autorun feature on the switch (autorun is enabled by default unless an operator or manager password has been set—See [“Autorun and configuring passwords” \(page 313\)](#).)
4. (If the AutoRun file has been signed or encrypted) Enable secure-mode on the switch:
 - a. Configure an encryption key and a valid trusted certificate
 - b. Enable secure-mode via the CLI.
See [“About downloading switch software” \(page 303\)](#).
5. Insert the USB flash drive into the switch's USB auxiliary port.
The switch processes the AutoRun file automatically and writes a result (.txt) file and report (.xml) file back to the USB flash drive, reporting on the command operations that were executed.
6. Remove the USB device from the USB port.
The switch executes any post-commands, such as rebooting the switch to apply any configuration updates.
7. (Optional) Transfer the 'result file' and 'report file' to a PCM+-enabled computer for report checking.
See [“Troubleshooting autorun operations” \(page 312\)](#).

Configuring autorun on the switch (CLI)

Syntax

```
[ no ] autorun [ encryption-key key-string | secure-mode ]
```

When executed from the configuration mode, enables or disables USB autorun on the switch.

Use the `encryption-key` keyword to configure or remove an encryption-key (a base-64 encoded string.) The encryption key is a prerequisite for enabling autorun in secure-mode. Encryption is regarded only when the AutoRun file is also signed by an authentic source.

Use the `secure-mode` keyword to enable or disable secure mode for autorun.
(Default: Enabled—or disabled if a password has been set)

For information about enabling secure mode on autorun, see “[Autorun secure mode](#)” (page 313).

Viewing autorun configuration information

The `show autorun` command displays autorun configuration status information, as shown in the following example.

```
HP Switch(config)# show autorun

Autorun configuration status

Enabled           : Yes
Secure-mode       : Disabled
Encryption-key    :
```

About 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 www.procurve.com 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.

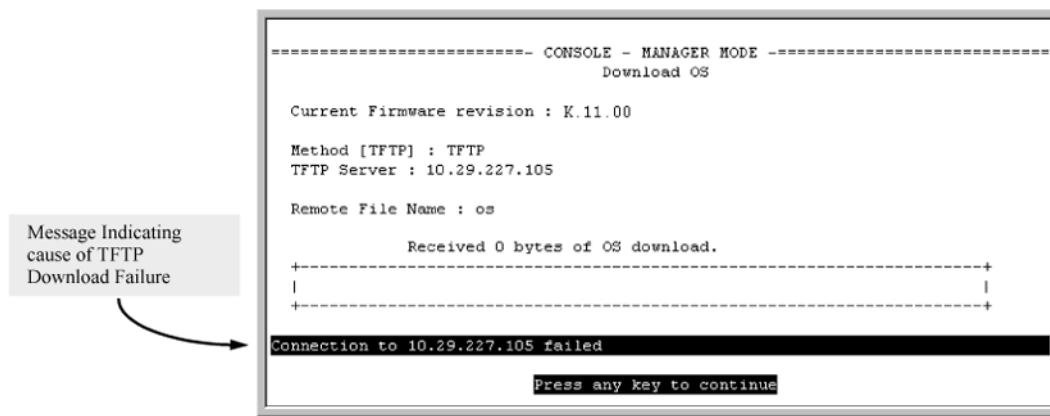
NOTE: Downloading new switch software does not change the current switch configuration. The switch configuration is contained in separate files that can also be transferred. See “[Transferring switch configurations](#)” (page 291).

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.

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 161 \(page 304\)](#).)

Figure 161 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 more on the Event Log, see [“Using the Event Log for troubleshooting switch problems” \(page 400\)](#).
- For descriptions of individual Event Log messages, see the *Event Log Message Reference Guide*.

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.

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.

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 [Figure 162 \(page 306\)](#).

Figure 162 Example of switch configuration with SFTP enabled

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

Running configuration:

```
; J8697 Configuration Editor; Created on release #K.11.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
no tftp-enable
password manager
password operator
```

```
ip ssh filetransfer
no tftp-enable
```

Enabling SFTP automatically disables TFTP and auto-tftp and displays this message.

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 163 (page 306)), or an SNMP application. Auto-TFTP is disabled by default and must be configured through the CLI.

Figure 163 Using the Menu interface to disable TFTP

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

System Name : ProCurve
System Contact :
System Location :

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

Time Zone [0] : 0
Daylight Time Rule [None] : None

Actions->  Cancel  Edit  Save  Help

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

Enables/Disables TFTP.
Note: If SFTP is enabled, this field will be set to **No**. You cannot use this field to enable TFTP if SFTP is enabled. Attempting to do so produces an **Inconsistent value** message in the banner below the **Actions** line.

- While SFTP is enabled, TFTP and auto-TFTP cannot be enabled from the CLI. Attempting to enable either non-secure TFTP option while SFTP is enabled produces one of the following messages in the CLI:

```
SFTP must be disabled before enabling tftp.
SFTP must be disabled before enabling auto-tftp.
```

Similarly, while SFTP is enabled, TFTP cannot be enabled using an SNMP management application. Attempting to do so generates an "inconsistent value" message. (An SNMP management application cannot be used to enable or disable auto-TFTP.)

- To enable SFTP by using an SNMP management application, you must first disable TFTP and, if configured, auto-TFTP on the switch. You can use either an SNMP application or the CLI to

disable TFTP, but you must use the CLI to disable auto-TFTP. The following CLI commands disable TFTP and auto-TFTP on the switch.

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

- 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 8212z1 only
| crash-data-e " "
| crash-data-f ""
| crash-data-g 8212z1 only
| 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 8212z1 only
| crash-log-e ""
| crash-log-f ""
| crash-log-g 8212z1 only
| 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 (this directory is not available on the 8212z1)
| mml.cor management module or management function
| im_a.cor interface module (chassis switches only)
| im_b.cor interface module (chassis switches only)
| im_1.cor interface module (chassis switches only)
| 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)

```

- When using SFTP to copy a software image onto the switch, the command return takes only a few seconds. However, this does not mean that the transfer is complete, because the switch requires additional time (typically more than one minute) to write the image to flash in the background. To verify the file transfer has been completed, you can use the `show flash` command or look for a confirmation message in the log, as in the following example:

```
I 01/09/09 16:17:07 00150 update: Primary Image updated.
```

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
```

About using USB to transfer files to and from the switch

The switch's USB port (labeled as *Auxiliary Port*) allows the use of a USB flash drive for copying configuration files to and from the switch. Beginning with software release K_12_XX or later, `copy` commands that used either `tftp` or `xmodem` now include an additional option for `usb` as a source or destination for file transfers.

Operating rules and restrictions on USB usage are:

- Unformatted USB flash drives must first be formatted on a PC (Windows FAT format.) For devices with multiple partitions, only the first partition is supported. Devices with secure partitions are not supported.
- If they already exist on the device, subdirectories are supported. When specifying a *filename*, you must enter either the individual file name (if at the root) or the full path name (for example, `/subdir/filename.`)

- To view the contents of a USB flash drive, use the `dir` command. This lists all files and directories at the root. To view the contents of a directory, you must specify the subdirectory name (that is, `dir subdirectory`.)
- The USB port supports connection to a single USB device. USB hubs to add more ports are not supported.

NOTE: Some USB flash drives may not be supported on your switch. Consult the latest *Release Notes* for information on supported devices.

Using PCM+ to update switch software

HP Switch Manager Plus includes a software update utility for updating on HP Switch switch products. For further information, see the *Getting Started Guide* and the *Administrator's Guide*, provided electronically with the application.

About transferring switch configurations

Using the CLI commands described in the section beginning with “[TFTP: Copying a configuration file to a remote host \(CLI\)](#)” (page 291), 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 304).

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 *Access Security Guide*.

About 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 *Access Security Guide*.

About copying diagnostic data to a remote host, USB device, 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`.
- Physically connected USB flash drive via the switch's USB port.
- Serially connected PC or UNIX workstation via **Xmodem**.

Behavior of autorun when USB port is disabled

Software versions K.13.XX operation

When using software version K.13.58, if the USB port is disabled (`no usb-port` command), the USB autorun function does not work in the USB port until the USB port is enabled, the config file is saved, and the switch is rebooted. The 5-volt power to the USB port remains on, even after the USB port has been disabled.

For software versions after K.13.58, the 5-volt power applied to the USB port is synchronized with the enabling of the USB port, that is, when the USB port is enabled, the 5 volts are supplied; when the USB port is disabled, the 5 volts are not supplied. For previous software versions, the power was supplied continuously. The autorun function does not require a switch reboot, but the USB device must be inserted at least once after the port is enabled so the switch recognizes that the device is present. If the USB device is inserted, and then the USB port is enabled, the switch does not recognize that a USB device is present.

Software version K.14.XX operation

For software versions K.14.XX, the USB port can be disabled and enabled without affecting the autorun feature. When the USB port is enabled, the autorun feature activates if a USB device is already inserted in the USB port.

Power is synchronized with the enabling and disabling of USB ports as described above for K.13.59 and later software.

About using USB autorun

USB autorun helps ease the configuration of HP Switch switches by providing a way to auto-execute CLI commands from a USB flash drive. Using this solution, you can create a command file (also known as an `AutoRun` file), write it to a USB storage device, and then execute the file simply by inserting the USB device into the switch's 'Auxiliary Port.' The `AutoRun` file is executed automatically when autorun is enabled on the switch and can be designed for various purposes, such as to configure the switch, to update software, or to retrieve diagnostic logs for troubleshooting purposes.

The overall USB autorun solution requires the following components:

- An HP Switch switch that can securely use USB autorun to load authorized configurations and write reporting information. This requires software versions K.13.01, T.13.01 or greater.
- The network management application *HP Switch Manager Plus* (PCM+.) PCM+ is required to create a valid `AutoRun` file and to view the results after the file has been executed on the switch.
- A non-proprietary USB flash drive.

NOTE: The ability to create a valid `AutoRun` file will be incorporated into an upcoming HP Switch Manager update; see the HP Switch Manager documentation for details. For guidelines on using the USB port for basic file copy capabilities, see [“About using USB to transfer files to and from the switch”](#) (page 309).

Security considerations

By default, the switch is unsecured when shipped (that is, USB autorun is enabled by default.) However, as soon as an operator or manager password is configured, autorun is disabled and must be re-enabled at the configuration level of the CLI before it can be used. The requirement to use PCM+ to create a valid `AutoRun` file helps prevent a nonauthorized command file from being created and processed by the switch.

In terms of physical security, access to the switch's console port and USB port are equivalent. Keeping the switch in a locked wiring closet or other secure space helps to prevent unauthorized

physical access. As additional precautions, you have the following configuration options via the CLI (see [Configuring autorun on the switch \(CLI\) \(page 302\)](#)):

- Disable autorun by setting an operator or manager password.
- Disable or re-enable the USB autorun function via the CLI.
- Enable autorun in secure mode to verify signatures in autorun command files and to decrypt encrypted command files.

Troubleshooting autorun operations

You can verify autorun operations by checking the following items:

USB auxiliary port LEDs

The following table shows LED indications on the Auxiliary Port that allow you to identify the different USB operation states.

Color	State	Meaning
Green	Slow blinking	Switch is processing USB AutoRun file.
Green	Solid	Switch has finished processing USB AutoRun file. This LED state indicates the AutoRun file was successfully executed and the report files were generated. You can review the report files on a USB-enabled computer for more details. Upon removal of the USB device, the LED turns OFF.
N/A	Off	Indicates one or more of the following: <ul style="list-style-type: none">• No USB device has been inserted.• A USB device that cannot be recognized as a USB storage device has been inserted.• No AutoRun file can be found on the inserted USB device.. If the USB device has just been removed from the port, the switch executes any post commands.
Amber	Fast blinking	Processing Error. The AutoRun file stops processing when an error is encountered (for example, no more disk space is available on the USB device to write the result and report files.) For more information on the error, remove the USB device and inspect its contents on a USB-enabled computer.

AutoRun status files.

The following files are generated during autorun operations and written to the USB flash drive:

- Report files (.xml file)—show which CLI commands have been run. The file name includes a serial number and datetime stamp to indicate when and on which device the AutoRun file was executed.
- Result files (.txt file)—contain the CLI output for each command that was run on the switch, allowing you to verify whether a command was executed successfully or not.

NOTE: PCM+ provides a mechanism to read these status files and capture the results of the commands executed. It also allows you to verify the report files for their authenticity and reject files that have not been signed (for details, see the HP Switch Manager documentation.)

The status files do not include any records of post commands that may have been executed after the USB flash drive was removed from the switch.

Event log or syslog

For details on how to use the switch's Event Log or syslog for help in isolating autorun-related problems, see [“Using the Event Log for troubleshooting switch problems” \(page 400\)](#).

Autorun secure mode

You can use autorun secure mode to verify the authenticity of autorun command files. Secure-mode is configured using the `autorun secure-mode` command and can be enabled under both of the following conditions:

- An encryption-key has already been configured using the `autorun encryption key` command.
- A trusted certificate for verifying autorun command files has been copied to the switch using the `copy [tftp | usb] autorun-cert-file` command.

There is an additional security option to install a valid key-pair for signing the result files that are generated during autorun operations. You can generate the key-pair on the switch using the `crypto key generate autorun [rsa]` command.

NOTE: You can also install the key-pair from a tftp server or via the USB port using the `copy [tftp | usb] autorun-key-file ipaddr filename` command. The filename must contain the private key and the matching public key in a X509 certificate structure. Both the private key and the X509 certificate must be in PEM format.

Operating notes and restrictions

- Autorun is enabled by default, until passwords are set on the device.
- Secure-mode and encryption-key are disabled by default.
- To enable secure mode, both an encryption key and trusted certificate must be set.
- If secure-mode is enabled, the following conditions apply:
 - The encryption-key cannot be removed or unconfigured.
 - The key-pair cannot be removed.
- If secure mode is disabled, the key-pair can be removed using the `crypto key zeroize autorun` command.
- When installing the autorun certificate file and/or the other key files, the files must be in PEM format.

Autorun and configuring passwords

When an operator or manager password is configured on a switch, autorun is disabled automatically, and a message is displayed on the screen, as shown in the following example:

```
HP Switch# password manager
New password for manager: *****
Please retype new password for manager: *****
Autorun is disabled as operator/manager is configured.
```

After passwords are set, you can re-enable autorun as needed using the `autorun` command. For more information on configuring passwords, see the *Access Security Guide*.

B Monitoring and Analyzing Switch Operation

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 164 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.

Accessing system information (CLI)

Syntax

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

Displays global system information and operational parameters for the switch.

chassislocate	Displays 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 figure Figure 165 (page 315) .)
information	Displays global system information and operational parameters for the switch. (See Figure 167 (page 315) .)
power-supply	Shows chassis power supply and settings.
temperature	Shows system temperature and settings.
fans	Shows system fan status. (See Figure 166 (page 315) .)

Example

Figure 165 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
```

Figure 166 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
```

Figure 167 Switch system information

```
HP Switch(config)# show system

Status and Counters - General System Information

System Name       : HP Switch 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                 Memory - Free   : 110,527,264

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

Locating a switch

To locate where a specific 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.

blink 1–1440	Blinks the chassis locate LED for a specified number of minutes (Default: 30 min.)
on 1–1440	Turns the chassis locate LED on for a specified number of minutes (Default: 20 min.)
off	Turns the chassis locate LED off.

Chassislocate at Boot

The chassislocate command has an optional parameter that configures it to run in the future instead of immediately.

Syntax

```
chassislocate <on|blink>[Minutes] [at startup] chassislocate  
off
```

at Specify when the command is applied (default immediately.)

now Turn on the chassis locate LED immediately.

startup Turn on the chassis locate LED at switch startup.

Example

```
chassislocate blink 10 at startup
```

Show command

Syntax

```
HP Switch# show system chassislocate
```

Display

```
Locator LED Status  
          Current   Time  
Member   State      Remaining  Configuration  
-----  -  
1        blink     00:27:05  blink 30 at startup  
2        on        01:05:27  
3        off
```

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)

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 available as well.)

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

Example

Figure 168 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
```

Accessing system information (Menu)

From the console Main Menu, select:

- 1. Status and Counters**
- 1. General System Information**

Figure 169 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 switch management address information (CLI)

Syntax

```
show management
```

Accessing switch management address information (Menu)

From the Main Menu, select:

1. Status and Counters ...

2. Switch Management Address Information

Figure 170 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 170 \(page 318\)](#), 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, see the *Advanced Traffic Management Guide*.

Viewing additional component information (CLI)

The CLI `show modules` command displays additional component information for the following:

- SSM—identification, including serial number
- Mini-GBICS—a list of installed mini-GBICs displaying the type, "J" number, and serial number (when available)

Syntax

```
show modules [details]
```

Displays information about the installed modules ([Figure 171 \(page 319\)](#)), including:

- The slot in which the module is installed
- The module description
- The serial number
- The SSM description, serial number, and status (8200zl switches only)

Additionally, this command displays the part number (J number) and serial number of the chassis. (See [Figure 172 \(page 319\)](#).)

Examples

Figure 171 The `show modules` command output

```
HP Switch(config)# show modules

Status and Counters - Module Information

Chassis: 5406z1 J8697A          Serial Number:  SG560TN124
Slot  Module Description          Serial Number
-----
A     HP Switch J8706A 24p SFP z1 Module  AD722BX88F
B     HP Switch J8702A 24p Gig-T z1 Module FE999CV77F
C     HP Switch J8707A 4p 10-Gbe z1 Module FB345DC99D
```

Figure 172 The `show modules details` command for the 8212z1, showing SSM and mini-GBIC information

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

Status and Counters - Module Information

Chassis: 8212z1 J8715A          Serial Number:  SG560TN124
Slot  Module Description          Serial Number  Status
-----
MM1   HP Switch J9092A Management Module 8200z1  AD722BX88F    Active
SSM   HP Switch J8784A System Support Module    AF988DC78G    Active
C     HP Switch J8750A 20p +4 Mini-GBIC Module  446S2BX007    Active
      GBIC 1: J4859B 1GB LX-LC             4720347DFED734
      GBIC 2: J4859B 1GB LX-LC             4720347DFED735
```

NOTE: On HP Switch 3500yl and 6200yl series switches, the mini-GBIC information does not display, because the ports are fixed and not part of any module.

Viewing port status (Menu)

From the Main Menu, select:

- 1. Status and Counters ...**
- 3. Module Information**

Enabling and Disabling Compatibility Mode for v2 zl and zl modules

NOTE: In the following context, v2 zl modules are the second version of the current zl modules. Both v2 zl and zl modules are supported in the 5400z1 and 8200z1 series chassis switches. Compatibility Mode allows the inter-operation of v2 zl modules with zl modules in a chassis switch. When in Compatibility Mode, the switch accepts either v2 zl or zl modules. The default is Compatibility Mode enabled. If Compatibility Mode is disabled by executing the `no allow-v1-modules` command, the switch will only power up v2 zl modules.

Syntax

```
[ no ] allow-v1-modules
```

Enables Compatibility Mode for interoperation of v2 zl and zl modules in the same chassis. (See [Figure 173 \(page 320\)](#).)

The `no` form of the command disables Compatibility Mode. Only the v2 zl modules are powered up. (See [Figure 174 \(page 320\)](#).)

(Default: Enabled.)

For information about how the switch modules behave in various situations and combinations, see [“Compatibility mode” \(page 352\)](#).

Examples

Figure 173 Enabling compatibility mode

```
HP Switch(config)# allow-vl-modules
This will erase the configuration and reboot the switch.
Continue [y/n]?
```

Figure 174 Disabling compatibility mode

```
HP Switch(config)# no allow-vl-modules
All V1 modules will be disabled. Continue [y/n]?
```

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 175 Example of port status on the menu interface

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

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.
```

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.

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 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.

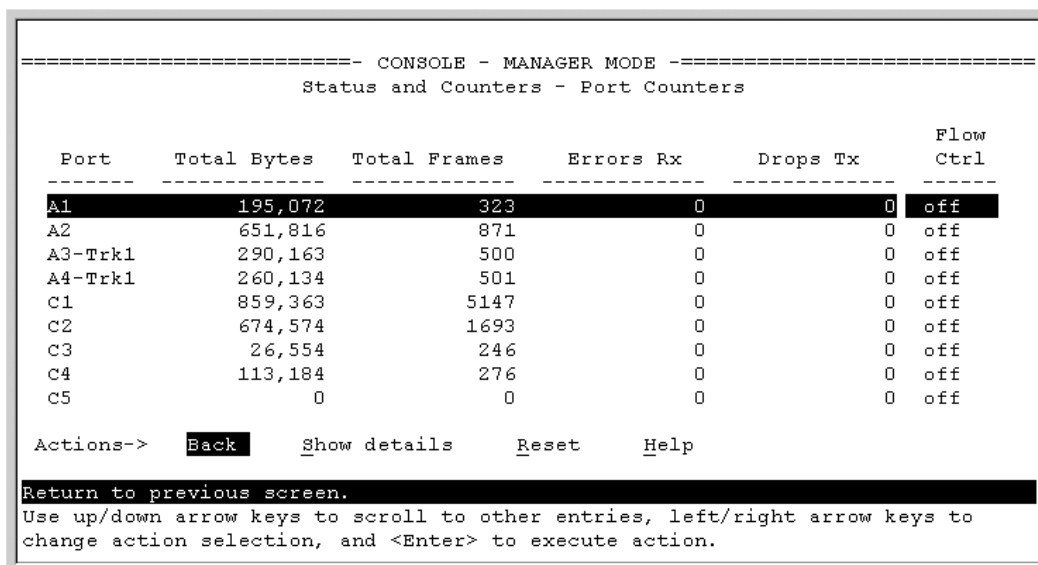
NOTE: to reset the port counters to zero, you must reboot the switch.

Accessing port and trunk statistics (Menu)

From the Main Menu, select:

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

Figure 176 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 177 \(page 322\)](#), below.

Figure 177 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” [\(page 353\)](#).)

Notes:

Once cleared, statistics cannot be reintroduced.

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:

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.

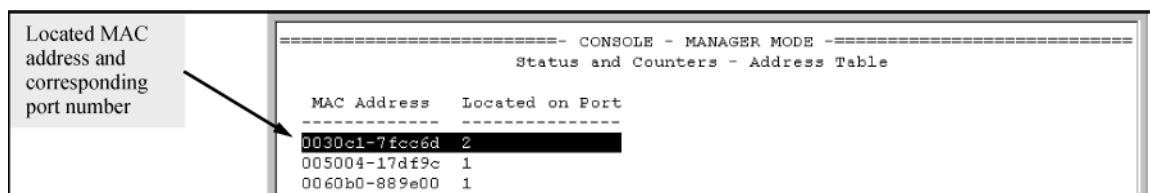
```

----- 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 178 (page 323).)

Figure 178 Example of the address table



```

----- CONSOLE - MANAGER MODE -----
                          Status and Counters - Address Table
-----
  MAC Address  Located on Port
-----
0030c1-7fcc6d  2
005004-17df9c  1
0060b0-889e00  1
  
```

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.

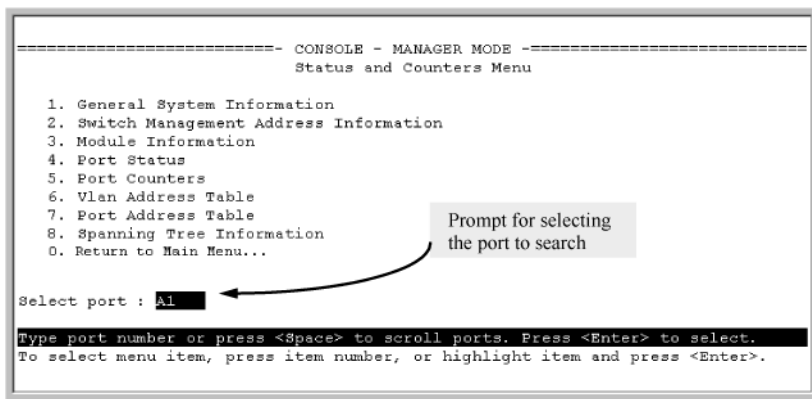
1. Proceeding from Figure 178 (page 323), press **[S]** (for **Search**), to display the following prompt:

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 179 (page 324).) If the switch does not find the MAC address on the currently selected VLAN, it leaves the MAC address listing empty.

Figure 179 Example of menu indicating located MAC address



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**

Figure 180 Listing MAC addresses for a specific port

```
Switch-1(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 : 000883-028300
IST Regional Root Priority    : 32768
IST Regional Root Path Cost   : 2000000
IST Remaining Hops            : 19
```

Port	Type	Cost	Priority	State	Designated Bridge	Hello Time	PtP	Edge
A1	10/100TX	Auto	128	Forwarding	000883-028300	9	Yes	No
A2	10/100TX	Auto	128	Blocking	0001e7-948300	9	Yes	No
A3	10/100TX	Auto	128	Forwarding	000883-02a700	2	Yes	No
A4	10/100TX	Auto	128	Disabled				
A5	10/100TX	Auto	128	Disabled				
.				
.				

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 \(page 324\)](#), 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.

- 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 181 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
|-----
```

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).

Port	Type	Cost	Priority	State	Designated Bridge	Hello Time	PtP	Edge
A1	100/1000T	Auto	128	Forwarding	000883-028300	9	Yes	No
A2	100/1000T	Auto	128	Blocked	0001e7-948300	9	Yes	No
A3	100/1000T	Auto	128	Forwarding	000883-02a700	2	Yes	No
A4	100/1000T	Auto	128	Disabled				
A5	100/1000T	Auto	128	Disabled				
.				
.				

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)

Show command	Output
<code>show ip igmp</code>	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
<code>show ip igmp config</code>	Displays the IGMP configuration information, including VLAN ID, VLAN name, status, forwarding, and Querier information.
<code>show ip igmp <i>vlan-id</i></code>	Per-VLAN command listing above, IGMP status for specified VLAN (VID)
<code>show ip igmp group <i>ip-addr</i></code>	Lists the ports currently participating in the specified group, with port type, Access type, Age Timer data and Leave Timer data.
<code>show ip igmp groups</code>	Displays VLAN-ID, group address, uptime, expiration time, multicast filter type, and the last reporter for IGMP groups.
<code>show ip igmp statistics</code>	Displays IGMP operational information, such as VLAN IDs and names, and filtered and flooding statistics.

Examples

Example 62 Output from show ip igmp config command

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

IGMP Service

VLAN ID VLAN Name      IGMP   Forward with Querier Querier
-----  -
1        DEFAULT_VLAN  No     No           Yes    125
2        VLAN2         Yes    No           Yes    125
12       New_Vlan     No     No           Yes    125
```

Example 63 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 <i>vlan-id</i>	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 figures show how you could list data on the above VLANs.

Figure 182 Listing the VLAN ID (vid) and status for specific ports

```
HP Switch# show vlan ports A1-A2

Status and Counters = VLAN Information - for ports A1,A2

802.1Q VLAN ID Name                Status
-----
1          DEFAULT_VLAN              Static
33         VLAN-33                   Static
```

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

Figure 183 Example of VLAN listing for the entire switch

```
HP Switch# show vlan

Status and Counters - VLAN Information

VLAN support : Yes
Maximum VLANs to support : 9
Primary VLAN: DEFAULT_VLAN

802.1Q VLAN ID Name                Status
-----
1          DEFAULT_VLAN              Static
33         VLAN-33                   Static
44         VLAN-44                   Static
```

Figure 184 Port listing for an individual VLAN

```
HP Switch(config)# show vlan 1

Status and Counters - VLAN Information - VLAN 1

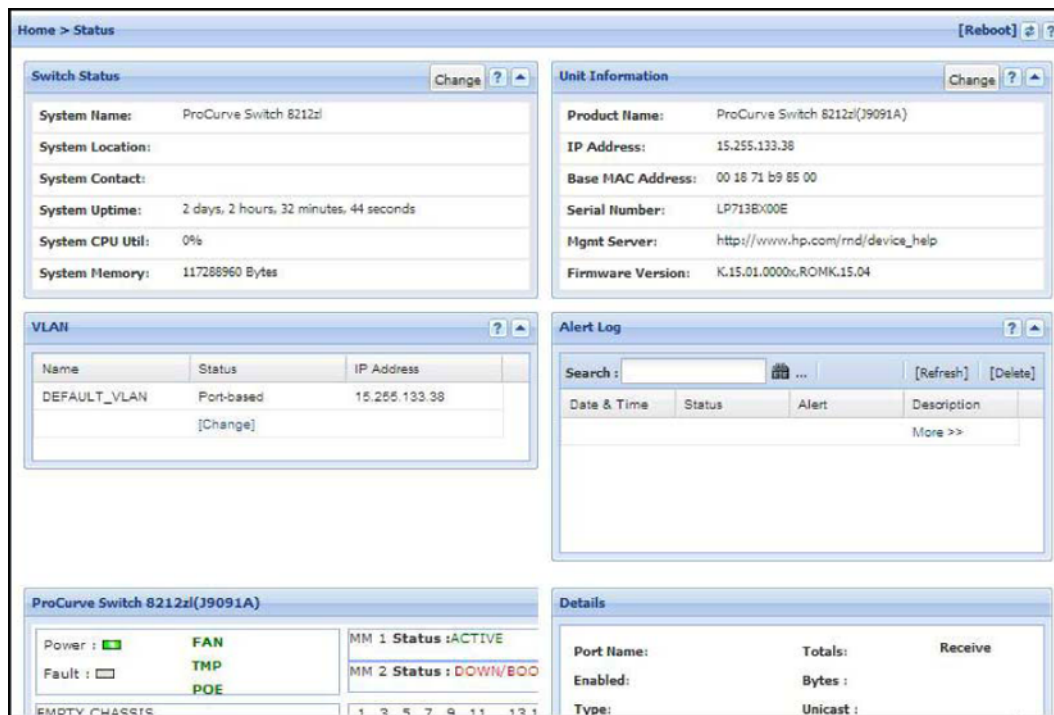
VLAN ID : 1
Name : DEFAULT_VLAN
Status : Static
Voice : Yes
Jumbo : No

Port Information Mode      Unknown VLAN Status
-----
A1          Untagged Learn      Up
A2          Tagged   Learn      Up
A3          Untagged Learn      Up
A4          Untagged Learn      Down
A5          Untagged Learn      Up
A6          Untagged Learn      Up
A7          Untagged Learn      Up
```

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 185 \(page 329\)](#).

Figure 185 Example of a WebAgent status screen



Configuring local mirroring (CLI)

Local mirroring overview

To configure a local mirroring session in which the mirroring source and destination are on the same switch, follow these general steps:

1. Determine the session and local destination port:
 - Session number (1-4) and (optional) alphanumeric name
 - Exit port (any port on the switch except a monitored interface used to mirror traffic)

⚠ 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.

2. Enter the mirror *session-#* [*name session-name*] port *port-#* command to configure the session.
 3. Determine the traffic to be selected for mirroring by any of the following methods and the appropriate configuration level (VLAN, port, mesh, trunk, switch):
 - a. Direction: inbound, outbound, or both
 - b. Classifier-based mirroring policy: inbound only for IPv4 or IPv6 traffic
 - c. MAC source and/or destination address: inbound, outbound, or both
 4. Enter the monitor command to assign one or more source interfaces to the session.
- After you complete step 4, the switch begins mirroring traffic to the configured exit port.

Setting up local mirroring

The following commands configure mirroring for a local session in which the mirroring source and destination are on the same switch.

- The `mirror` command identifies the destination in a mirroring session.
- The `interface` and `vlan` commands identify the mirroring source, including source interface, traffic direction, and traffic-selection criteria for a specified session.

Configuring a local mirroring session

For more information, see [“Configure a mirroring session on the source switch”](#) (page 362).

Syntax

```
[ no ] mirror 1 - 4 port exit-port-# [name name-str ]
```

The `no mirror session-# port` command removes the mirroring session and any mirroring source previously assigned to that session by the following commands.

Configuring traffic-direction criteria to select traffic

For more information, see [“About selecting all inbound/outbound traffic to mirror”](#) (page 364).

Syntax

```
[ no ] [ interface port/trunk/mesh | vlan vid-# ] monitor all  
in | out | both mirror session [ session ]... [no-tag-added]
```

Configuring ACL criteria to select inbound traffic — deprecated

Deprecated command:

Syntax

```
[ no ] [ interface port/trunk/mesh | vlan vid-# ] monitor ip  
access-group acl-name in mirror session [ session ... ]
```

Configuring a mirroring policy to select inbound traffic

For more information, see [“About selecting inbound traffic using advanced classifier-based mirroring”](#) (page 367).

Syntax

```
class ipv4 | ipv6 classname [ no ] [seq-number] [ match |  
ignore ip-protocol source-address destination-address ]  
[precedence precedence-value] [tos tos-value] [ip-dscp  
codepoint] [vlan vlan-id]
```

Syntax

```
policy mirror policy-name [no] [seq-number] [ class ipv4 | ipv6  
classname action mirror session ] [ action mirror session  
]... [ no ]  
default-class action mirror session [ no ] [ interface  
port/trunk | vlan vid-# ]  
service-policy mirror-policy-name in
```

In the `policy mirror` command, the `mirror session` parameter accepts a number (1 to 4) or name, if the specified mirroring session has already been configured with the name `name-str` option in the `mirror` command.

The `[no] interface | vlan service-policy` in command removes the mirroring policy from a port, VLAN, trunk, or mesh interface for a specified session, but leaves the session available for other assignments.

Configuring MAC-based criteria to select traffic

For more information, see [“About selecting inbound/outbound traffic using a MAC address” \(page 366\)](#).

Syntax

```
[ no ] monitor mac mac-addr [ src | dst | both ] mirror  
session
```

Enter the `monitor mac mirror` command at the global configuration level.

Use the `no` form of the complete command syntax (for example, `no monitor mac 112233-445566 src mirror 3`) to remove a MAC address as mirroring criteria from an active session on the switch without removing the session itself.

Configuring a remote mirroring destination (CLI)

For more information, see [“Configure a mirroring destination on a remote switch” \(page 362\)](#).

On the remote switch

Syntax

```
mirror endpoint ip src-ip src-udp-port dst-ip exit-port  
[truncation]
```

On the local switch

For more information, see [“Configure a source switch in a remote mirroring session” \(page 363\)](#).

Syntax

```
mirror session remote ip src-ip src-udp-port dst-ip
```

Configuring a local mirroring destination on the local switch (CLI)

For more information, see [“Configure a source switch in a remote mirroring session” \(page 363\)](#).

Syntax

```
mirror session port exit-port
```

Configuring monitored traffic (CLI)

Deprecation of ACL-based traffic selection:

In release K.14.01 and greater, the use of ACLs to select inbound traffic in a mirroring session `interface | vlan monitor ip access-group` in `mirror` command has been deprecated and is replaced with classifier-based mirroring policies. For more information, see [“About selecting inbound traffic using advanced classifier-based mirroring” \(page 367\)](#).

Syntax

```
interface port/trunk/mesh
```

Syntax

```
monitor all [ in | out | both ] mirror session [no-tag-added]
```

Syntax

```
monitor ip access-group acl-name in mirror session ()
```

For more information, see [“About selecting inbound traffic using an ACL \(deprecated\)” \(page 365\)](#).

Syntax

```
service-policy mirror-policy-name in
```

For more information, see [“Configuring classifier-based mirroring \(CLI\)” \(page 342\)](#).

Syntax

```
vlan vid-#
```

Syntax

```
monitor all[ in | out | both ]mirror session
```

For more information, see [“About selecting inbound traffic using an ACL \(deprecated\)” \(page 365\)](#).

Syntax

```
monitor ip access-group acl-name in mirror session ()
```

For more information, see [“About selecting inbound traffic using an ACL \(deprecated\)” \(page 365\)](#).

Syntax

```
service-policy mirror-policy-name in
```

For more information, see [“About selecting inbound traffic using an ACL \(deprecated\)” \(page 365\)](#).

Syntax

```
monitor mac mac-addr src | dest | both mirror
```

For more information, see [“Configuring classifier-based mirroring \(CLI\)” \(page 342\)](#).

Syntax

```
show monitor[ endpoint | session-number | name session-name ]
```

For more information, see [“Viewing a classifier-based mirroring configuration \(CLI\)” \(page 343\)](#).

Configuring local mirroring (Menu)

For information about local mirroring, see [“Remote mirroring overview” \(page 360\)](#).

NOTE: If mirroring has already been enabled on the switch, the Menu screens appear different from the one shown in this section.

Procedure 1

1. From the Main Menu, select:
 1. **Switch Configuration ...**
 3. **Network Monitoring Port**

Figure 186 The default network mirroring configuration screen

```
----- CONSOLE - MANAGER MODE -----
Switch Configuration - Network Monitoring Port

Monitoring Enabled [No] : No
                          ^
                          |
                          | Enable mirroring by
                          | setting this parameter
                          | to "Yes".

Actions->  Cancel    Edit    Save    Help

Select whether to enable traffic monitoring.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

2. In the Actions menu, press **[E]** (for **Edit**.)
3. If mirroring is currently disabled for session 1 (the default), enable it by pressing the Space bar (or **[Y]**) to select **Yes**.
4. Press the down arrow key to display a screen similar to [Figure 187 \(page 333\)](#), and move the cursor to the **Monitoring Port** parameter.

Figure 187 How to select a local exit port

```
----- CONSOLE - MANAGER MODE -----
Switch Configuration - Network Monitoring Port

Monitoring Enabled [No] : Yes
Monitoring Port :
Monitor : Ports

Port  Type  Action  |  Port  Type  Action
-----+-----+-----|-----+-----+-----
1    1000T    |      | 31    1000T    |
2    1000T    |      | 32    1000T    |
3    1000T    |      | 33    1000T    |
4    1000T    |      | 34    1000T    |
5    1000T    |      | 35    1000T    |
6    1000T    |      | 36    1000T    |
7    1000T    |      | 37    1000T    |
8    1000T    |      | 38    1000T    |

Actions->  Cancel    Edit    Save    Help

Select the port that will act as the Monitoring Port.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

5. Use the Space bar to select the port to use for sending mirrored traffic to a locally connected traffic analyzer or IDS.
(The selected interface must be a single port. It cannot be a trunk or mesh.) In this example, port 5 is selected as the local exit port.
6. Highlight the Monitor field and use the Space bar to select the interfaces to mirror:
 - Ports:** Use for mirroring ports, static trunks, or the mesh.
 - VLAN:** Use for mirroring a VLAN.
7. Do one of the following:
 - If you are mirroring ports, static trunks, or the mesh, go to [step 8 \(page 334\)](#).
 - If you are mirroring a VLAN:
 - i. Press **[Tab]** or the down arrow key to move to the **VLAN** field.

```

----- CONSOLE - MANAGER MODE -----
Switch Configuration - Network Monit
Monitoring Enabled [No] : Yes
Monitoring Port : 5
Monitor : VLAN
VLAN : ██████████

```

Use the Space bar to select a VLAN to mirror.

- ii. Use the Space bar to select the **VLAN** you want to mirror.
 - iii. Go to [step 10 \(page 334\)](#).
8. Use the down arrow key to move the cursor to the **Action** column for the individual port interfaces and position the cursor at a port, trunk, or mesh you want to mirror.

```

----- CONSOLE - MANAGER MODE -----
Switch Configuration - Network Monitoring Port
Monitoring Enabled [No] : Yes
Monitoring Port : 5
Monitor : Ports

```

Port	Type	Action	Port	Type
1	1000T	██████████	31	1000T
2	1000T		32	1000T
3	1000T		33	1000T
4	1000T		34	1000T
5	1000T		35	1000T
6	1000T		36	1000T
7	1000T		37	1000T
8	1000T		38	1000T

Use the down arrow key to select the interface(s) whose traffic you want to mirror to the local exit port.

Actions-> Cancel Edit Save Help

Select whether to monitor the selected port.
Use arrow keys to change field selection, <Space> to toggle field choices, and <Enter> to go to Actions.

9. Press the Space bar to select **Monitor** for the ports, trunks, mesh, or any combination of these that you want mirrored.
- Use the down arrow key to move from one interface to the next in the **Action** column. (If the mesh or any trunks are configured, they appear at the end of the port listing.)
10. When you finish selecting interfaces to mirror, press **[Enter]**, then press **[S]** (for **Save**) to save your changes and exit from the screen.
11. Return to the Main Menu.

Using the CLI, you can configure a mirroring session for a destination device connected to an exit port on either:

- The same switch as the source interface (local mirroring.)
- A different switch (remote mirroring.) The remote switch must be an HP switch offering the full mirroring capabilities described in this chapter.

CAUTION: After you configure a mirroring session with traffic-selection criteria and a destination, the switch immediately starts to mirror traffic to each destination device connected to an exit port. In a remote mirroring session that uses IPv4 encapsulation, if the exit switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic. For this reason, HP Switch strongly recommends that you configure the exit switch for a remote mirroring session before configuring the source switch for the same session.

Configuring the mirroring destination on a remote switch (CLI)

For more information, see [“Configure a mirroring destination on a remote switch” \(page 362\)](#).

Syntax

```

mirror endpoint ip src-ip-addr src-udp-port dst-ip-addr port
exit-port

```

Enter this command on a remote switch to configure the exit port to use in a remote mirroring session. You will configure the mirroring source on the local switch in the next step (see [“Configuring the mirroring source on the local switch \(CLI\)”](#) (page 335).)

The `mirror endpoint ip` command configures:

- The unique UDP port number to be used for the mirroring session on the source switch. The recommended port range is from 7933 to 65535.
- The IP address of the source switch to use in the session.
- The IP address and exit-port number on the remote (endpoint) switch.

In a remote mirroring endpoint, the IP address of exit port and the remote destination switch can belong to different VLANs.

Configuring the mirroring source on the local switch (CLI)

For more information, see [“Configure a mirroring session on the source switch”](#) (page 362).

Syntax

```
mirror 1 - 4 [name name-str ] remote ip src-ip src-udp-port  
dst-ip [truncation]
```

The `no mirror 1 - 4` command form removes both the mirroring session and any mirroring sources previously assigned to the session by the following commands.

Configuring traffic-direction criteria to select traffic (CLI)

Syntax

```
[ no ] [ interface port/trunk/mesh | vlan vid-# ] monitor  
all in | out | both mirror 1 - 4 | name-str [ 1 - 4 |  
name-str ...]
```

Configuring ACL criteria to select inbound traffic

Deprecated command:

For more information, see [“About selecting inbound traffic using an ACL \(deprecated\)”](#) (page 365).

Syntax

```
[ no ] [ interface port/trunk/mesh | vlan vid-# ] monitor ip  
access-group acl-name in  
mirror [ 1 - 4 | name-str ] [ 1 - 4 | name-str ...]
```

Configuring a mirroring policy to select inbound traffic (CLI)

For more information, see [“About selecting inbound traffic using advanced classifier-based mirroring”](#) (page 367).

Syntax

```
class ipv4 | ipv6 classname [ no ] [seq-number] [ match |  
ignore ip-protocol source-address destination-address ]  
[precedence precedence-value] [tos tos-value] [ip-dscp  
codepoint] [vlan vlan-id]
```

Syntax

```
policy mirror policy-name [no][seq-number][ class ipv4 | ipv6
classname action mirror session ][ action mirror session
]... [ no ]
default-class action mirror session [ no ][ interface
port/trunk | vlan vid-# ]
service-policy mirror-policy-name in
```

In the `policy mirror` command, the `mirror session` parameter accepts a number (1 to 4) or name, if the specified mirroring session has already been configured with the name `name-str` option in the `mirror` command.

The

```
no[ interface | vlan ] service-policy in
```

command removes the mirroring policy from a port, VLAN, trunk, or mesh interface for a specified session, but leaves the session available for other assignments.

Configuring the MAC-based criteria to select traffic (CLI)

For more information, see [“About selecting inbound/outbound traffic using a MAC address” \(page 366\)](#).

Syntax

```
[ no ] monitor mac mac-addr [ src | dst | both ] mirror
session
```

NOTE: If you have already configured session 1 with a destination, you can enter the `vlan vid monitor` or `interface port monitor` command without traffic-selection criteria and session identifier to:

- Overwrite the existing session 1 configuration.
 - Automatically configure mirroring in session 1 for inbound and outbound traffic on specified VLAN or port interfaces with the preconfigured destination.
-

Configuring a destination switch in a remote mirroring session (CLI)

- △ **CAUTION:** When configuring a remote mirroring session, *always* configure the destination switch first. Configuring the source switch first can result in a large volume of mirrored, IPv4-encapsulated traffic arriving at the destination without an exit path, which can slow switch performance.
-

Syntax

```
mirror endpoint ip src-ip src-udp-port dst-ip exit-port-#
no mirror endpoint ip src-ip src-udp-port dst-ip
```

Used on a destination switch to configure the remote endpoint of a mirroring session. The command uniquely associates the mirrored traffic from the desired session on a monitored source with a remote exit port on the destination switch. You must use the same set of source and destination parameters used when you configure the same session on both the source and destination switches.

For a given mirroring session, the same `src-ip`, `src-udp-port` and `dst-ip` values must be entered with the `mirror endpoint ip` command on the destination switch, and later with the `mirror remote ip` command on the

source switch. For more information, see the [mirror remote ip \(page 363\)](#) command syntax.

- ⚠ **CAUTION:** Do not remove the configuration of a remote mirroring endpoint support for a given session if there are source switches currently configured to mirror traffic to the endpoint.

<i>src-ip</i>	Must exactly match the <i>src-ip</i> address you configure on the source switch for the remote session.
<i>src-udp-port</i>	Must exactly match the <i>src-udp-port</i> value you configure on the source switch for the remote session. The recommended port range is 7933 to 65535. This setting associates the monitored source with the desired remote endpoint in the remote session by using the same, unique UDP port number to identify the session on the source and remote switches.
<i>dst-ip</i>	Must exactly match the <i>dst-ip</i> setting you configure on the source switch for the remote session.
<i>exit-port-#</i>	Exit port for mirrored traffic in the remote session, to which a traffic analyzer or IDS is connected.

The `no` form of the command deletes the mirroring endpoint for the configured session on the remote destination switch.

Configuring a source switch in a local mirroring session (CLI)

Enter the `mirror port` command on the source switch to configure an exit port on the same switch. To create the mirroring session, use the information gathered in [“High-level overview of the mirror configuration process” \(page 361\)](#).

Syntax

```
mirror 1 - 4 port exit-port-# [name name-str ]
no mirror 1- 4
```

Assigns the exit port to use for the specified mirroring session and must be executed from the global configuration level.

1 - 4	Identifies the mirroring session created by this command. (Multiple sessions on the switch can use the same exit port.)	
name <i>name-str</i>	Optional alphanumeric name string used to identify the session (up to 15 characters)	
port <i>exit-port-#</i>	Exit port for mirrored traffic in the remote session. This is the port to which a traffic analyzer or IDS is connected.	

The `no` form of the command removes the mirroring session and any mirroring source previously assigned to that session. To preserve the session while deleting a mirroring source assigned to it, see the `no` command descriptions under [“Configure the monitored traffic in a mirror session” \(page 363\)](#).

Configuring a source switch in a remote mirroring session (CLI)

Syntax

```
[ no ] mirror 1 - 4 [name name-str ] remote ip src-ip
src-udp-port dst-ip [truncation]
```

Used on the source switch to uniquely associate the mirrored traffic in the specified session with a remote destination switch. You must configure the same source and

destination parameters when you configure the same session on both the source and destination switches. (If multiple remote sessions use the same source and destination IP addresses, each session must use a unique UDP port value.)

When you execute this command, the following message is displayed:

Caution: Please configure destination switch first.
Do you want to continue [y/n]?

- If you have not yet configured the session on the remote destination switch, follow the configuration procedure in [“Configure a mirroring destination on a remote switch” \(page 362\)](#) before using this command.
- If you have already configured the session on the remote destination switch, enter **y** (for "yes") to complete this command.

1 - 4	Identifies the mirroring session created by this command.	
name <i>name-str</i>	Optional alphanumeric name string used as an additional session identifier (up to 15 characters.)	
<i>src-ip</i>	The IP address of the VLAN or subnet on which the traffic to be mirrored enters or leaves the switch.	
<i>src-udp-port</i>	<p>Associates the remote session with a UDP port number. When multiple sessions have the same source IP address <i>src-ip</i> and destination IP address <i>dst-ip</i>, the UDP port number must be unique in each session. The UDP port number used for a given session should be in the range of 7933 to 65535.</p> <p>CAUTION: UDP port numbers below 7933 are reserved for various IP applications. Using them for mirroring can result in the interruption of other IP functions and in non-mirrored traffic being received on the destination switch and sent to a device connected to the remote exit port.</p> <p>The configured UDP port number is included in the frames mirrored from the source switch to the remote destination switch (<i>mirror endpoint</i>), and enables the remote switch to match the frames to the exit port configured for the combined UDP port number, source IP address, and destination IP address. For more information, see the <i>mirror endpoint ip</i> command syntax in “Configure a mirroring destination on a remote switch” (page 362).</p>	
<i>dst-ip</i>	For the remote session specified in the command, this is the IP address of the VLAN or subnet on which the remote exit port exists. (The exit port to which a traffic analyzer or IDS is connected is configured on the remote switch in section “Configure a mirroring destination on a remote switch” (page 362) .)	
[truncation]	Enables truncation of oversize frames, causing the part of the frame in excess of the MTU size to be truncated. Unless truncation is enabled, oversize frames are dropped. The frame size is truncated to a multiple of 18 bytes—for example, if the MTU is 1000 bytes, the frame is truncated to 990 bytes (55 * 18 bytes.)	

The `no` form of the command removes the mirroring session and any mirroring source previously assigned to the session. To preserve the session while deleting a monitored source assigned to it, see the `no` command descriptions in [“Configure the monitored traffic in a mirror session”](#) (page 363).

Selecting all traffic on a port interface for mirroring according to traffic direction (CLI)

Syntax

```
[ no ] interface port/trunk/mesh monitor all [ in | out | both ] [ mirror 1 - 4 | name-str ] [ 1 - 4 | name-str 1 - 4 | name-str 1 - 4 | name-str ] [no-tag-added]
```

Assigns a mirroring source to a previously configured mirroring session on a source switch by specifying the port, trunk, and/or mesh sources to use, the direction of traffic to mirror, and the session.

<pre>interface <i>port/trunk/mesh</i></pre>	<p>Identifies the source ports, static trunks, and/or mesh on which to mirror traffic.</p> <p>Use a hyphen for a range of consecutive ports or trunks (a5-a8, Trk2-Trk4.)</p> <p>Use a comma to separate non-contiguous interfaces (b11, b14, Trk4, Trk7.)</p>
<pre>monitor all [in out both]</pre>	<p>For the interface specified by <i>port/trunk/mesh</i>, selects traffic to mirror based on whether the traffic is entering or leaving the switch on the interface:</p> <ul style="list-style-type: none"> • <code>in</code>: Mirrors entering traffic. • <code>out</code>: Mirrors exiting traffic. • <code>both</code>: Mirrors traffic entering and exiting. <p>If you enter the <code>monitor all</code> command without selection criteria or a session identifier, the command applies by default to session 1</p>
<pre>mirror [1 - 4 <i>name-str</i>]</pre>	<p>Assigns the traffic specified by the interface and direction to a session by number or—if configured—by name. The session must have been previously configured, as described in “Configure a mirroring session on the source switch” (page 362).</p> <p>Depending on how many sessions are already configured on the switch, you can use the same command to assign the specified source to up to four sessions, for example, <code>interface a1 monitor all in mirror 1 2 4</code>.</p> <p>For limits on configuring mirroring sources in a session, see “Mirroring-source restrictions” (page 363).</p> <ul style="list-style-type: none"> • <code>1 - 4</code>: Configures the selected port traffic to be mirrored in the specified session number. • <code>[name <i>name-str</i>]</code>: Optional; configures the selected port traffic to be mirrored in the specified session name. The string can be used interchangeably with the session number when using this command to assign a mirroring source to a session. To configure an alphanumeric name for a mirror session, see the command description in “Configure a source switch in a remote mirroring session” (page 363).
<pre>[no-tag-added]</pre>	<p>Prevents a VLAN tag from being added to the mirrored copy of an outbound packet sent to a local or remote mirroring destination. For more information, see “Untagged mirrored packets” (page 364).</p>

The `no` form of the command removes a mirroring source assigned to the session, but does not remove the session itself. This enables you to repurpose a session by removing an unwanted mirroring source and adding another in its place.

Selecting all traffic on a VLAN interface for mirroring according to traffic direction (CLI)

Syntax

```
vlan vid-# monitor all [ in | out | both ] [ mirror 1 - 4 |
name-str ] [ 1 - 4 | name-str 1 - 4 | name-str 1
- 4 | name-str ]
```

This command assigns a monitored VLAN source to a previously configured mirroring session on a source switch by specifying the VLAN ID, the direction of traffic to mirror, and the session.

<code>vlan vid-#</code>	Identifies the VLAN on which to mirror traffic.
<code>monitor all [in out both]</code>	Uses the direction of traffic on the specified <code>vid-#</code> to select traffic to mirror. See the syntax description on (page 339) . (If you enter the <code>monitor all</code> command without selection criteria or a session identifier, the command applies by default to session 1.)
<code>mirror [1 - 4 name-str]</code>	<p>Assigns the VLAN traffic defined by the VLAN ID and traffic direction to a session number or name. (The session must have been previously configured as described in “Configure a mirroring session on the source switch” (page 362).)</p> <p>Depending on how many sessions are already configured on the switch, you can use the same command to assign the specified VLAN source to up to four sessions, for example, <code>interface a1 monitor all in mirror 1 2 4</code>. For limits on configuring mirroring sources in a session, see “Mirroring-source restrictions” (page 363).</p> <ul style="list-style-type: none"> <code>1 - 4</code> : Configures the selected VLAN traffic to be mirrored in the specified session number. <code>[name name-str]</code> : Optional; configures the selected port traffic to be mirrored in the specified session name. The string can be used interchangeably with the session number when using this command to assign a mirroring source to a session. To configure an alphanumeric name for a mirroring session, see the command description under “Configuring a source switch in a remote mirroring session (CLI)” (page 337).

Assigning a VLAN to a mirroring session precludes assigning any other mirroring sources to the same session. If a VLAN is already assigned to a given mirroring session, using this command to assign another VLAN to the same mirroring session results in the second assignment replacing the first. Also, if there are other (port, trunk, or mesh) mirroring sources already assigned to a session, the switch displays a message similar to:

```
Mirror source port exists on session N. Can not add mirror
source VLAN.
```

The `no` form of the command removes a mirroring source assigned to the session, but does not remove the session itself. This allows you to repurpose a session by removing an unwanted mirroring source and adding another in its place.

Configuring a MAC address to filter mirrored traffic on an interface (CLI)

For more information, see [“About selecting inbound/outbound traffic using a MAC address” \(page 366\)](#).

Enter the `monitor mac mirror` command at the global configuration level.

Syntax

```
[ no ] monitor mac mac-addr [ src | dest | both ] mirror 1 - 4 | name-str [ 1 - 4 | name-str ] [ 1 - 4 | name-str ] [ 1 - 4 | name-str ]
```

Use this command to configure a source and/or destination MAC address as criteria for selecting traffic in one or more mirroring sessions on the switch. The MAC address you enter is configured to mirror inbound (`src`), outbound (`dest`), or both inbound and outbound (`both`) traffic on any port or learned VLAN on the switch.

<code>monitor mac mac-addr</code> Configures the MAC address as selection criteria for mirroring traffic on any port or learned VLAN on the switch.	
<code>src dest both</code>	Specifies how the MAC address is used to filter and mirror packets in inbound and/or outbound traffic on the interfaces on which the mirroring session is applied: <ul style="list-style-type: none"><code>src</code>: Mirrors all packets in inbound traffic that contain the specified MAC address as source address.<code>dest</code>: Mirrors all packets in outbound traffic that contain the specified MAC address as destination address. NOTE: The MAC address of the switch is not supported as either the source or destination MAC address used to select mirrored traffic. <ul style="list-style-type: none"><code>both</code>: Mirrors all packets in both inbound and outbound traffic that contain the specified MAC address as either source or destination address.
<code>mirror [1 - 4 name-str]</code>	Assigns the inbound and/or outbound traffic filtered by the specified MAC address to a previously configured mirroring session. The session is identified by a number or (if configured) a name. Depending on how many sessions are configured on the switch, you can use the same command to configure a MAC address as mirroring criteria in up to four sessions. To identify a session, you can enter either its name or number; for example: <code>mirror 1 2 3 traffsrc4</code> For the restrictions on how many mirroring source criteria you can configure in the same session, see "Mirroring-source restrictions" (page 363) . <code>1 - 4</code> : Specifies a mirroring session by number, for which the configured MAC address is used to select and mirror inbound and/or outbound traffic.

Packets that are sent or received on an interface configured with a mirroring session and that contain the MAC address as source and/or destination address are mirrored to a previously configured destination device.

To remove a MAC address as selection criteria in a mirroring session, you must enter the complete command syntax, for example, `no monitor mac 998877-665544 dest mirror 4`.

The `no` form of the command removes the MAC address as a mirroring criteria from an active session, but does not remove the session itself. This enables you to repurpose a session by removing an unwanted mirroring criteria and adding another in its place.

Configuring classifier-based mirroring (CLI)

For more information and a list of general steps for the process beginning with this command, see “Classifier-based mirroring configuration” (page 368). For information about restrictions on classifier-based mirroring, see “Classifier-based mirroring restrictions” (page 369).

Context: Global configuration

Syntax

```
[ no ] class [ ipv4 | ipv6 classname ]
```

Defines the name of a traffic class and specifies whether a policy is to be applied to IPv4 or IPv6 packets, where *classname* is a text string (64 characters maximum.)

After you enter the `class` command, you enter the class configuration context to specify match criteria. A traffic class contains a series of `match` and `ignore` commands, which specify the criteria used to classify packets.

To configure a default traffic class, use the `default-class` command as described below. A default class manages the packets that do not match the `match/ignore` criteria in any other classes in a policy.

Context: Class configuration

Syntax

```
[ no ] [seq-number] [ match | ignore ip-protocol  
source-address destination-address ] [ip-dscp codepoint]  
[precedence precedence-value] [tos tos-value] [vlan vlan-id]
```

For detailed information about how to enter `match` and `ignore` commands to configure a traffic class, the *Advanced Traffic Management Guide*.

Context: Global configuration

Syntax

```
[ no ] policy mirror policy-name
```

Defines the name of a mirroring policy and enters the policy configuration context. (For more information, see “Classifier-based mirroring configuration” (page 368).)

A traffic policy consists of one or more classes and one or more mirroring actions configured for each class of traffic. The configured actions are executed on packets that match a `match` statement in a class. No policy action is performed on packets that match an `ignore` statement.

Context: Policy configuration

Syntax

```
[ no ] [seq-number] class [ ipv4 | ipv6 classname ]  
action mirror session
```

Defines the mirroring action to be applied on a pre-configured IPv4 or IPv6 traffic class when a packet matches the `match` criteria in the traffic class. You can enter multiple `class action mirror` statements in a policy.

<code>[<i>seq-number</i>]</code>	The (optional) <code>seq-number</code> parameter sequentially orders the mirroring actions that you enter in a policy configuration. Actions are executed on matching packets in numerical order.
----------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

	Default: Mirroring action statements are numbered in increments of 10, starting at 10.
<code>class [ipv4 ipv6 class-name]</code>	Defines the preconfigured traffic class on which the mirroring actions in the policy are executed and specifies whether the mirroring policy is applied to IPv4 or IPv6 traffic in the class. The classname is a text string (64 characters maximum.)
<code>action mirror session</code>	Configures mirroring for the destination and session specified by the <code>session</code> parameter.

Context: Policy configuration

Syntax

```
[ no ] default-class action mirror session [action mirror  
session ]..
```

Configures a default class that allows packets that are not matched nor ignored by any of the class configurations in a mirroring policy to be mirrored to the destination configured for the specified session.

Applying a mirroring policy on a port or VLAN interface

Enter one of the following `service-policy` commands from the global configuration context. (For more information, see [“Classifier-based mirroring configuration” \(page 368\)](#).)

Context: Global configuration

Syntax

```
interface port-list service-policy policy-name in
```

Configures the specified ports with a mirroring policy that is applied to inbound traffic on each interface.

Separate individual port numbers in a series with a comma, for example, a1, b4, d3. Enter a range of ports by using a dash, for example, a1-a5.

The mirroring policy name you enter must be the same as the policy name you configured with the `policy mirror` command in the syntax [\(page 342\)](#).

Syntax

```
vlan vlan-id service-policy policy-name in
```

Configures a mirroring policy on the specified VLAN that is applied to inbound traffic on the VLAN interface.

Valid VLAN ID numbers range from 1 to 4094.

The mirroring policy name you enter must be the same as the policy name you configured with the `policy mirror` command in the syntax [\(page 330\)](#).

For more information about how to apply a mirroring policy to an interface, see the *Advanced Traffic Management Guide*.

Viewing a classifier-based mirroring configuration (CLI)

To display information about a classifier-based mirroring configuration or statistics on one or more mirroring policies, enter one of the following commands:

- `show class [ipv4 class-name | ipv6 class-name | config]`
- `show policy [policy-name | config]`

- show policy resources
- show statistics policy [policy-name] [interface port-num | vlan vid in]

For examples of classifier-based show command output, see “Classifier-based mirroring configuration” (page 368).

Viewing all mirroring sessions configured on the switch (CLI)

Syntax

```
show monitor
```

If a monitored source for a remote session is configured on the switch, the following information is displayed. Otherwise, the output displays: **Mirroring is currently disabled**.

Sessions	Lists the four configurable sessions on the switch.
Status	Displays the current status of each session : <ul style="list-style-type: none"> • active: The session is configured. • inactive: Only the destination has been configured; the mirroring source is not configured. • not defined: Mirroring is not configured for this session.
Type	Indicates whether the mirroring session is local (port), remote (IPv4), or MAC-based (mac) for local or remote sessions.
Sources	Indicates how many monitored source interfaces are configured for each mirroring session.
Policy	Indicates whether the source is using a classifier-based mirroring policy to select inbound IPv4 or IPv6 traffic for mirroring.

If a remote mirroring endpoint is configured on the switch, the following information is displayed. Otherwise, the output displays: **There are no Remote Mirroring endpoints currently assigned**.

Type	Indicates whether the mirroring session is local (port), remote (IPv4), or MAC-based (mac) for local or remote sessions.
UDP Source Addr	The IP address configured for the source VLAN or subnet on which the monitored source interface exists. In the configuration of a remote session, the same UDP source address must be configured on the source and destination switches.
UDP port	The unique UDP port number that identifies a remote session. In the configuration of a remote session, the same UDP port number must be configured on the source and destination switches.
UDP Dest Addr	The IP address configured for the destination VLAN or subnet on which the remote exit port exists. In the configuration of a remote session, the same UDP destination address must be configured on the source and destination switches.
Dest Port	Identifies the exit port for a remote session on a remote destination switch.

Figure 188 Displaying the currently configured mirroring sessions on the switch

HP Switch# show monitor						Local and Remote Mirroring Sources: <ul style="list-style-type: none"> • Session 1 is performing local mirroring using a classifier-based policy as traffic-selection criteria. • Session 2 is performing remote mirroring using MAC-based traffic-selection criteria. • Session 3 is not configured. • Session 4 is configured for remote mirroring from a non-policy source (for example, traffic direction), but is currently not mirroring any traffic.
Network Monitoring						
Sessions	Status	Type	Sources	Policy		
1	active	port	1	yes		
2	active	mac	2	no		
3	not defined					
4	inactive	IPv4	0	no		
Remote Mirroring - Remote Endpoints						
Type	UDP Source Addr	UDP port	UDP Dest Addr			
IPv4	10.10.30.1	7950	10.10.20.1			

Viewing the remote endpoints configured on the switch (CLI)

Syntax

```
show monitor endpoint
```

Displays the remote mirroring endpoints configured on the switch. Information on local sessions configured on the switch is not displayed. (To view the configuration of a local session, use the

```
show monitor [ 1-4 | name name-str ] ]
```

command, as described on page 74 and page 77.)

Type	Indicates whether the session is a <code>port</code> (local) or <code>IPv4</code> (remote) mirroring session.
UDP Source Addr	The IP address configured for the source VLAN or subnet on which the monitored source interface exists. In the configuration of a remote session, the same UDP source address must be configured on the source and destination switches.
UDP port	The unique UDP port number that identifies a remote session. In the configuration of a remote session, the same UDP port number must be configured on the source and destination switches.
UDP Dest Addr	The IP address configured for the destination VLAN or subnet on which the remote exit port exists. In the configuration of a remote session, the same UDP destination address must be configured on the source and destination switches.
Dest Port	ifies the exit port for a remote session on a remote destination switch.

Example


In [Figure 189 \(page 346\)](#), the `show monitor endpoint` output shows that the switch is configured as the remote endpoint (destination) for two remote sessions from the same monitored source interface.

Figure 189 Displaying the configuration of remote mirroring endpoints on the switch

```

HP Switch(config)# show monitor endpoint
Remote Mirroring - Remote Endpoints

Type   UDP Source Addr  UDP port  UDP Dest Addr  Dest Port
-----
IPv4   10.10.10.1       8001      10.10.30.2     4
IPv4   10.10.10.1       8003      10.10.30.2     5
    
```


These two sessions monitor traffic from the same source switch, but use different UDP port numbers.

Viewing the mirroring configuration for a specific session (CLI)

Syntax

```
show monitor [ 1 - 4 | name name-str ]
```

Displays detailed configuration information for a specified local or remote mirroring session on a source switch.

Session	Displays the number of the specified session.
Session Name	Displays the name of the session, if configured.
Policy	Indicates whether the source is using a classifier-based mirroring policy to select inbound IPv4 or IPv6 traffic for mirroring.
Mirroring Destination	For a local mirroring session, displays the port configured as the exit port on the source switch. For a remote mirroring session, displays IPv4, which indicates mirroring to a remote (endpoint) switch.
UDP Source Addr	The IP address configured for the source VLAN or subnet on which the monitored source interface exists. In the configuration of a remote session, the same UDP source address must be configured on the source and destination switches.
UDP port	The unique UDP port number that identifies a remote session. In the configuration of a remote session, the same UDP port number must be configured on the source and destination switches.
UDP Dest Addr	The IP address configured for the destination VLAN or subnet on which the remote exit port exists. In the configuration of a remote session, the same UDP destination address must be configured on the source and destination switches.
Status	For a remote session, displays current session activity: <ul style="list-style-type: none"> • active: The session is configured and is mirroring traffic. A remote path has been discovered to the destination. • inactive: The session is configured, but is not currently mirroring traffic. A remote path has <i>not</i> been discovered to the destination. • not defined: Mirroring is not configured for this session.
Monitoring Sources	For the specified local or remote session, displays the source (port, trunk, or VLAN) interface and the MAC address (if configured) used to select mirrored traffic.
Direction	For the selected interface, indicates whether mirrored traffic is entering the switch (<i>in</i>), leaving the switch (<i>out</i>), or <i>both</i> .

Viewing a remote mirroring session (CLI)

After you configure session 2 for remote mirroring (Figure 190 (page 347)), you can enter the `show monitor 2` command to verify the configuration (Figure 191 (page 347)).

Figure 190 Configuring a remote mirroring session and monitored source

```
HP Switch(config)# mirror 2 name test-10 remote ip 10.10.10.1 8010 10.10.30.2
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
HP Switch(config)# interface b1 monitor all both mirror 2
```

Figure 191 Displaying the Configuration of a Remote Mirroring Session

```
HP Switch(config)# show monitor 2
Network Monitoring

Session: 2      Session Name: test-10
Policy: no policy relationship exists

Mirror Destination: IPv4
  UDP Source Addr  UDP port  UDP Dest Addr  Status
-----
  10.10.10.1      8010     10.10.30.2    active

Monitoring Sources  Direction
-----
Port: B1            Both
```

If no monitored (source) interface is configured for a mirroring session, no information is displayed in the Monitoring Sources and Direction columns.

Viewing a MAC-based mirroring session (CLI)

After you configure a MAC-based mirroring session (Figure 192 (page 347)), you can enter the `show monitor 3` command to verify the configuration (Figure 193 (page 347).)

Figure 192 Configuring a MAC-based mirroring session

```
HP Switch(config)# mirror 3 port a1
HP Switch# monitor mac 112233-445566 src mirror 3
```

Figure 193 Displaying a MAC-based mirroring session

```
HP Switch(config)# show monitor 3
Network Monitoring

Session: 3      Session Name:
Policy: no policy relationship exists

Mirror Destination: A1 (Port)

Monitoring Sources  Direction
-----
MAC: 112233-445566 Source
```

The MAC address used to select packets in a local mirroring session is displayed in these columns.

Viewing a local mirroring session (CLI)

When used to display the configuration of a local session, the `show monitor` command displays a subset of the information displayed for a remote mirroring session.

Example

Figure 194 (page 348) displays a local mirroring configuration for a session configured as follows:

- Session number: 1
- Session name: Detail
- Classifier-based mirroring policy, "MirrorAdminTraffic", is used to select inbound traffic on port B1.
- Mirrored traffic is sent to exit port B3.

Figure 194 Displaying the configuration of a local mirroring session

```
HP Switch(config)# show monitor 1
Network Monitoring

Session: 1      Session Name: Detail
Policy: MirrorAdminTraffic

Mirror Destination: B3      (Port)

Monitoring Sources  Direction
-----
Port: B1           In
```

Viewing information on a classifier-based mirroring session (CLI)

In the following example, a classifier-based mirroring policy (`mirrorAdminTraffic`) mirrors selected inbound IPv4 packets on VLAN 5 to the destination device configured for mirroring session 3.

Figure 195 Configuring a classifier-based mirroring policy in a local mirroring session

```
HP Switch(config)# mirror 3 port c1
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
HP Switch(config)# class ipv4 AdminTraffic
HP Switch(config-class)# match ip 15.29.61.1 0.63.255.255 0.0.0.0
255.255.255.255
HP Switch(config-class)# match ip 0.0.0.0 255.255.255.255 15.29.61.1
0.63.255.255
HP Switch(config-class)# exit
HP Switch(config)# policy mirror MirrorAdminTraffic
HP Switch(config-policy)# class ipv4 AdminTraffic action mirror 3
HP Switch(config-policy)# exit
HP Switch(config)# vlan 5 service-policy MirrorAdminTraffic in
```

Example 64 Displaying a classifier-based policy in a local mirroring session

```
HP Switch(config)# show monitor 3

Network Monitoring

Session: 3      Session Name:
Policy: MirrorAdminTraffic

Mirror Destination: C1      (Port)

Monitoring Sources  Direction
-----
VLAN: 5           Source
```

Viewing information about a classifier-based mirroring configuration (CLI)

Syntax

```
show class ipv4 classname
show class ipv6 classname
show class config
```

<code>ipv4 <i>classname</i></code>	Lists the statements that make up the IPv4 class identified by <i>classname</i> .
<code>ipv6 <i>classname</i></code>	Lists the statements that make up the IPv6 class identified by <i>classname</i> .
<code>config</code>	Displays all classes, both IPv4 and IPv6, and lists the statements that make up each class.

Additional variants of the `show class ...` command provide information on classes that are members of policies that have been applied to ports or VLANs.

Figure 196 `show class` output for a mirroring policy

```
HP Switch(config)# show class ipv4 AdminTraffic
Statements for Class ipv4 "AdminTraffic"
 10 match ip 15.29.16.1 0.63.255.255 0.0.0.0 255.255.255.255
 20 match ip 0.0.0.0 255.255.255.255 15.29.16.1 0.63.255.255
```

Viewing information about a classifier-based mirroring configuration (CLI)

Syntax

```
show policy policy-name
show policy config
```

<code>policy-name</code>	Lists the statements that make up the specified policy.
<code>config</code>	Displays the names of all policies defined for the switch and lists the statements that make up each policy.

Additional variants of the `show policy...` command provide information on policies that have been applied to ports or VLANs.

Figure 197 `show policy` output for a mirroring policy

```
HP Switch(config)# show policy MirrorAdminTraffic
Statements for Policy "MirrorAdminTraffic"
 10 class ipv4 "AdminTraffic" action mirror 3
```

Viewing information about statistics on one or more mirroring policies (CLI)

Syntax

```
[ show | clear ]statistics policy policy-name port port-num
[ show | clear ]statistics policy policy-name vlan vid in
```

<code>show</code>	Displays the statistics for a specified policy applied to a specified port or VLAN.
<code>clear</code>	Clears statistics for the specified policy and port or VLAN.
<code>policy-name</code>	The name of the policy.
<code>port-num</code>	The number of the port on which the policy is applied (single port only, not a range.)
<code>vid</code>	The number or name of the vlan on which the policy is applied. VLAN ID numbers range from 1 to 4094.
<code>in</code>	Indicates that statistics are shown for inbound traffic only.

Figure 198 (page 350) shows the number of packets (in parentheses) that have been mirrored for each match/ignore statement in the mirroring policy.

Figure 198 show statistics policy output for a mirroring policy

```
HP Switch# show statistics policy MirrorAdminTraffic vlan 30 in
HitCounts for Policy MirrorAdminTraffic
10 class ipv4 "AdminTraffic" action mirror 3
(5244) 10 match ip 15.29.16.1 0.63.255.255 0.0.0.0 255.255.255.255
(9466) 20 match ip 0.0.0.0 255.255.255.255 15.29.16.1 0.63.255.255
```

Viewing resource usage for mirroring policies (CLI)

Syntax

```
show policy resources
```

Displays the number of hardware resources (rules, meters, and application port ranges) used by classifier-based mirroring policies (local and remote) that are currently applied to interfaces on the switch, as well as QoS policies and other software features.

NOTE: The information displayed is the same as the output of the `show qos resources` and `show access-list resources` commands.

For a detailed explanation of the information displayed with the `show [qos | access-list | policy resources]` command, see the [“Displaying current resource usage” \(page 466\)](#) section in the [“Monitoring Resources” \(page 466\)](#) appendix.

Figure 199 Displaying the hardware resources used by currently configured mirroring policies

```

HP Switch# show policy resources

Resource usage in Policy Enforcement Engine
-----
Ports | Rules Available | Rules Used | ACL | QoS | IDM | VT | Mirror | Other |
-----+-----+-----+-----+-----+-----+-----+-----+-----+
1-24 | 3014 | 15 | 11 | 0 | 1 | 0 | 3 |
25-48 | 3005 | 15 | 10 | 10 | 1 | 0 | 3 |
A | 3017 | 15 | 8 | 0 | 1 | 0 | 3 |

Ports | Meters Available | Meters Used | ACL | QoS | IDM | VT | Mirror | Other |
-----+-----+-----+-----+-----+-----+-----+-----+-----+
1-24 | 250 | 5 | 0 | 0 | 0 | 0 | 0 |
25-48 | 251 | 4 | 0 | 0 | 0 | 0 | 0 |
A | 253 | 2 | 0 | 0 | 0 | 0 | 0 |

Ports | Application Port Ranges Used | ACL | QoS | IDM | VT | Mirror | Other |
-----+-----+-----+-----+-----+-----+-----+-----+-----+
1-24 | 3014 | 2 | 0 | 0 | 0 | 0 | 0 |
25-48 | 3005 | 2 | 0 | 0 | 0 | 0 | 0 |
A | 3017 | 2 | 0 | 0 | 0 | 0 | 0 |

0 of 8 Policy Engine management resources used.
Key:
ACL = Access Control Lists
QoS = Device & Application Port Priority, QoS Policies, ICMP rate limits
IDM = Identity Driven Management
VT = Virus Throttling blocks
Mirror = Mirror Policies, Remote Intelligent Mirror endpoints
Other = Management VLAN, DHCP Snooping, ARP Protection, Jumbo IP-MTU.

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.

```

Includes the hardware resources used by classifier-based local and remote mirroring policies that are currently applied to interfaces on the switch.

Viewing the mirroring configurations in the running configuration file (CLI)

Use the `show run` command to view the current mirroring configurations on the switch. In the `show run` command output, information about mirroring sources in configured sessions begins with the `mirror` keyword; monitored source interfaces are listed per-interface.

Example

Figure 200 Displaying mirroring sources and sessions in the running configurations

```

HP Switch(config)# show run
Running configuration:
; J8697A Configuration Editor; Created on release #K.12.XX
max-vlans 300
ip access-list extended "100"
  10 permit icmp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 0
  exit
  no ip address
  exit
. . .
mirror 1 port B3
mirror 2 name "test-10" remote ip 10.10.10.1 8010 10.10.30.2
. . .
interface B1
  monitor ip access-group "100" In mirror 1
  monitor all Both mirror 2
  exit
. . .

```

Mirroring sessions with exit ports configured on the switch: B3 is an exit port for a local session; session 2 has a remote destination and exit port.

Selection criteria used to monitor traffic on port B1 for mirroring sessions 1 (ACL-based) and 2 (direction-based)

Information about remote endpoints configured for remote sessions on the switch begin with the `mirror endpoint` keywords. In the following example, two remote sessions use the same exit port:

Figure 201 Displaying remote mirroring endpoints in the running configuration


```

HP Switch(config)# show run
Running configuration:
; J8693A Configuration Editor; Created on release #K.12.XX
module 3 type J8694A
. . .

mirror endpoint ip 10.10.20.1 8010 10.10.30.2 port 4
mirror endpoint ip 10.10.51.10 7955 10.10.30.2 port 4
. . .

```

Remote endpoints configured on the switch, including source IP address, UDP port number, destination IP address, and remote exit port. Each remote session is identified by a unique UDP port number.



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 (page B-6 (page 314).)
- **Counters**
Display details of traffic volume on individual ports (“Accessing port and trunk statistics (Menu)” (page 321).)
- **Event Log**
Lists switch operating events (“Using the Event Log for troubleshooting switch problems” (page 400) .)
- **Configurable trap receivers**
Uses SNMP to enable management stations on your network to receive SNMP traps from the switch.
- **Port monitoring (mirroring)**
Copy all traffic from the specified ports to a designated monitoring port .

NOTE: Link test and ping test—analysis tools in troubleshooting situations—are described in Appendix C, “Troubleshooting” (page 379). See “Diagnostic tools” (page 439).

Compatibility mode

Table 22 (page 352) shows how the v2 zl and zl modules behave in various combinations and situations when Compatibility mode is enabled and when it is disabled.

Table 22 Compatibility mode enabled/disabled comparisons

Modules	Compatibility mode enabled	Compatibility mode disabled
v2 zl modules only	Can insert zl module and the module will come up. Any v2 zl modules are limited to the zl configuration capacities.	v2 zl modules are at full capacity. ZL modules are not allowed to power up.
Mixed v2 zl and zl modules	Can insert zl module and the module will come up. Any v2 zl modules are limited to the zl configuration capacities.	ZL modules are not allowed to power up.

Table 22 Compatibility mode enabled/disabled comparisons *(continued)*

Modules	Compatibility mode enabled	Compatibility mode disabled
	If compatibility mode is disabled, the zl modules go down.	
ZL modules only	Same as exists already. If a v2 zl module is inserted, it operates in the same mode as the zl module, but with performance increases.	The Management Module is the only module that powers up.
	In Compatibility Mode, no v2 zl features are allowed, whether the modules are all v2 zl or not.	If Compatibility Mode is disabled and then enabled, the startup config is erased and the chassis reboots.

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 [\(page 353\)](#), 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.

About traffic mirroring

Starting in software release K.12.xx, traffic mirroring (Intelligent Mirroring) allows you to mirror (send a copy of) network traffic received or transmitted on a switch interface to a local or remote destination, such as a traffic analyzer or IDS.)

Traffic mirroring provides the following benefits:

- Allows you to monitor the traffic flow on specific source interfaces.
- Helps in analyzing and debugging problems in network operation resulting from a misbehaving network or an individual client. The mirroring of selected traffic to an external device makes it easier to diagnose a network problem from a centralized location in a topology spread across a campus.
- Supports remote mirroring to simultaneously mirror switch traffic on one or more interfaces to multiple remote destinations. (In remote mirroring, you must first configure the remote mirroring endpoint—remote switch and exit port—before you specify a mirroring source for a session.)

Mirroring

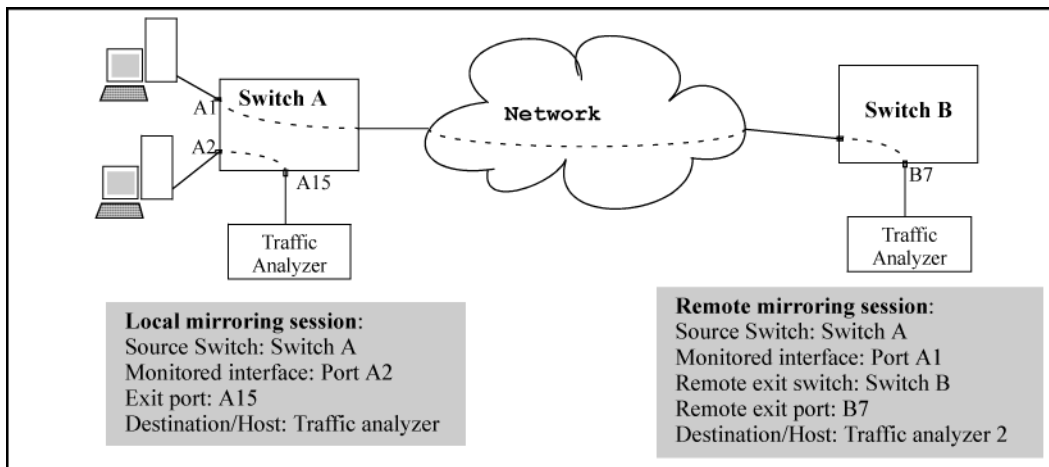
Figure 202 (page 354) shows an example of the terms used to describe the configuration of a sample local and remote mirroring session:

- In the local session, inbound traffic entering Switch A is monitored on port A2 and mirrored to a destination (host), traffic analyzer 1, through exit port A15 on the switch.
A local mirroring session means that the monitored interface (A2) and exit port (A15) are on the same switch.
- In the remote session, inbound traffic entering Switch A is monitored on port A1. A mirrored copy of monitored traffic is routed through the network to a remote mirroring endpoint: exit port B7 on Switch B. A destination device, traffic analyzer 2, is connected to the remote exit port.

A remote mirroring session means that:

- The monitored interface (A1) and exit port (B7) are on different switches.
- Mirrored traffic can be bridged or routed from a source switch to a remote switch.

Figure 202 Local and remote sessions showing mirroring terms



Mirroring destinations

Traffic mirroring supports destination devices that are connected to the local switch or to a remote switch:

- Traffic can be copied to a destination (host) device connected to the same switch as the mirroring source in a *local* mirroring session. You can configure up to *four* exit ports to which destination devices are connected.
- Traffic can be bridged or routed to a destination device connected to a different switch in a *remote* mirroring session. You can configure up to 32 remote mirroring endpoints (IP address and exit port) to which destination devices are connected.

Mirroring sources and sessions

Traffic mirroring supports the configuration of port and VLAN interfaces as mirroring sources in up to *four* mirroring sessions on a switch. Each session can have one or more sources (ports and/or static trunks, a mesh, or a VLAN interface) that monitor traffic entering and/or leaving the switch.

NOTE: Using the CLI, you can make full use of the switch's local and remote mirroring capabilities. Using the Menu interface, you can configure only local mirroring for either a single VLAN or a group of ports, static trunks, or both.

In remote mirroring, a 54-byte remote mirroring tunnel header is added to the front of each mirrored frame for transport from the source switch to the destination switch. This may cause some frames that were close to the MTU size to exceed the MTU size. Mirrored frames exceeding the allowed MTU size are dropped, unless the optional `[truncation]` parameter is set in the `mirror` command. For more information, including the size limitation for jumbo and non-jumbo frames, see [“Maximum supported frame size” \(page 375\)](#). For information on the `[truncation]` parameter, see [“Configure a source switch in a remote mirroring session” \(page 363\)](#).

Mirroring sessions

A mirroring session consists of a mirroring source and destination (endpoint.) Although a mirroring source can be one of several interfaces, as mentioned above, for any session, the destination must be a single (exit) port. The exit port cannot be a trunk, VLAN, or mesh interface.

You can map multiple mirroring sessions to the same exit port, which provides flexibility in distributing hosts, such as traffic analyzers or an IDS. In a remote mirroring endpoint, the IP address of the exit port and the remote destination switch can belong to different VLANs.

Mirroring sessions can have the same or a different destination. You can configure an exit port on the local (source) switch and/or on a remote switch as the destination in a mirroring session. When configuring a mirroring destination, consider the following options:

- Mirrored traffic belonging to different sessions can be directed to the same destination or to different destinations.
- You can reduce the risk of oversubscribing a single exit port by:
 - Directing traffic from different session sources to multiple exit ports.
 - Configuring an exit port with a higher bandwidth than the monitored source port.
- You can segregate traffic by type, direction, or source.

Mirroring session limits

A switch running software release K.12.xx or greater supports the following:

- A maximum of four mirroring (local and remote) sessions.
- A maximum of 32 remote mirroring endpoints (exit ports connected to a destination device that receive mirrored traffic originating from monitored interfaces on a different switch.)

Selecting mirrored traffic

You can use any of the following options to select the traffic to be mirrored on a port, trunk, mesh, or VLAN interface in a local or remote session:

- **All traffic**
Monitors all traffic entering or leaving the switch on one or more interfaces (inbound and outbound.)
- **Direction-based traffic selection**
Monitors traffic that is either entering or leaving the switch (*inbound or outbound.*) Monitoring traffic in only one direction improves operation by reducing the amount of traffic sent to a mirroring destination.

- **MAC-based traffic selection**

Monitors only traffic with a matching source and/or destination MAC address in packet headers entering and/or leaving the switch on one or more interfaces (*inbound and/or outbound*.)

- **Classifier-based service policy**

Provides a finer granularity of match criteria to zoom in on a subset of a monitored port or VLAN traffic (IPv4 or IPv6) and select it for local or remote mirroring (*inbound only*.)

Deprecation of ACL-based traffic selection

In software release **K.14.01 or greater**, the use of ACLs for selecting traffic in a mirroring session has been deprecated and is replaced by the use of advanced classifier-based service policies (see [“About selecting inbound traffic using advanced classifier-based mirroring”](#) (page 367).)

As with ACL criteria, classifier-based match/ignore criteria allow you to limit a mirroring session to selected inbound packets on a given port or VLAN interface (instead of mirroring all inbound traffic on the interface.)

The following commands have been deprecated:

- `interface port/trunk/mesh monitor ip access-group acl-name in mirror [1 - 4 | name-str]`
- `vlan vid-# monitor ip access-group acl-name in mirror [1 - 4 | name-str]`

After you install and boot release K.14.01 or greater, ACL-based local and remote mirroring sessions configured on a port or VLAN interface are automatically converted to classifier-based mirroring policies. For more information, see [“Migration to release K.14.01 or greater”](#) (page 359).

If you are running software release **K.13.xx** or earlier, ACL permit/deny criteria are supported to select IP traffic entering a switch to mirror in a local or remote session, using specified source and/or destination criteria.

Mirrored traffic destinations

Local destinations

A local mirroring traffic destination is a port on the same switch as the source of the traffic being mirrored.

Remote destinations

A *remote* mirroring traffic destination is an HP switch configured to operate as the exit switch for mirrored traffic sessions originating on other HP switches. As of June, 2007, switches capable of this operation include the following HP switches:

- 3500yl
- 5400zl
- 6200yl
- 8200zl



CAUTION:

After you configure a mirroring session with traffic-selection criteria and a destination, the switch immediately starts to mirror traffic to each destination device connected to an exit port. In a remote mirroring session that uses IPv4 encapsulation, if the intended exit switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic. For this reason, HP Switch strongly recommends that you configure the exit switch for a remote mirroring session before configuring the source switch for the same session.

Monitored traffic sources

You can configure mirroring for traffic entering or leaving the switch on:

- **Ports and static trunks**
Provides the flexibility for mirroring on individual ports, groups of ports, static port trunks, or any combination of these..
- **Meshed ports**
Enables traffic mirroring on all ports configured for meshing on the switch.
- **Static VLANs**
Supports traffic mirroring on static VLANs configured on the switch. This option enables easy mirroring of traffic from all ports on a VLAN. It automatically adjusts mirroring to include traffic from newly added ports and to exclude traffic from ports removed from the VLAN.

Criteria for selecting mirrored traffic

On the monitored sources listed above, you can configure the following criteria to select the traffic you want to mirror:

- Direction of traffic movement (entering or leaving the switch, or both.)
- Type of IPv4 or IPv6 traffic entering the switch, as defined by a classifier-based service policy (see [“About selecting inbound traffic using advanced classifier-based mirroring” \(page 367\).](#))
In software release **K.14.01 or greater**, classifier-based service policies replace ACL-based traffic selection in mirroring sessions.
- Source and/or destination MAC addresses in packet headers.

Mirroring configuration

Table 23 (page 357) shows the different types of mirroring that you can configure using the CLI, Menu, and SNMP interfaces.

Table 23 Mirroring configuration options

Monitoring interface and configuration level	Traffic selection criteria	Traffic direction		
		CLI config	Menu and web i/f config ¹	Snmp config
VLAN	All traffic	Inbound only Outbound only Both directions	All traffic (inbound and outbound combined)	Inbound only Outbound only Both directions
	ACL (IP traffic) ²	See “About selecting inbound traffic using advanced classifier-based mirroring” (page 367).		
	Classifier-based policy (IPv4 or IPv6 traffic)	Inbound only	Not available	Not available
Port(s) Trunk(s) Mesh	All traffic	Inbound only Outbound only Both directions	All traffic (inbound and outbound combined)	Inbound only Outbound only Both directions
	ACL (IP traffic) ³	See “About selecting inbound traffic using advanced classifier-based mirroring” (page 367).		
	Classifier-based policy (IPv4 or IPv6 traffic)	Inbound only	Not available	Not available
Switch (global)	MAC source/destination address	Inbound only	Not available	Inbound only Outbound only

Table 23 Mirroring configuration options (continued)

Monitoring interface and configuration level	Traffic selection criteria	Traffic direction		
		CLI config	Menu and web i/f config ¹	Snmp config
		Outbound only Both directions		Both directions

¹ Configures only session 1, and only for local mirroring.

² In release K.14.01 and greater, the use of ACLs to select inbound traffic in a mirroring session (using the `[interface | vlan]monitor ip access-group in mirror` command) has been deprecated and is replaced with classifier-based mirroring policies.

³ In release K.14.01 and greater, the use of ACLs to select inbound traffic in a mirroring session (using the `[interface | vlan]monitor ip access-group in mirror` command) has been deprecated and is replaced with classifier-based mirroring policies.

Configuration notes

Using the CLI, you can configure all mirroring options on a switch.

Using the Menu, you can configure only session 1 and only local mirroring in session 1 for traffic in both directions on specified interfaces. (If session 1 has been already configured in the CLI for local mirroring for inbound-only or outbound-only traffic, and you use the Menu to modify the session 1 configuration, session 1 is *automatically* reconfigured to monitor both inbound and outbound traffic on the assigned interfaces. If session 1 has been configured in the CLI with a classifier-based mirroring policy or as a remote mirroring session, an error message is displayed if you try to use the Menu to configure the session.)

You can use the CLI can configure sessions 1 to 4 for local or remote mirroring in any combination, and override a Menu configuration of session 1.

You can also use SNMP configure sessions 1 to 4 for local or remote mirroring in any combination and override a Menu configuration of session 1, *except* that SNMP cannot be used to configure a classifier-based mirroring policy.

Remote mirroring endpoint and intermediate devices

The remote mirroring endpoint that is used in a remote mirroring session must be an HP switch that supports the mirroring functions described in this chapter. (A remote mirroring endpoint consists of the remote switch and exit port connected to a destination device.) Because remote mirroring on an HP switch uses IPv4 to encapsulate mirrored traffic sent to a remote endpoint switch, the intermediate switches and routers in a layer 2/3 domain can be from any vendor if they support IPv4.

The following restrictions apply to remote endpoint switches and intermediate devices in a network configured for traffic mirroring:

- The exit port for a mirroring destination must be an individual port and *not* a trunk, mesh, or VLAN interface.
- A switch mirrors traffic on static trunks, but *not* on dynamic LACP trunks.
- A switch mirrors traffic at line rate. When mirroring multiple interfaces in networks with high-traffic levels, it is possible to copy more traffic to a mirroring destination than the link supports. However, some mirrored traffic may not reach the destination. If you are mirroring a high-traffic volume, you can reduce the risk of oversubscribing a single exit port by:
 - Directing traffic from different session sources to multiple exit ports.
 - Configuring an exit port with a higher bandwidth than the monitored source port.

Migration to release K.12.xx

On a switch that is running a software release earlier than K.12.xx with one or more mirroring sessions configured, when you download and boot release K.12.xx, the existing mirroring configurations are managed as follows:

- A legacy mirroring configuration on a port or VLAN interface maps to session 1.
- Traffic-selection criteria for session 1 is set to `both`; both inbound and outbound traffic (traffic entering *and* leaving the switch) on the configured interface is selected for mirroring.
- In a legacy mirroring configuration, a local exit port is applied to session 1.

Booting from software versions earlier than K.12.xx

If it is necessary to boot the switch from a legacy (pre-K.12.xx) software version after using version K.12.xx or greater to configure mirroring, remove mirroring from the configuration before booting with the earlier software.

Maximum supported frame size

The IPv4 encapsulation of mirrored traffic adds a 54-byte header to each mirrored frame. If a resulting frame exceeds the MTU allowed in the path from the mirroring source to the mirroring destination, the frame is dropped, unless the optional `[truncation]` parameter is set in the `mirror` command. For more information, see [“Maximum supported frame size” \(page 375\)](#). For information on the `[truncation]` parameter, see [“Configure a source switch in a remote mirroring session” \(page 363\)](#).

Frame truncation

Mirroring does not truncate frames unless the `[truncation]` parameter in the `mirror` command is set. If that parameter is not set, oversized mirroring frames are dropped. Also, remote mirroring does not allow downstream devices in a mirroring path to fragment mirrored frames.

Migration to release K.14.01 or greater

NOTE: If a switch is running software release K.12.xx, you must first upgrade to release K.13.xx before migrating the switch to release K.14.01 or greater.

When you download and boot software release K.14.01 or greater on a switch that is running release K.13.xx and has one or more mirroring sessions configured, an ACL-based mirroring configuration on a port or VLAN interface is mapped to a class and policy configuration based on the ACL.

The new mirroring policy is automatically configured on the same port or VLAN interface on which the mirroring ACL was assigned. The behavior of the new class and mirroring-policy configuration exactly matches the traffic-selection criteria and mirroring destination used in the ACL-based session .)

[Figure 203 \(page 360\)](#) and [Figure 204 \(page 360\)](#) show how ACL-based selection criteria in a mirroring session are converted to a classifier-based policy and class configuration when you install release K.14.01 or greater on a switch.

Figure 203 Mirroring configuration in show run output in release K.13.xx

```
HP Switch(config)# show run
Running configuration:
. . .
ip access-list extended "100"
 10 permit icmp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 0
  exit
. . .
mirror 1 port C1
mirror 2 name "test-10" remote ip 10.10.10.1 8010 10.10.30.2
. . .
interface C1
  monitor ip access-group "100" In mirror 1
  exit
. . .
```

Configuration of ACL 100 that is used to select mirrored traffic in session 1

Existing mirror sessions configured on the switch for a local (port C1 in session 1) and remote (session 2) monitored interface

ACL-based traffic selection on monitored interface C1 in session 1

Figure 204 Mirroring configuration in show run output in release K.14.01 or greater

```
HP Switch(config)# show run
Running configuration:
. . .
mirror 1 port B3
mirror 2 name "test-10" remote ip 10.10.10.1 8010 10.10.30.2
. . .
class ipv4 "100MirrorClass"
 10 match icmp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 0
  exit
policy mirror "100MirrorPolicy"
 10 class ipv4 "100" action mirror 1
  exit
. . .
interface C1
  service-policy "100MirrorPolicy" In
  exit
. . .
```

After migration to release K.14.01 or greater, the existing mirroring configurations for sessions 1 (local) and 2 (remote) on the switch remain the same.

The traffic-selection criteria in ACL 100 (Figure B-B-27) applied to inbound traffic on port C1 in session 1 are converted to a class and policy configuration with the names, "100MirrorClass" and "100MirrorPolicy", which are applied to inbound traffic on port C1 in session 1 with the **service-policy** command.

Using the Menu to configure local mirroring

Menu and WebAgent limits

You can use the Menu and WebAgent to quickly configure or reconfigure local mirroring on session 1 and allow one of the following two mirroring source options:

- Any combination of source ports, trunks, and a mesh.
- One static, source VLAN interface.

The Menu and WebAgent also has these limits:

- Configure and display only session 1 and only as a local mirroring session for traffic in *both* directions on the specified interface. (Selecting inbound-only or outbound-only is not an option.)
- If session 1 has been configured in the CLI for local mirroring for inbound-only or outbound-only traffic on one or more interfaces, using the Menu to change the session 1 configuration *automatically reconfigures the session* to monitor both inbound and outbound traffic on the designated interface(s.)
- If session 1 has been configured in the CLI with an ACL/classifier-based mirroring policy or as a remote mirroring session, the Menu is not available for changing the session 1 configuration.
- The CLI (and SNMP) can be used to override any Menu configuration of session 1.

Remote mirroring overview

To configure a remote mirroring session in which the mirroring source and destination are on different switches, follow these general steps:

1. Determine the IP addressing, UDP port number, and destination (exit) port number for the remote session:
 - a. Source VLAN or subnet IP address on the source switch.
 - b. Destination VLAN or subnet IP address on the destination switch.
 - c. Random UDP port number for the session (7933-65535.)
 - d. Remote mirroring endpoint: Exit port and IP address of the remote destination switch (In a remote mirroring endpoint, the IP address of the exit port and remote switch can belong to different VLANs. Any loopback IP address can be used except the default loopback address 127.0.0.1.)

Requirement: For remote mirroring, the same IP addressing and UDP port number must be configured on both the source and destination switches.
2. On the remote *destination* (endpoint) switch, enter the `mirror endpoint` command with the information from [step 1 \(page 361\)](#) to configure a mirroring session for a specific exit port.
3. Determine the session (1 to 4) and (optional) alphanumeric name to use on the *source* switch.
4. Determine the traffic to be filtered by any of the following selection methods and the appropriate configuration level (VLAN, port, mesh, trunk, global):
 - a. Direction: inbound, outbound, or both.
 - b. Classifier-based mirroring policy: inbound only for IPv4 or IPv6 traffic.
 - c. MAC source and/or destination address: inbound, outbound, or both.
5. On the *source* switch:
 - a. Enter the `mirror` command with the session number (1 to 4) and the IP addresses and UDP port number from [step 1 \(page 361\)](#) to configure a mirroring session. If desired, enter the `[truncation]` parameter to allow oversized packets to be truncated rather than dropped.
 - b. Enter one of the following commands to configure one or more of the traffic-selection methods in [step 4 \(page 361\)](#) for the configured session:


```
interface port/trunk/mesh [ monitor | service-policy policy-name
in ]
vlan vid [ monitor | service-policy policy-name in ]
monitor mac mac-addr
```

After you complete **b**, the switch begins mirroring traffic to the remote destination (endpoint) configured for the session.

Quick reference to remote mirroring setup

The commands beginning with “[Configuring the mirroring destination on a remote switch \(CLI\)](#)” ([page 334](#)), configure mirroring for a remote session in which the mirroring source and destination are on different switches:

- The `mirror` command identifies the destination in a mirroring session.
- The `interface` and `vlan` commands identify the monitored interface, traffic direction, and traffic-selection criteria for a specified session.

⚠ CAUTION: When configuring a remote mirroring session, always configure the destination switch first. Configuring the source switch first can result in a large volume of mirrored, IPv4-encapsulated traffic arriving at the destination without an exit path, which can slow switch performance.

High-level overview of the mirror configuration process

Determine the mirroring session and destination

For a local mirroring session

Determine the port number for the exit port (such as A5, B10, and so forth), then go to “[Configure the monitored traffic in a mirror session](#)” ([page 363](#)).

For a remote mirroring session

Determine the following information and then go to [“Configure a mirroring destination on a remote switch”](#) (page 362).

- The IP address of the VLAN or subnet on which the exit port exists on the destination switch.
- The port number of the remote exit port on the remote destination switch. (In a remote mirroring endpoint, the IP address of the exit port and the remote destination switch can belong to different VLANs.)
- The IP address of the VLAN or subnet on which the mirrored traffic enters or leaves the source switch.

⚠ CAUTION: Although the switch supports the use of UDP port numbers from 1 to 65535, UDP port numbers below 7933 are reserved for various IP applications. Using these port numbers for mirroring can result in an interruption of other IP functions, and in non-mirrored traffic being received on the destination (endpoint) switch and sent to the device connected to the remote exit port.

- The unique UDP port number to use for the session on the source switch. (The recommended port range is from 7933 to 65535.)

Configure a mirroring destination on a remote switch

This step is required only if you are configuring a remote mirroring session in which the exit port is on a different switch than the monitored (source) interface. If you are configuring local mirroring, go to [“Configure a mirroring session on the source switch”](#) (page 362).

For remote mirroring, you must configure the *destination* switch to recognize each mirroring session and forward mirrored traffic to an exit port before you configure the *source* switch. Configure the destination switch with the values you determined for remote mirroring in [“High-level overview of the mirror configuration process”](#) (page 361).

NOTE: A remote destination switch can support up to 32 remote mirroring endpoints (exit ports connected to a destination device in a remote mirroring session.)

Configure a destination switch in a remote mirroring session

Enter the `mirror endpoint ip` command on the remote switch to configure the switch as a remote endpoint for a mirroring session with a different source switch. (For information on how to do this, see [“Configuring a destination switch in a remote mirroring session \(CLI\)”](#) (page 336).)

Configure a mirroring session on the source switch

To configure local mirroring, only a session number and exit port number are required. For more information, see [“Configuring a source switch in a local mirroring session \(CLI\)”](#) (page 337).

If the exit port for a mirroring destination is on a remote switch instead of the local (source) switch, you must enter the source IP address, destination IP address, and UDP port number for the remote mirroring session (see [mirror remote ip](#) (page 363).) You may also wish to enable frame truncation to allow oversized frames to be truncated rather than dropped.

Frames that exceed the maximum size (MTU) are either dropped or truncated, according to the setting of the `[truncation]` parameter in the `mirror` command. Frames that are near the MTU size may become oversized when the 54-byte remote mirroring tunnel header is added for transport between source switch and destination switch. (The addition of the header is a frequent cause for frames becoming oversized, but note that all oversized frames, whatever the cause of their excess size, are dropped or truncated.) If a frame is truncated, bytes are removed from the end of the frame. This may cause the checksum in the original frame header to fail. Some protocol analyzers may flag such a checksum mismatch as an alert.

NOTE: Note that if you enable jumbo frames to allow large frames to be transmitted, you must enable jumbo frames on all switches in the path between source and destination switches.

Configure a source switch in a remote mirroring session

Enter the `mirror remote ip` command on the source switch to configure a remote destination switch for a mirroring session on the source switch. (For information on how to do this, see [“Configuring a source switch in a remote mirroring session \(CLI\)”](#) (page 337). The source IP address, UDP port number, and destination IP address that you enter must be the same values that you entered with the `mirror endpoint ip` command in [“Configure a mirroring destination on a remote switch”](#) (page 362).

- △ **CAUTION:** After you configure a mirroring session with traffic-selection criteria and a destination, the switch immediately starts to mirror traffic to the destination device connected to each exit port. In a remote mirroring session that uses IPv4 encapsulation, if the remote (endpoint) switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic. For this reason, HP Switch strongly recommends that you configure the endpoint switch in a remote mirroring session, as described in [“Configure a mirroring destination on a remote switch”](#) (page 362), before using the `mirror remote ip` command in this section to configure the mirroring source for the same session.

Configure the monitored traffic in a mirror session

This step configures one or more interfaces on a source switch with traffic-selection criteria to select the traffic to be mirrored in a local or remote session configured in section [“Configure a mirroring session on the source switch”](#) (page 362).

Traffic selection options

To configure traffic mirroring, specify the source interface, traffic direction, and criteria to be used to select the traffic to be mirrored by using the following options:

- Interface type
 - Port, trunk, and/or mesh
 - VLAN
 - Switch (global configuration level)
- Traffic direction and selection criteria
 - All inbound and/or outbound traffic on a port or VLAN interface
 - Only inbound IP traffic selected with an ACL (deprecated in software release **K.14.01 and greater**)
 - Only inbound IPv4 or IPv6 traffic selected with a classifier-based mirroring policy
 - All inbound and/or outbound traffic selected by MAC source and/or destination address

The different ways to configure traffic-selection criteria on a monitored interface are described in the following sections.

Mirroring-source restrictions

In a mirroring session, you can configure any of the following sources of mirrored traffic:

- Multiple port and trunk, and/or mesh interfaces
- One VLAN

If you configure a VLAN as the source interface in a mirroring session and assign a second VLAN to the session, the second VLAN overwrites the first VLAN as the source of mirrored traffic.

- One classifier-based policy
If you configure a mirroring policy on a port or VLAN interface to mirror inbound traffic in a session, you cannot configure a port, trunk, mesh, ACL, or VLAN as an additional source of mirrored traffic in the session.
- Up to 320 MAC addresses (used to select traffic according to source, destination MAC address, or both) in all mirroring sessions configured on a switch

About selecting all inbound/outbound traffic to mirror

Use the commands beginning with “[Selecting all traffic on a port interface for mirroring according to traffic direction \(CLI\)](#)” (page 339) and ending with “[Selecting all traffic on a VLAN interface for mirroring according to traffic direction \(CLI\)](#)” (page 340) to configure all inbound and/or outbound traffic on a specified VLAN, port, or trunk interfaces for a local or remote mirroring session. For an example of a mirroring configuration that selects all inbound or outbound traffic on a monitored interface, see:

- “[Example: Local mirroring using traffic-direction criteria](#)” (page 371)
- “[Example: Remote mirroring using a classifier-based policy](#)” (page 371)

NOTE: If you have already configured session 1 with a local or remote destination (as described in “[Configure a mirroring session on the source switch](#)” (page 362)), you can enter the `vlan vid monitor` or `interface port monitor` command without additional parameters for traffic-selection criteria and session number to configure mirroring for all inbound and outbound traffic on the specified VLAN or port interfaces in session 1 with the preconfigured destination.

Untagged mirrored packets

Although a VLAN tag is added (by default) to the mirrored copy of untagged outbound packets to indicate the source VLAN of the packet, it is sometimes desirable to have mirrored packets look exactly like the original packet. The `no-tag-added` parameter gives you the option of not tagging mirrored copies of outbound packets, as shown in [Figure 205](#) (page 364) and [Figure 206](#) (page 364).

Figure 205 Mirroring commands with the `no-tag-added` option

```
HP Switch(config)#interface 3 monitor all in mirror 1 no-tag-added
HP Switch(config)#interface mesh monitor all both mirror 1 no-tag-added
```

Figure 206 Displaying a mirror session configuration with the `no-tag-added` option

```
HP Switch# show monitor 1

Network Monitoring

Session: 1   Session Name:
ACL: no ACL relationship exists

Mirror Destination: 48
Untagged traffic  : untagged ← Indicates the no-tag-added option is configured.
Monitoring Sources Direction
-----
Port: 3           Both
```

About using SNMP to configure `no-tag-added`

The MIB object `hpicfBridgeDontTagWithVlan` is used to implement the `no-tag-added` option, as shown below:

```
hpicfBridgeDontTagWithVlan OBJECT-TYPE
    SYNTAX INTEGER
```

```

    {
      enabled(1),
      disabled(2)
    }
MAX-ACCESS read-write
STATUS current
DESCRIPTION
  "This oid mentions whether VLAN tag is part of the
  mirror'ed copy of the packet. The value 'enabled'
  denotes that the VLAN tag shouldn't be part
  of the mirror'ed copy; 'disabled' does put
  the VLAN tag in the mirror'ed copy. Only one
  logical port is allowed.
  This object is persistent and when written
  the entity
  SHOULD save the change to non-volatile storage."
DEFVAL { 2 }
 ::= { hpicfBridgeMirrorSessionEntry 2 }

```

Operating notes

The following conditions apply for the `no-tag-added` option:

- The specified port can be a physical port, trunk port, or mesh port.
- Only a single logical port (physical port or trunk) can be associated with a mirror session when the `no-tag-added` option is specified. No other combination of ACL mirroring, VLAN mirroring, or port mirroring can be associated with the mirror session. If more than one logical port is specified, the following error message is displayed:
Cannot monitor more than one logical port with no-tag-added option
- If a port changes its VLAN membership and/or untagged status within the VLAN, the "untagged port mirroring" associated with that port is updated when the configuration change is processed.
- Only four ports or trunks can be monitored at one time when all four mirror sessions are in use (one logical port per mirror session) without VLAN tags being added to a mirrored copy.
- The `no-tag-added` option can also be used when mirroring is configured with SNMP.
- A VLAN tag is still added to the copies of untagged packets obtained via VLAN-based mirroring.

About selecting inbound traffic using an ACL (deprecated)

Deprecation of ACL-based traffic selection

In release K.14.01 or greater, the use of ACLs to select inbound traffic in a mirroring session has been replaced with classifier-based mirroring policies (see ["About selecting inbound traffic using advanced classifier-based mirroring" \(page 367\)](#).)

The following commands have been deprecated:

- `interface port/trunk/mesh monitor ip access-group acl-name in mirror 1 - 4 | name-str`
- `vlan vid-# monitor ip access-group acl-name in mirror 1 - 4 | name-str`

After you install and boot release K.14.01 or greater, ACL-based local and remote mirroring sessions configured on a port or VLAN interface are automatically converted to classifier-based mirroring policies. For more information, see ["Migration to release K.14.01 or greater" \(page 359\)](#).

About selecting inbound/outbound traffic using a MAC address

Use the `monitor mac mirror` command at the global configuration level to apply a source and/or destination MAC address as the selection criteria used in a local or remote mirroring session.

While classifier-based mirroring allows you to mirror traffic using a policy to specify IP addresses as selection criteria, MAC-based mirroring allows you monitor switch traffic using a source and/or destination MAC address. You can apply MAC-based mirroring in one or more mirroring sessions on the switch to monitor:

- Inbound traffic
- Outbound traffic
- Both inbound and outbound traffic

MAC-based mirroring is useful in HP Switch Network Immunity security solutions that provide detection and response to malicious traffic at the network edge. After isolating a malicious MAC address, a security administrator can mirror all traffic sent to and received from the suspicious address for troubleshooting and traffic analysis.

The MAC address that you enter with the `monitor mac mirror` command is configured to select traffic for mirroring from all ports and learned VLANs on the switch. Therefore, a suspicious MAC address used in wireless applications can be continuously monitored as it re-appears in switch traffic on different ports or VLAN interfaces.

You can configure MAC-based mirroring from the CLI or an SNMP management station and use it to mirror:

- All inbound and outbound traffic from a group of hosts to one destination device.
- Inbound and/or outbound traffic from each host to a different destination device.
- Inbound and outbound traffic from all monitored hosts separately on two destination devices: mirroring all inbound traffic to one device and all outbound traffic to another device.

Restrictions

The following restrictions apply to MAC-based mirroring:

- Up to 320 different MAC addresses are supported for traffic selection in all mirroring sessions configured on the switch.
- A destination MAC address is not supported as mirroring criteria for routed traffic, because in routed packets, the destination MAC address is changed to the next-hop address when the packet is forwarded. Therefore, the destination MAC address that you want to mirror will not appear in routed packet headers.

This restriction also applies to the destination MAC address of a host that is directly connected to a routing switch. (Normally, a host is connected to an edge switch, which is directly connected to the router.)

To mirror routed traffic, we recommend that you use classifier-based policies to select IPv4 or IPv6 traffic for mirroring, as described in [“About selecting inbound traffic using advanced classifier-based mirroring”](#) (page 367).

- On a switch, you can use a MAC address only once as a source MAC address and only once as a destination MAC address to filter mirrored traffic.

For example, after you enter the following commands:

```
monitor mac 111111-222222 src mirror 1
monitor mac 111111-222222 dest mirror 2
```

The following commands are not supported:

```
monitor mac 111111-222222 src mirror 3
monitor mac 111111-222222 dest mirror 4
```

In addition, if you enter the `monitor mac 111111-222222 both mirror 1` command, you cannot use the MAC address `111111-222222` in any other `monitor mac mirror` configuration commands on the switch.

- To re-use a MAC address that has already been configured as a source and/or destination address for traffic selection in a mirror session, you must first remove the configuration by entering the `no` form of the command and then re-enter the MAC address in a new `monitor mac mirror` command.

For example, if you have already configured MAC address `111111-222222` to filter inbound and outbound mirrored traffic, and you decide to use it to filter only inbound traffic in a mirror session, you could enter the following commands:

```
monitor mac 111111-222222 both mirror 1
no monitor mac 111111-222222 both mirror 1
monitor mac 111111-222222 src mirror 1
```

- A mirroring session in which you configure MAC-based mirroring is not supported on a port, trunk, mesh, or VLAN interface on which a mirroring session with a classifier-based mirroring policy is configured.

About selecting inbound traffic using advanced classifier-based mirroring

In software release K.14.01 or greater, in addition to the traffic selection options described in [“Configure the monitored traffic in a mirror session” \(page 363\)](#), traffic mirroring supports the use of advanced classifier-based functions that provide:

- A finer granularity for selecting the inbound IP traffic that you want to mirror on an individual port or VLAN interface (instead of mirroring all inbound traffic on the interface)
- Support for mirroring both IPv4 and IPv6 traffic
- The ability to re-use the same traffic classes in different software-feature configurations; for example, you can apply both a QoS rate-limiting and mirroring policy on the same class of traffic.

Deprecation of ACL-based traffic selection

In software release K.14.01 or greater, advanced classifier-based policies replace ACL-based traffic selection in mirroring configurations.

Like ACL-based traffic-selection criteria, classifier-based service policies apply only to inbound traffic flows and are configured on a per-port or per-VLAN basis. In a mirroring session, classifier-based service policies do not support:

- The mirroring of outbound traffic exiting the switch
- The use of meshed ports as monitored (source) interfaces

Classifier-based mirroring is *not* designed to work with other traffic-selection methods in a mirroring session applied to a port or VLAN interface:

- If a mirroring session is already configured with one or more traffic-selection criteria (MAC-based or all inbound and/or outbound traffic), the session does not support the addition of a classifier-based policy.
- If a mirroring session is configured to use a classifier-based mirroring policy, no other traffic-selection criteria (MAC-based or all inbound and/or outbound traffic) can be added to the session on the same or a different interface.

Classifier-based mirroring policies provide greater precision when analyzing and debugging a network traffic problem. Using multiple match criteria, you can finely select and define the classes of traffic that you want to mirror on a traffic analyzer or IDS device.

For more information on how to configure and use classifier-based service policies, see the *Advanced Traffic Management Guide*.

For an example of a mirroring configuration that uses a classifier-based service policy to select traffic on a monitored interface, see [“Example: Remote mirroring using a classifier-based policy” \(page 371\)](#).

Classifier-based mirroring configuration

1. Evaluate the types of traffic in your network and identify the traffic types that you want to mirror.
2. Create an IPv4 or IPv6 traffic class using the `class` command to select the packets that you want to mirror in a session on a preconfigured local or remote destination device. (See [“Configuring classifier-based mirroring \(CLI\)” \(page 342\)](#).)

A traffic class consists of match criteria, which consist of match and ignore commands.

- `match` commands define the values that header fields must contain for a packet to belong to the class and be managed by policy actions.
- `ignore` commands define the values which, if contained in header fields, exclude a packet from the policy actions configured for the class.

NOTE: Be sure to enter match/ignore statements in the *precise order* in which you want their criteria to be used to check packets.

The following match criteria are supported in match/ignore statements for inbound IPv4/IPv6 traffic:

- IP source address (IPv4 and IPv6)
- IP destination address (IPv4 and IPv6)
- IP protocol (such as ICMP or SNMP)
- Layer 3 IP precedence bits
- Layer 3 DSCP codepoint
- Layer 4 TCP/UDP application port (including TCP flags)
- VLAN ID

Enter one or more match or ignore commands from the class configuration context to filter traffic and determine the packets on which policy actions will be performed. (See [\(page 342\)](#).)

3. Create a mirroring policy to configure the session and destination device to which specified classes of inbound traffic are sent by entering the `policy mirror` command from the global configuration context. (See [\(page 330\)](#).)

NOTE: Be sure to enter each class and its associated mirroring actions in the *precise order* in which you want packets to be checked and processed.

To configure the mirroring actions that you want to execute on packets that match the criteria in a specified class, enter one or more class action mirror commands from the policy configuration context. (See [\(page 342\)](#).)

You can configure only one mirroring session (destination) for each class. However, you can configure the same mirroring session for different classes.

A packet that matches the match criteria in a class is mirrored to the exit (local or remote) port that has been previously configured for the session, where session is a value from **1** to **4** or a text string (if you configured the session with a name when you entered the `mirror` command.)

Prerequisite: The local or remote exit port for a session must be already configured before you enter the `mirror session` parameter in a class action statement:

- In a local mirroring session, the exit port is configured with the `mirror session-number port` command.
- In a remote mirroring session, the remote exit port is configured with the `mirror endpoint ip` and `mirror session-number remote ip` commands.

For more information, see [“Configure a mirroring destination on a remote switch” \(page 362\)](#) and [“Configure a mirroring session on the source switch” \(page 362\)](#).

Restriction: In a policy, you can configure only one mirroring session per class. However, you can configure the same session for different classes.

Mirroring is not executed on packets that match ignore criteria in a class.

The execution of mirroring actions is performed in the order in which the classes are numerically listed in the policy.

The complete no form of the `class action mirror` command or the `no seq-number` command removes a class and mirroring action from the policy configuration.

To manage packets that do not match the match or ignore criteria in any class in the policy, and therefore have no mirroring actions performed on them, you can enter an optional default class. The default class is placed at the end of a policy configuration and specifies the mirroring actions to perform on packets that are neither matched nor ignored.

4. (Optional) To configure a default-class in a policy, enter the `default-class` command at the end of a policy configuration and specify one or more actions to be executed on packets that are not matched and not ignored. (See [“Syntax” \(page 343\)](#).)

Prerequisite: The local or remote exit port for a session must be already configured with a destination device before you enter the `mirror session` parameter in a default-class action statement. For more information, see [“Configure a mirroring destination on a remote switch” \(page 362\)](#) and [“Configure a mirroring session on the source switch” \(page 362\)](#).

For general information about how to configure and manage a service policy, see the *Advanced Traffic Management Guide*.

5. Apply the mirroring policy to inbound traffic on a port (`interface service-policy in` command) or VLAN (`vlan service-policy in` command) interface.

CAUTION: After you apply a mirroring policy for one or more preconfigured sessions on a port or VLAN interface, the switch immediately starts to use the traffic-selection criteria and exit port to mirror traffic to the destination device connected to each exit port.

In a remote mirroring session that uses IPv4 encapsulation, if the remote switch is not already configured as the destination for the session, its performance may be adversely affected by the stream of mirrored traffic.

For this reason, HP Switch strongly recommends that you first configure the exit switch in a remote mirroring session, as described in [“Configure a mirroring destination on a remote switch” \(page 362\)](#) and [“Configure a mirroring session on the source switch” \(page 362\)](#), before you apply a mirroring service policy on a port or VLAN interface.

Restrictions: The following restrictions apply to a mirroring service policy:

- Only one mirroring policy is supported on a port or VLAN interface.
- If you apply a mirroring policy to a port or VLAN interface on which a mirroring policy is already configured, the new policy replaces the existing one.
- A mirroring policy is supported only on inbound traffic.

Because only one mirroring policy is supported on a port or VLAN interface, ensure that the policy you want to apply contains all the required classes and actions for your configuration.

Classifier-based mirroring restrictions

The following restrictions apply to mirroring policies configured with the classifier-based model:

- A mirroring policy is supported only on *inbound* IPv4 or IPv6 traffic.
- A mirroring policy is not supported on a meshed port interface. (Classifier-based policies are supported only on a port, VLAN, or trunk interface.)
- Only one classifier-based mirroring policy is supported on a port or VLAN interface. You can, however, apply a classifier-based policy of a different type, such as QoS.

- You can enter multiple `class action mirror` statements in a policy.
 - You can configure only one mirroring session (destination) for each class.
 - You can configure the same mirroring session for different classes.
- If a mirroring session is configured with a classifier-based mirroring policy on a port or VLAN interface, no other traffic-selection criteria (MAC-based or all inbound and/or outbound traffic) can be added to the session.

Figure 207 Mirroring configuration in which only a mirroring policy is supported

```
Switch-B(config)# mirror endpoint 10.10.40.4 9200 10.10.50.5 port al
...
Switch-A(config)# mirror 1 remote ip 10.10.40.4 9200 10.10.50.5
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
Switch-A(config)# class ipv4 Data2
Switch-A(config-class)# match ip 10.28.31.1 any
Switch-A(config-class)# match ip any host 10.28.31.0/24
Switch-A(config-class)# exit
Switch-A(config)# policy mirror SalesData
Switch-A(config-policy)# class ipv4 Data2 action mirror 1
Switch-A(config-policy)# exit
Switch-A(config)# vlan 10 service-policy SalesData in
Switch-A(config)# vlan 10 monitor all out mirror 1
A prior mirror policy relationship exists with mirror session 1. Please remove.
```

Classifier-based policy used to select mirrored traffic in session 1

The configuration of additional traffic-direction criteria to select mirrored traffic is not supported in session 1.

- If a mirroring session is already configured with one or more traffic-selection criteria (MAC-based or all inbound and/or outbound traffic), the session does not support the addition of a classifier-based policy.

Figure 208 Mirroring configuration in which only traffic-selection criteria are supported

```
Switch-B(config)# mirror endpoint 10.10.40.4 9200 10.10.50.5 port al
...
Switch-A(config)# mirror 1 remote ip 10.10.40.4 9200 10.10.50.5
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
Switch-A(config)# vlan 10 monitor all out mirror 1
Switch-A(config)# class ipv4 Data2
Switch-A(config-class)# match ip 10.28.31.1 any
Switch-A(config-class)# match ip any host 10.28.31.0/24
Switch-A(config-class)# exit
Switch-A(config)# policy mirror SalesData
Switch-A(config-policy)# class ipv4 Data2 action mirror 1
Switch-A(config-policy)# exit
Switch-A(config)# vlan 10 service-policy SalesData in
Mirror source VLAN exists on mirror session 1. Cannot add this mirror source.
```

Configuration of traffic-direction criteria to select all outbound traffic on VLAN 10 in mirror session 1

The configuration of an additional classifier-based policy to select mirrored traffic on VLAN 10 is not supported in session 1.

About applying multiple mirroring sessions to an interface

You can apply a mirroring policy to an interface that is already configured with another traffic-selection method (MAC-based or all inbound and/or outbound traffic) for a different mirroring session.

The classifier-based policy provides a finer level of granularity that allows you to zoom in on a subset of port or VLAN traffic and select it for local or remote mirroring.

In the following example, traffic on Port b1 is used as the mirroring source for two different, local mirroring sessions:

- All inbound and outbound traffic on Ports b1, b2, and b3 is mirrored in session 4.
- Only selected voice traffic on Port b1 is mirrored in session 2.

Figure 209 Example of applying multiple sessions to the same interface

```
HP Switch(config)# mirror 4 port a2
HP Switch(config)# interface b1-b3 monitor all both mirror 4
HP Switch(config)# mirror 2 port b4
HP Switch(config)# class ipv4 voice
HP Switch(config-class)# match ip any any ip-dscp ef
HP Switch(config-class)# exit
HP Switch(config)# policy mirror IPphones
HP Switch(config-policy)# class ipv4 voice action mirror 2
HP Switch(config-policy)# exit
HP Switch(config)# interface b1 service-policy IPphones in
```

Mirroring configuration examples

Example: Local mirroring using traffic-direction criteria

An administrator wants to mirror the inbound traffic from workstation "X" on port A5 and workstation "Y" on port B17 to a traffic analyzer connected to port C24 (see [Figure 210 \(page 371\)](#).) In this case, the administrator chooses "1" as the session number. (Any unused session number from 1 to 4 is valid.) Because the switch provides both the source and destination for the traffic to monitor, local mirroring can be used. In this case, the command sequence is:

1. Configure the local mirroring session, including the exit port.
2. Configure the monitored source interfaces for the session.

Figure 210 Local mirroring topology




Figure 211 Configuring a local mirroring session for all inbound and outbound port traffic

```
HP Switch(config)# mirror 1 port c24
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
HP Switch(config)# interface a5,b17 monitor all in mirror
1
```

Configures port C24 as the mirroring destination (exit port) for session 1.

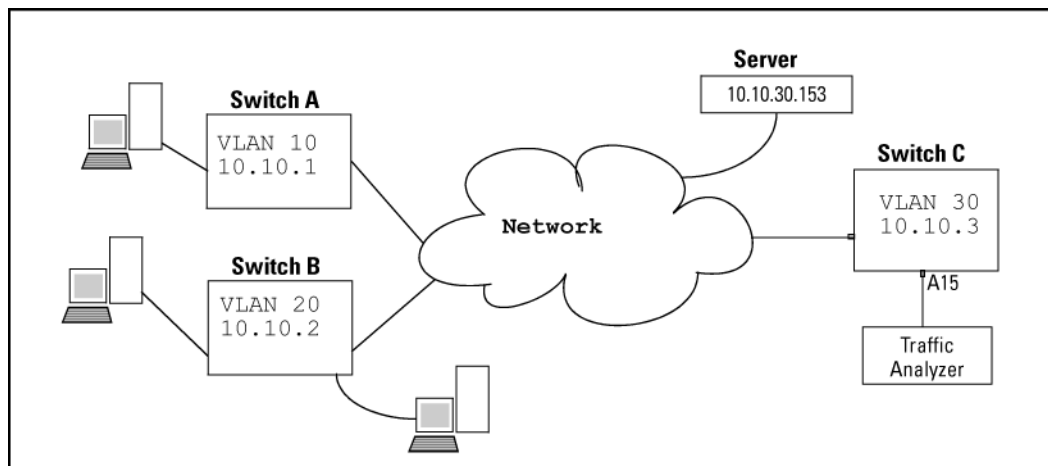
Reminder to configure mirroring destination before configuring source.

Mirrors all inbound and outbound traffic on ports A5 and B17 to the mirroring destination configured for session 1.

Example: Remote mirroring using a classifier-based policy

In the network shown in [Figure 212 \(page 372\)](#), an administrator has connected a traffic analyzer to port A15 (in VLAN 30) on switch C to monitor the TCP traffic to the server at 10.10.30.153 from workstations connected to switches A and B. Remote mirroring sessions are configured on switches A and B, and a remote mirroring endpoint on switch C. TCP traffic is routed through the network to the server from VLANs 10 and 20 on VLAN 30.

Figure 212 Sample topology in a remote mirroring session



To configure this remote mirroring session using a classifier-based policy to select inbound TCP traffic on two VLAN interfaces, take the following steps:

1. On remote switch C, configure a remote mirroring endpoint using port A15 as the exit port (as described in “Configure a mirroring destination on a remote switch” (page 362).)

Figure 213 Configuring a remote mirroring endpoint: remote switch and exit port

On the remote (endpoint) switch, configures the IP address and UDP port of each source switch.	On the remote (endpoint) switch, configures the destination IP address and remote exit port.
Switch-C(config)# mirror endpoint 10.10.10.119 9300 10.10.30.2 port a15	
Switch-C(config)# mirror endpoint 10.10.20.145 9300 10.10.30.2 port a15	

2. On source switch A, configure an association between the remote mirroring endpoint on switch C and a mirroring session on switch A (as described in “Configure a mirroring session on the source switch” (page 362).)
3. On switch A, configure a classifier-based mirroring policy to select inbound TCP traffic destined to the server at 10.10.30.153, and apply the policy to the interfaces of VLAN 10 (as described in “About selecting inbound traffic using advanced classifier-based mirroring” (page 367).)

Figure 214 Configuring a classifier-based policy on source switch A

```
On a source switch, associates session number 1 with a source IP address and UDP port, and a remote destination IP address.
1 Switch-A(config)# mirror 1 remote ip 10.10.10.119 9300 10.10.30.2
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
2 Switch-A(config)# class ipv4 tcp7
Switch-A(class-config)# match tcp any 10.10.30.153
Switch-A(class-config)# match tcp any host 10.10.20.153/24
Switch-A(class-config)# match tcp any any eq 80
Switch-A(class-config)# exit
Switch-A(config)# policy mirror mirrorTCP
Switch-A(policy-config)# class ipv4 tcp7 action mirror 1
Switch-A(policy-config)# exit
3 Switch-A(config)# vlan 10 service-policy mirrorTCP in
```

1 The source IP address and UDP port number identify the mirroring source in session 1; the destination IP address identifies the remote switch to which traffic is mirrored. (The exit port for mirrored traffic, configured in Figure B-B-54, and the remote switch can belong to different VLANs.)

2 Configures a class that selects IPv4 TCP traffic destined to: the server at 10.10.30.153, a device in subnet 10.10.20.0, and any TCP traffic on port 80. (A packet that does not match these criteria is transmitted without being mirrored.)

3 Configures VLAN 10 as the source interface, and the mirroring policy as the selection criteria for inbound traffic on VLAN 10 in session 1.

4. On source switch B, repeat steps 2 and 3:
 - a. Configure an association between the remote mirroring endpoint on switch C and a mirroring session on switch B.
 - b. Configure a classifier-based mirroring policy to select inbound TCP traffic destined to the server at 10.10.30.153, and apply the policy to a VLAN interface for VLAN 20.

Because the remote session has mirroring sources on different switches, you can use the same session number (1) for both sessions.

Figure 215 Configuring a classifier-based policy on source switch B

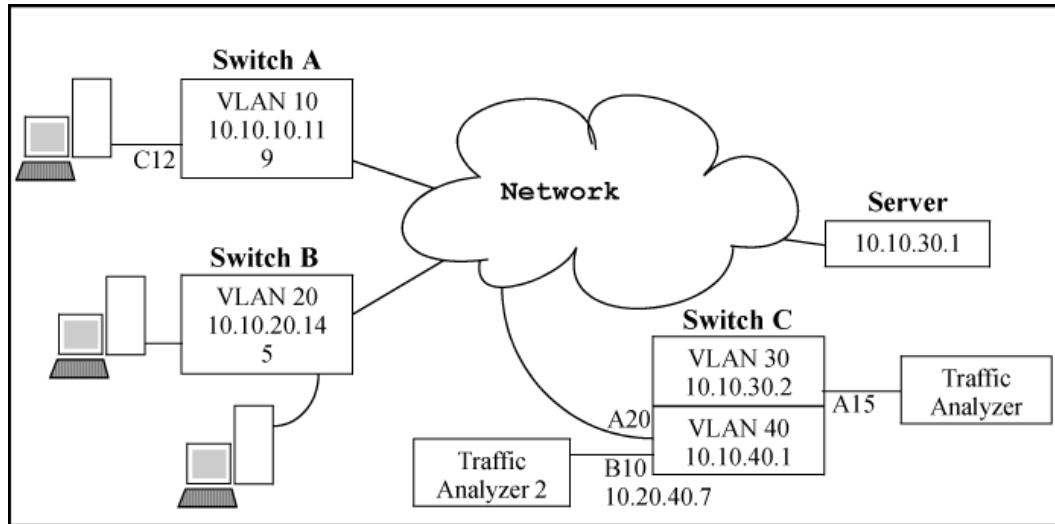
```
The configuration of remote-mirroring session 1 on Switch B is the same as on Switch A (figure B-55), except for the difference in source VLAN and source IP address. Note that on different switches, the UDP port number (9300) can be the same.
Switch-B(config)# mirror 1 remote ip 10.10.20.145 9300 10.10.30.2
Caution: Please configure destination switch first.
Do you want to continue [y/n]? y
Switch-B(config)# class ipv4 tcp7
Switch-B(class-config)# match tcp any 10.10.30.153
Switch-B(class-config)# match tcp any host 10.10.20.153/24
Switch-B(class-config)# match tcp any any eq 80
Switch-B(class-config)# exit
Switch-B(config)# policy mirror mirrorTCP
Switch-B(policy-config)# class ipv4 tcp7 mirror 1
Switch-B(policy-config)# exit
Switch-B(config)# vlan 20 service-policy mirrorTCP in
```

Example: Remote mirroring using traffic-direction criteria

In the network shown in [Figure 216 \(page 374\)](#), the administrator connects another traffic analyzer to port B10 (in VLAN 40) on switch C to monitor all traffic entering switch A on port C12. For this mirroring configuration, the administrator configures a mirroring destination (with a remote exit port of B10) on switch C, and a remote mirroring session on switch A.

If the mirroring configuration in the preceding example is enabled, it is necessary to use a different session number (2) and UDP port number (9400.) (The IP address of the remote exit port [10.10.40.7] connected to traffic analyzer 2 [exit port B10] can belong to a different VLAN than the destination IP address of the VLAN used to reach remote switch C [10.20.40.1]).


Figure 216 Sample topology for remote mirroring from a port interface



To configure this remote mirroring session using a directional-based traffic selection on a port interface, the operator must take the following steps:

1. On remote switch C, configure the remote mirroring endpoint using port B10 as the exit port for a traffic analyzer (as described in [“Configure a mirroring destination on a remote switch”](#) (page 362)):

Figure 217 Configuring a remote mirroring endpoint



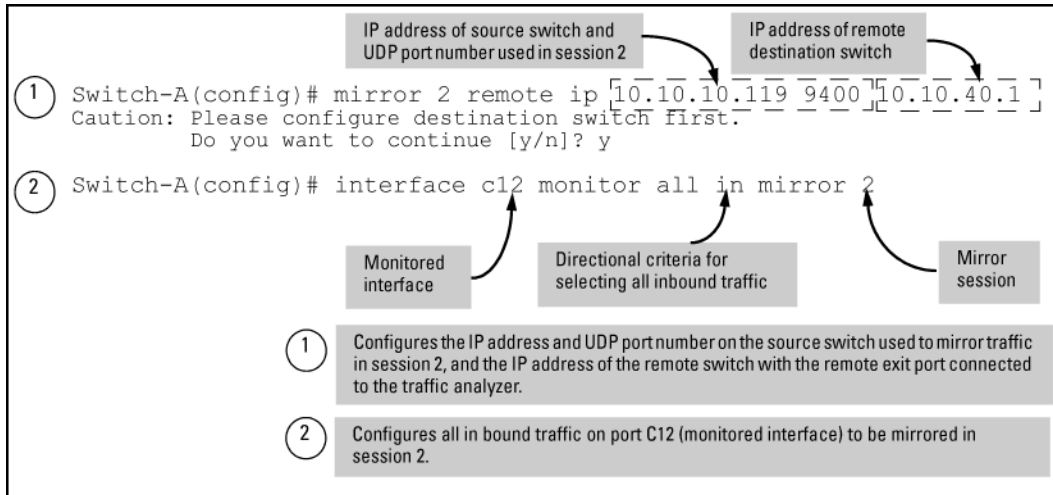
2. On source switch A, configure session 2 to use UDP port 9400 to reach the remote mirroring endpoint on switch C (10.10.40.1):

```
mirror 2 remote ip 10.10.10.119 9400 10.10.40.1
```

- On source switch A, configure the local port C12 to select all inbound traffic to send to the preconfigured mirroring destination for session 2:

```
interface c12 monitor all in mirror 2
```

Figure 218 Configuring a remote mirroring session for inbound port traffic



Maximum supported frame size

The IPv4 encapsulation of mirrored traffic adds a 54-byte header to each mirrored frame. If a resulting frame exceeds the MTU allowed in the network, the frame is dropped or truncated.

NOTE: Oversized mirroring frames are dropped or truncated, according to the setting of the [truncation] parameter in the mirror command. Also, remote mirroring does not allow downstream devices in a mirroring path to fragment mirrored frames.

If jumbo frames are enabled on the mirroring source switch, the mirroring destination switch and all downstream devices connecting the source switch to the mirroring destination must be configured to support jumbo frames.

Enabling jumbo frames to increase the mirroring path MTU

On 1-Gbps and 10-Gbps ports in the mirroring path, you can reduce the number of dropped frames by enabling jumbo frames on all intermediate switches and routers. (The MTU on the switches covered by this manual is 9220 bytes for frames having an 802.1Q VLAN tag, and 9216 bytes for untagged frames.)

Table 24 Maximum frame sizes for mirroring

	Frame type configuration	Maximum frame size	VLAN tag	Frame mirrored to local port	Frame mirrored to remote port	
				Data	Data	IPv4 header
Untagged	Non-jumbo (default config.)	1518	0	1518	1464	54
	Jumbo ¹ on all VLANs	9216	0	9216	9162	54
	Jumbo ¹ On all but source VLAN	1518	0	n/a ²	1464	54
Tagged	Non-jumbo	1522	4	1522	1468	54

Table 24 Maximum frame sizes for mirroring *(continued)*

	Frame type configuration	Maximum frame size	VLAN tag	Frame mirrored to local port	Frame mirrored to remote port	
				Data	Data	IPv4 header
	Jumbo ¹ on all VLANs	9220	4	9218	9164	54
	Jumbo ¹ On all but source VLAN	1522	4	n/a ²	1468	54

¹ Jumbo frames are allowed on ports operating at or above 1 Gbps

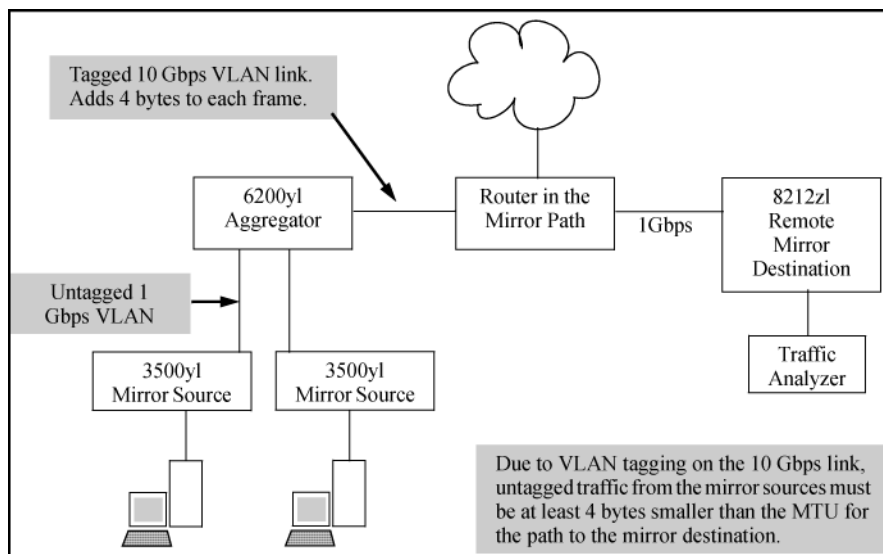
² For local mirroring, a non-jumbo configuration on the source VLAN dictates an MTU of 1518 bytes for untagged frames, and an MTU of 1522 for tagged frames, regardless of the jumbo configuration on any other VLANs on the switch.

Effect of downstream VLAN tagging on untagged, mirrored traffic

In a remote mirroring application, if mirrored traffic leaves the switch without 802.1Q VLAN tagging, but is forwarded through a downstream device that adds 802.1Q VLAN tags, the MTU for untagged mirrored frames leaving the source switch is reduced below the values shown in Table 24 (page 375).

For example, if the MTU on the path to the destination is 1522 bytes, untagged mirrored frames leaving the source switch cannot exceed 1518 bytes. Likewise, if the MTU on the path to the destination is 9220 bytes, untagged mirrored frames leaving the source switch cannot exceed 9216 bytes.

Figure 219 Effect of downstream VLAN tagging on the MTU for mirrored traffic



Operating notes for traffic mirroring

- **Mirroring dropped traffic**

When an interface is configured to mirror traffic to a local or remote destination, packets are mirrored regardless of whether the traffic is dropped while on the interface. For example, if an ACL is configured on a VLAN with a `deny` ACE that eliminates packets from a Telnet application, the switch still mirrors the Telnet packets that are received on the interface and subsequently dropped.

- **Mirroring and spanning tree**

Mirroring is performed regardless of the STP state of a port or trunk. This means, for example, that inbound traffic on a port blocked by STP can still be monitored for STP packets during the STP setup phase.

- **Tagged and untagged frames**

For a frame entering or leaving the switch on a mirrored port, the mirrored copy retains the tagged or untagged state the original frame carried when it entered into or exited from the switch. (The tagged or untagged VLAN membership of ports in the path leading to the mirroring destination does not affect the tagged or untagged status of the mirrored copy itself.)

Thus, if a tagged frame arrives on a mirrored port, the mirrored copy is also tagged, regardless of the status of ports in the destination path. If a frame exits from the switch on a mirrored port that is a tagged member of a VLAN, the mirrored copy is also tagged for the same reason.

To prevent a VLAN tag from being added to the mirrored copy of an outbound packet sent to a mirroring destination, you must enter the `no-tag-added` parameter when you configure a port, trunk, or mesh interface to select mirrored traffic. For more information, see [“Selecting all traffic on a port interface for mirroring according to traffic direction \(CLI\)”](#) (page 339) and [“Untagged mirrored packets”](#) (page 364).

- **Effect of IGMP on mirroring**

If both inbound and outbound mirroring is operating when IGMP is enabled on a VLAN, two copies of mirrored IGMP frames may appear at the mirroring destination.

- **Mirrored traffic not encrypted**

Mirrored traffic undergoes IPv4 encapsulation, but mirrored encapsulated traffic is not encrypted.

- **IPv4 header added**

The IPv4 encapsulation of mirrored traffic adds a 54-byte header to each mirrored frame. If a resulting frame exceeds the maximum MTU allowed in the network, it is dropped or truncated (according to the setting of the `[truncation]` parameter in the `mirror` command.)

To reduce the number of dropped frames, enable jumbo frames in the mirroring path, including all intermediate switches and/or routers. (The MTU on the switch is 9220 bytes, which includes 4 bytes for the 802.1Q VLAN tag.) For more information, see [“Maximum supported frame size”](#) (page 375). To configure the switch for jumbo frames, see [“Viewing the current jumbo configuration”](#) (page 148).

- **Intercepted or injected traffic**

The mirroring feature does not protect against either mirrored traffic being intercepted or traffic being injected into a mirrored stream by an intermediate host.

- **Inbound mirrored IPv4-encapsulated frames are not mirrored**

The switch does not mirror IPv4-encapsulated mirrored frames that it receives on an interface. This prevents duplicate mirrored frames in configurations where the port connecting the switch to the network path for a mirroring destination is also a port whose inbound or outbound traffic is being mirrored.

For example, if traffic leaving the switch through ports B5, B6, and B7 is being mirrored through port B7 to a network analyzer, the mirrored frames from traffic on ports B5 and B6 will not be mirrored a second time as they pass through port B7.

- **Switch operation as both destination and source**

A switch configured as a remote destination switch can also be configured to mirror traffic to one of its own ports (local mirroring) or to a destination on another switch (remote mirroring.)

- **Monitor command note**

If session 1 is already configured with a destination, you can enter the `[no] vlan vid monitor` or `[no] interface port monitor` command without mirroring criteria and a mirror session number. In this case, the switch automatically configures or removes mirroring

for inbound and outbound traffic from the specified VLAN or ports to the destination configured for session 1.

- **Loss of connectivity suspends remote mirroring**

When a remote mirroring session is configured on a source switch, the switch sends an ARP request to the configured destination approximately every 60 seconds. If the source switch fails to receive the expected ARP response from the destination for the session, transmission of mirrored traffic in the session halts. However, because the source switch continues to send ARP requests for each configured remote session, link restoration or discovery of another path to the destination enables the source switch to resume transmitting the session's mirrored traffic after a successful ARP response cycle occurs.

Note that if a link's connectivity is repeatedly interrupted ("link toggling"), little or no mirrored traffic may be allowed for sessions using that link. To verify the status of any mirroring session configured on the source switch, use the `show monitor` command.

Troubleshooting traffic mirroring

If mirrored traffic does not reach the configured remote destination (endpoint) switch or remote exit port, check the following configurations:

- In a remote mirroring session, the `mirror remote ip` command parameters configured on the source switch for source IP address, source UDP port, and destination IP address must be identical to the same parameters configured with the `mirror endpoint ip` command on the remote destination switch.
- The configured remote exit port must not be a member of a trunk or mesh.
- If the destination for mirrored traffic is on a different VLAN than the source, routing must be correctly configured along the path from the source to the destination.
- On the remote destination (endpoint) switch, the IP addresses of the remote exit port and the switch can belong to different VLANs.
- All links on the path from the source switch to the destination switch must be active.



CAUTION:

A mirroring exit port should be connected only to a network analyzer, IDS, or other network edge device that has no connection to other network resources. Configuring a mirroring exit port connection to a network can result in serious network performance problems, and is strongly discouraged by HP Switch Networking.

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/support
- 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:
 - 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 [“Port Status and Configuration” \(page 48\)](#). 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*.
- 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 the *Multicast and Routing Guide*.
- 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*.

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 switch console 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 nondefault 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:

Figure 220 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
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
```

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.

Figure 221 Examples of correctly and incorrectly specifying a single host

```
Switch(config)# access-list 6 permit host 10.28.100.100 ← Correct.
Switch(config)# access-list 6 permit host 10.28.100.100 [255.255.255.255]
Invalid input: 255.255.255.255
Switch(config)# access-list 6 permit host 10.28.100.100 [32]
Invalid input: 10.28.100.100/32
```

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 the *Access Security Guide*.

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.

Examples

Remote gateway case

Configuring ACL "101" (Example 65 (page 384)) and applying it outbound on VLAN 1 in Figure 222 (page 384) 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.

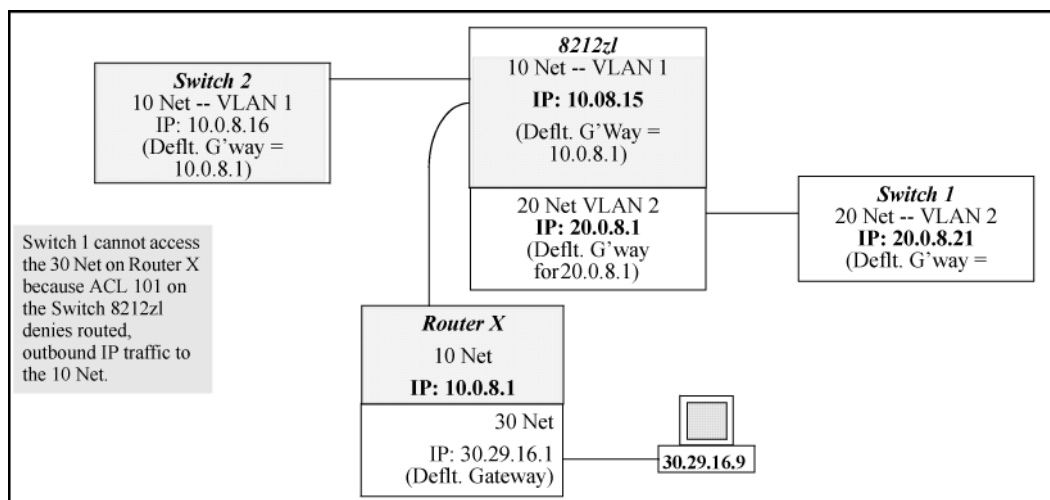
Example 65 ACE blocking an entire subnet

In [Figure 222 \(page 384\)](#), 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.)

```
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 0.0.0.0 255.255.255.255  
exit
```

Figure 222 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 \(page 384\)](#).

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 387).

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 the *Access Security Guide*.

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 the *Access Security Guide*.

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 the *Access Security Guide*.

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.

Figure 223 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  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
```

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.

Figure 224 Displaying encryption keys

```
HP Switch(config)# show radius
Status and Counters - General RADIUS Information
  Deadtme(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key : My-Global-Key
  Dynamic Authorization UDP Port : 3799

Server IP Addr      Auth Acct DM/ Time
Port Port CoA Window Encryption Key
-----
10.33.18.119      1812 1813
                               119-only-key
```

Also, ensure that the switch port used to access the RADIUS server is not blocked by an 802.1X configuration on that port. For example, `show port-access authenticator port-list` gives you the status for the specified ports. Also, ensure that other factors, such as port security or any 802.1X configuration on the RADIUS server are not blocking the link.

The authorized MAC address on a port that is configured for both 802.1X and port security either changes or is re-acquired after execution of `aaa port-access authenticator port-list initialize`

If the port is force-authorized with `aaa port-access authenticator port-list control authorized` command and port security is enabled on the port, then executing `initialize` causes the port to clear the learned address and learn a new address from the first packet it receives after you execute `initialize`.

A trunked port configured for 802.1X is blocked

If you are using RADIUS authentication and the RADIUS server specifies a VLAN for the port, the switch allows authentication, but blocks the port. To eliminate this problem, either remove the port from the trunk or reconfigure the RADIUS server to avoid specifying a VLAN.

QoS-related problems

Loss of communication when using VLAN-tagged traffic

If you cannot communicate with a device in a tagged VLAN environment, ensure that the device either supports VLAN tagged traffic or is connected to a VLAN port that is configured as `untagged`.

Radius-related problems

The switch does not receive a response to RADIUS authentication requests

In this case, the switch attempts authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH.)

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use `ping` to ensure that the switch has access to the configured RADIUS server.
- Verify that the switch is using the correct encryption key for the designated server.
- Verify that the switch has the correct IP address for the RADIUS server.
- Ensure that the `radius-server timeout` period is long enough for network conditions.
- Verify that the switch is using the same UDP port number as the server.

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.

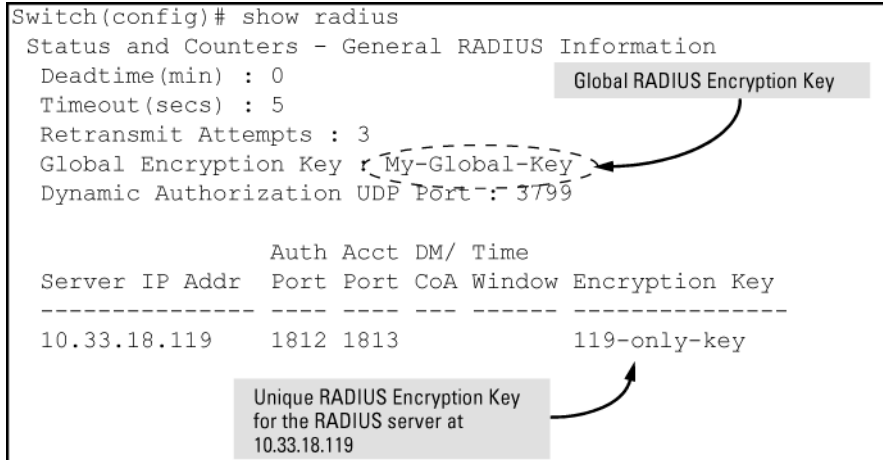
RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch

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.

Figure 225 Examples of 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
  Dynamic Authorization UDP Port : 3799

Server IP Addr      Auth Acct DM/ Time
Port Port CoA Window Encryption Key
-----
10.33.18.119      1812 1813
                               119-only-key
```



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 the *Advanced Traffic Management Guide*.

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 the *Access Security Guide*.)

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 CRLF.
- 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 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.

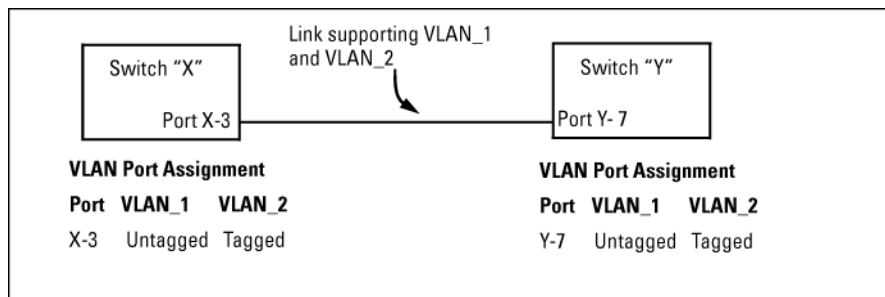
None of the devices assigned to one or more VLANs on an 802.1Q-compliant switch are being recognized

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 226 \(page 392\)](#).

Figure 226 Example of correct VLAN port assignments on a link



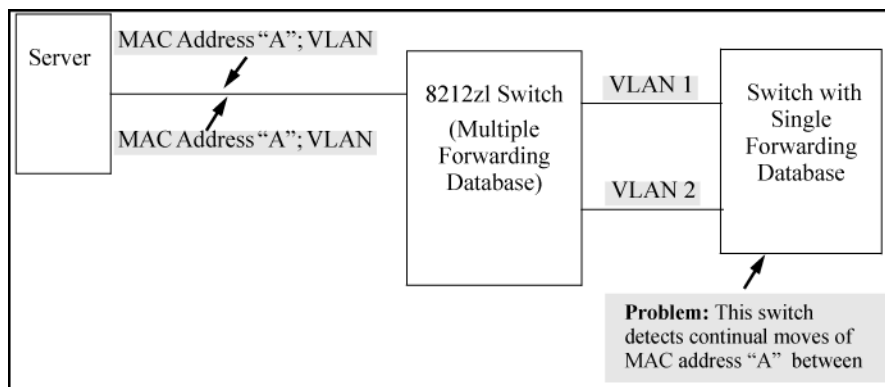
- If VLAN_1 (VID=1) is configured as "Untagged" on port 3 on switch "X," it must also be configured as "Untagged" on port 7 on switch "Y." Make sure that the VLAN ID (VID) is the same on both switches.
- Similarly, if VLAN_2 (VID=2) is configured as "Tagged" on the link port on switch "A," it must also 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 227 Example of duplicate MAC address



Fan failure

When two or more fans fail, a two-minute timer starts. After two minutes, the switch is powered down and must be rebooted to restart it. This protects the switch from possible overheating.

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.
```

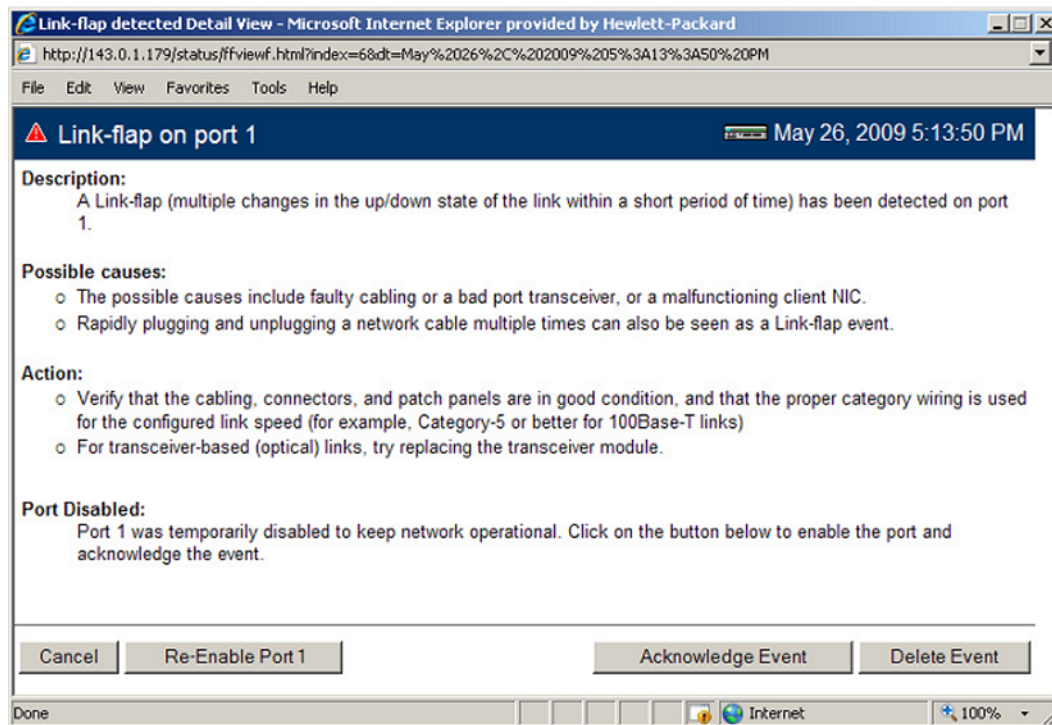
```
02674 FFI: port number-Administrator action required to re-enable.
```

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 228 \(page 394\)](#). 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 228 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 transceivers are supported:

Product #	Description
J8436A	10GbE X2-SC SR Optic
J8437A	10GbE X2-SC LR Optic
J4858A	Gigabit-SX-LC Mini-GBIC
J4858B	Gigabit-SX-LC Mini-GBIC
J4858C ¹	Gigabit-SX-LC Mini-GBIC

¹ The J4858C transceivers that have the letters "VM" in their serial numbers do not support DOM.

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.

Example 66 Example of output for a specified transceiver

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.

```
HP Switch(config)# show interfaces transceiver 21
```

Transceiver Technical information:

Port	Type	Product Number	Serial Number	Part Number
21	1000SX	J4858C	MY050VM9WB	1990-3657

Example 67 Example of output when no transceiver is present in specified interface

If there is no transceiver in the port specified in the command, the output displays as shown below.

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

```
No Transceiver found on interface 22
```

Example 68 Example of output when no ports are specified

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

```
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	

Example 69 Example of output when "all" is specified

You can specify all for `port-list` as shown below.

```
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 25 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 26 \(page 396\)](#), [Table 27 \(page 396\)](#), and [Table 28 \(page 397\)](#) is only displayed when the transceiver supports DOM.

Table 26 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 27 \(page 396\)](#).

Table 27 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

Table 27 Alarm and error information (GBIC/SFP transceivers only) (continued)

Alarm	Description
Voltage High	Voltage is high
Voltage Low	Voltage is low

The alarm information for XENPAK transceivers is shown in [Table 28 \(page 397\)](#).

Table 28 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
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

Example 70 Example of detailed information for a 1000SX Mini-GBIC transceiver

An example of the output for the show interfaces transceiver [port-list] detail for a 1000SX transceiver is shown below.

```
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
```

Example 71 Example of detailed information for a 10GbE-LR transceiver

An example of the output for a 10GbE-LR transceiver is shown below.

```
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 2011
```

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 72 Example of 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 73 Example of 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 29 General transceiver information

Parameter	Description
Interface Index	The switch interface number
Transceiver-type	Pluggable transceiver type
Transceiver model	Pluggable transceiver model

Table 29 General transceiver information (continued)

Parameter	Description
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.
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

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.

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.

Starting in software release K.13.xx, the maximum number of entries supported in the Event Log is increased from 1000 to 2000. Entries are listed in chronological order, from the oldest to the most recent.

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 229 \(page 401\)](#), each Event Log entry is composed of six or seven fields, depending on whether numbering is turned on or not:

Figure 229 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 <code>[no] log-number</code> command.
System module	The internal module (such as "ports:" for port manager) that generated a log entry. If VLANs are configured, a VLAN name also appears for an event that is specific to an individual VLAN. Table C-2 (page 402) lists the different system modules with a description of each one.
Management module	(8200zl switches) is either the active management module represented by AM1 or AM2, or the standby management module, represented by SM1 or SM2.
Event message	A brief description of the operating event.

Table 30 Event Log system modules

System module	Description	Documented in HP Switch hardware/software guide
802.1x	<p>802.1X authentication: Provides access control on a per-client or per-port basis:</p> <ul style="list-style-type: none"> • Client-level security that allows LAN access to 802.1X clients (up to 32 per port) with valid user credentials • Port-level security that allows LAN access only on ports on which a single 802.1X-capable client (supplicant) has entered valid RADIUS user credentials 	<i>Access Security Guide</i>
acl	<p>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.</p>	<i>Advanced Traffic Management Guide</i>
addrmgr	<p>Address Table Manager: Manages MAC addresses that the switch has learned and are stored in the switch's address table.</p>	<i>Management and Configuration Guide</i>
arp-protect	<p>Dynamic ARP Protection: Protects the network from ARP cache poisoning. Only valid ARP requests and 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.</p>	<i>Access Security Guide</i>
auth	<p>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.</p>	<i>Access Security Guide</i>
cdp	<p>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.</p>	<i>Management and Configuration Guide</i>
chassis	<p>Hardware operation, including modules and ports, power supply, fans, transceivers, CPU interrupt errors, switch temperature, and so on. Chassis messages include events on POE operation.</p>	<i>Installation Guides 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</p>	<i>Access Security Guide</i>

Table 30 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
	<p>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>	
console	<p>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.</p>	<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	<p>Dynamic Configuration Arbiter (DCA) determines the client-specific parameters that are assigned in an authentication session.</p>	<i>Access Security Guide</i>
dhcp	<p>Dynamic Host Configuration Protocol (DHCP) server configuration: Switch is automatically configured from a DHCP (Bootp) server, including IP address, subnet mask, default gateway, Timep Server address, and TFTP server address.</p>	<i>Management and Configuration Guide</i>
dhcp v6c	<p>DHCP for IPv6 prefix assignment</p>	<i>IPv6 Configuration Guide</i>
dhcpr	<p>DHCP relay: Forwards client-originated DHCP packets to a DHCP network server.</p>	<i>Advanced Traffic Management Guide</i>
download	<p>Download operation for copying a software version or files to the switch.</p>	<i>Management and Configuration Guide</i>
dhcp-snoop	<p>DHCP snooping: Protects your network from common DHCP attacks, such as address spoofing and repeated address requests.</p>	<i>Access Security Guide</i>
dma	<p>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.</p>	—
fault	<p>Fault Detection facility, including response policy and the sensitivity level at which a network problem should generate an alert.</p>	<i>Management and Configuration Guide</i>

Table 30 Event Log system modules (continued)

System module	Description	Documented in HP Switch hardware/software guide
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 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>
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>

Table 30 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
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>
l2l3	<p>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.</p> <p>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.</p>	<i>Management and Configuration Guide</i> <i>Advanced Traffic Management Guide</i>
lldp	Link-Layer Discovery Protocol: Supports transmitting LLDP packets to neighbor devices and reading LLDP packets received from neighbor devices, enabling a switch to advertise itself to adjacent devices and to learn about adjacent LLDP devices.	<i>Management and Configuration Guide</i>
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	<p>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:</p> <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	<p>MAC lockdown and MAC lockout</p> <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 	<i>Access Security Guide</i>

Table 30 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
	<p>assigned port on the switch. MAC Lockdown also restricts the client device to a specific VLAN.</p> <ul style="list-style-type: none"> • MAC lockout blocks a specific MAC address so that the switch drops all traffic to or from the specified address. 	
mgr	<p>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.</p>	<i>Management and Configuration Guide</i>
mld	<p>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.</p>	<i>Multicast and Routing Guide</i>
mtm	<p>Multicast Traffic Manager (MTM): Controls and coordinates L3 multicast traffic for upper layer protocols.</p>	<i>Multicast and Routing Guide</i>
netinet	<p>Network Internet: Monitors the creation of a route or an Address Resolution Protocol (ARP) entry and sends a log message in case of failure.</p>	<i>Advanced Traffic Management Guide</i>
ospf	<p>Open Short Path First (OSPF): 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.</p>	<i>Multicast and Routing Guide</i>
pagp	<p>Ports Aggregation Protocol (PAgP): Obsolete. Replaced by LACP (802.3ad.) Not used for logging messages in software release K.13.xx.</p>	—
pim	<p>Protocol-independent multicast (PIM) routing: Enables IP multicast traffic to be transmitted for multimedia applications throughout a network without being blocked at routed interface (VLAN) boundaries.</p>	<i>Multicast and Routing Guide</i>
ports	<p>Port status and port configuration features, including mode (speed and duplex), flow control, broadcast limit, jumbo packets, and security settings.</p> <p>Port messages include events on POE operation and transceiver connections with other network devices.</p>	<i>Installation and Getting Started Guide</i> <i>Management and Configuration Guide</i> <i>Access Security Guide</i>

Table 30 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
QinQ	IEEE 802.1ad specification, known as QinQ (provider bridging), provides a second tier of VLANs in a bridged network. QinQ supports the forwarding of traffic from multiple customers over a provider network using service VLANs (S-VLANs.)	<i>Advanced Traffic Management 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 station clients capable of SSL/TLS operation.	<i>Access Security Guide</i>
stack	Stack management: Uses a single IP address and standard network cabling to manage a group (up to 16) of	<i>Advanced Traffic Management Guide</i>

Table 30 Event Log system modules (continued)

System module	Description	Documented in HP Switch hardware/software guide
	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.	
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>
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>

Table 30 Event Log system modules *(continued)*

System module	Description	Documented in HP Switch hardware/software guide
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>
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.	<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: Routing protocol not used for logging messages in software release K.13.xx.	—

Displaying and navigating in the Event Log

Displaying 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.
<i>option-str</i>	Displays all Event Log entries that contain the specified text. Use an <i>option-str</i> 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
```

Using the Menu

To display the Event Log from the Main Menu, select Event Log. [Figure 230 \(page 411\)](#) shows a sample event log display.

Figure 230 Example of an event log display

```
HP Switch 5406z1                               25-Oct-2007  18:02:52
=====CONSOLE - MANAGER MODE -
=====
M 10/25/07 16:30:02 sys: 'Operator cold reboot from CONSOLE session.'
I 10/25/07 17:42:51 00061 system: -----
-
I 10/25/07 17:42:51 00063 system: System went down: 10/25/07 16:30:02
I 10/25/07 17:42:51 00064 system: Operator cold reboot from CONSOLE session.
W 10/25/07 17:42:51 00374 chassis: WARNING: SSC is out of Date: Load 8.2 or
newer
I 10/25/07 17:42:51 00068 chassis: Slot D Inserted
I 10/25/07 17:42:51 00068 chassis: Slot E Inserted
I 10/25/07 17:42:51 00068 chassis: Slot F Inserted
I 10/25/07 17:42:51 00690 udpf: DHCP relay agent feature enabled
I 10/25/07 17:42:51 00433 ssh: Ssh server enabled
I 10/25/07 17:42:52 00400 stack: Stack Protocol disabled
I 10/25/07 17:42:52 00128 tftp: Enable succeeded
I 10/25/07 17:42:52 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.

Navigating in the Event Log

Using the Menu

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 \(page 411\)](#).

Table 31 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.

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.

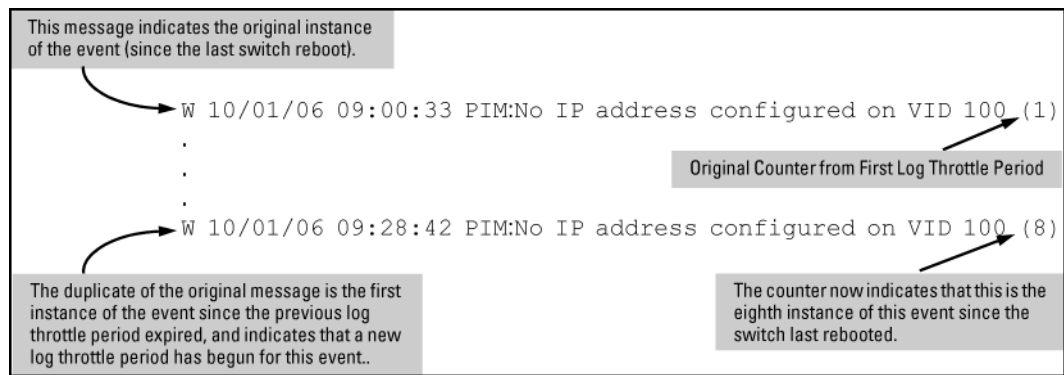
Example 74 The first instance of an event message and counter

The counter (1) indicates that this is the first instance of this event since the switch last rebooted.

```
W 10/01/12 09:00:33 PIM:No IP address configured on VID 100 (1)
```

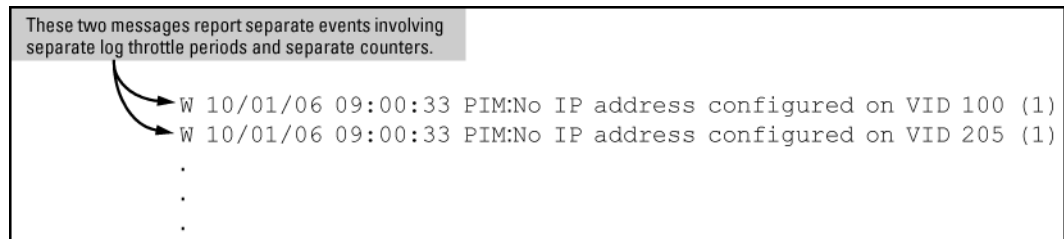
If PIM operation causes the same event to occur six more times during the initial log throttle period, there are no further entries in the Event Log. However, if the event occurs again after the log throttle period has expired, the switch repeats the message (with an updated counter) and starts a new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

Figure 231 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 232 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 \(page 413\)](#). (The same operation applies for messages sent to any configured SNMP trap receivers.)

Table 32 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.

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:
 - ACL "deny" matches
 - Dynamic ARP protection events
 - DHCP snooping events
 - DIPLD events
 - Events recorded in the switch's Event Log
 - IPv4 routing events
 - LACP events
 - LLDP events
 - OSPF events
 - PBR class events
 - SNMP events
 - SSH events
 - VRRP 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:
 - 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 [“Enabling or disabling syslog messaging” \(page 426\)](#) and [“Configuring a syslog server” \(page 428\)](#).

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.
	neighbor-adjacency [detail]	Enables or disables OSPFv3 (IPv6) adjacency logging. Must be executed in OSPFv3 context. The detail option displays all the adjacency state transitions and adjacency-related errors.
	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.)
	system-module	Sends Event Log messages from the specified system module to configured debug destinations. The severity filter is also applied to the system-module messages you select. The default setting is to send Event Log messages from all system modules. To restore the default setting, enter the no logging system-module system-module or logging system-module all-pass commands.
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 syslog-ip-addr 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 debug-type command to a buffer in switch memory. For more information on these options, see "Enabling or disabling syslog messaging" (page 426) .
	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	<p>fib: Displays IP Forwarding Information Base messages and events.</p> <p>forwarding: Sends IPv4 forwarding messages to the debug destinations.</p> <p>ospf: Sends OSPF event logging to the debug destinations.</p> <p>ospfv3: Enables debug messages for OSPFv3.</p> <p>packet: Sends IPv4 packet messages to the debug destinations.</p> <p>pim [packet [filter source ip-addr vlan vid]] : 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>pbr: Logs a message when a PBR policy is applied, when the action in a class goes active or when it goes inactive.</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>
		<p>ospf3 [adj event flood lsa-generation packet retransmission spf] : Sends OSPFv3 events to the debug destinations. Must be executed in</p>

		OSPFv3 context. Selecting an option filters the debug messages by that option. <code>packet</code> : Sends IPv6 packet messages to the debug destinations.
	<code>lldp</code>	Sends LLDP debug messages to the debug destinations.
	<code>lacp</code>	<code>event</code> : Sends messages related to change events. <code>packet</code> : Sends messages when BPDUs are exchanged.
	<code>security</code>	Sends security messages to the debug destination.
	<code>services</code>	Displays debug messages on the services module.
	<code>snmp</code>	Sends snmp messages to the debug destination.
	<code>vrrp</code>	Turns on tracing of the incoming and outgoing VRRP packets and sends debug logging to the debug destination.
	<code>wireless-services</code>	Sends wireless service module debug 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 417)" 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 418) 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 C-2 \(page 402\)](#).

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.)

Figure 233 Sample output of 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 234 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 234 \(page 419\)](#), 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 235 \(page 420\)](#).

Figure 235 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
```

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” \(page 428\)](#).

Configuring the types of debug messages that the switch can send to configured debug destinations

Syntax

```
[ no ] debug debug-type
```

acl	<p>When a match occurs on an ACL "deny" ACE (with <code>log</code> configured), the switch sends an ACL message to configured debug destinations. For information on ACLs, see the "Access Control Lists (ACLs)" chapter in the latest version of the following guides:</p> <p>IPv4 ACLs: <i>Access Security Guide</i></p> <p>IPv6 ACLs: <i>IPv6 Configuration Guide</i></p> <p>NOTE: Beginning with software release K.14.01, 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><code>logging</code>—Disables or re-enables syslog logging on one or more syslog servers configured with the <code>logging syslog-ip-addr</code> command.</p> <p><code>session</code>—Assigns or re-assigns destination status to the terminal device that was most recently used to request debug output.</p> <p><code>buffer</code>—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>For more information on these options, see “Enabling or disabling syslog messaging” (page 426).</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>
<code>ip [fib forwarding ospf ospfv3 packet pim rip]</code>	Sends IP messages to configured destinations.
<code>ip [fib [events]]</code>	For the configured debug destinations: <code>events</code> —Sends IP forwarding information base events.
<code>ip [ospf [adj event flood lsa-generation packet <i>packet-type</i> retransmission spf]]</code>	<p>For the configured debug destinations:</p> <p><code>ospf</code>—Enables the specified IP-OSPF message type.</p> <p><code>adj</code>—Adjacency changes.</p> <p><code>event</code>—OSPF events.</p> <p><code>flood</code>—Information on flood messages.</p> <p><code>lsa-generation</code>—New LSAs added to database.</p> <p><code>packet [packet-type]</code> — All OSPF packet messages sent and received on the switch, where <code>packet-type</code> enables only the specified OSPF packet type. Valid values are:</p> <p><code>dd</code>—Database descriptions</p> <p><code>hello</code>—Hello messages</p> <p><code>lsa</code>—Link-state advertisements</p> <p><code>lsr</code>—Link-state requests</p> <p><code>lsu</code>—Link-state updates</p> <p><code>retransmission</code>—Retransmission timer messages.</p> <p><code>spf</code>—Path recalculation messages.</p>
<code>ip [ospfv3]</code>	Enables OSPFv3 debug messages.
<code>ip [pim [packet filter source <i>ip-addr</i> vlan <i>vid</i>]]</code>	<p>For the configured debug destinations:</p> <p><code>packet</code>— Enables the specified PIM message type</p> <p><code>filter</code>— Enables or disables tracing of PIM messages filtered on VLAN or source group information.</p>

	<p>source <i>ip-addr</i> — Displays packets from a specific source to a specific group. Only a single source/group filter is supported.</p> <p>vlan <i>vlan-id</i>— Enables or disables tracing on a specified VLAN for PIM</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 high-stress production network.</p>
<pre>ip[rip [database event trigger]]</pre>	
	<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 forwarding nd ospfv3 packet]</pre>	
	<p>NOTE: See <i>IPv6 Configuration Guide</i>.</p> <p>When no debug options are included, displays debug messages for all IPv6 debug options.</p> <p>dhcpv6-client [events packe] —Displays DHCPv6 client event and packet data.</p> <p>dhcpv6-relay [events packet] —Displays DHCPv6 relay event and relay packet data.</p> <p>forwarding— Displays IPv6 Forwarding Information Base messages.</p> <p>nd—Displays debug messages for IPv6 neighbor discovery.</p>

	<p>ospfv3—Enables the specified IPv6-OSPF message type.</p> <ul style="list-style-type: none"> adj— Adjacency changes. event— OSPFv3 events. flood— Information on flood messages. lsa-generation— New link state advertisements added to database. packet [packet-type] —All OSPFv3 packet messages sent and received on the switch, where packet-type enables only the specified OSPFv3 packet type. Valid values are: <ul style="list-style-type: none"> dd— Database descriptions hello— Hello messages lsa— Link-state advertisements lsr— Linkstate requests lsu— Link-state updates retransmission—Retransmission timer messages. spf—Path recalculation messages. packet—Displays IPv6 packet messages.
lldp	Enables all LLDP message types for the configured destinations.
<pre>security[arp-protect dhcp-snooping dynamic-ip-lockdown port-access port-security radius-server ra-guard 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> <ul style="list-style-type: none"> agent—Displays DHCP snooping agent messages. event—Displays DHCP snooping event messages. packet—Displays DHCP snooping packet messages. <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>ra-guard—Sends blocked RAs and redirects 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><code>tacacs-server</code>—Sends TACACS debug messages to the debug destination.</p> <p><code>user-profile-mib</code>—Sends user profile MIB debug messages to the debug destination.</p>
<code>services slot-id-range</code>	
	Displays debug messages on the services module. Enter an alphabetic module ID or range of module IDs for the <code>slot-id-range</code> parameter.
<code>snmp event pdu routines</code>	
	<p>Displays the SNMP debug messages.</p> <p><code>event</code>—Displays SNMP event debug messages.</p> <p><code>pdu</code>—Displays SNMP pdu debug messages.</p> <p><code>routines</code>—Displays SNMP routines debug messages</p>
<code>vrrp</code>	Displays VRRP debug messages on the configured destinations.
<code>wireless-services slot-id-range</code>	
	Displays wireless-services debug messages on the wireless services module. Enter an alphabetic module ID or range of IDs for the <code>slot-id-range</code> parameter.

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

Figure 236 Example of 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
```

Figure 237 Example of 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
```

Figure 238 Example of 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
```

Enabling or disabling syslog messaging

When a syslog server is configured, the forwarding of events begins immediately. The commands **no debug event** or **no debug all** have no effect. The only way to disable the forwarding of events to the Syslog server is by removing the server with the no logging <ip-address> command, or the **no logging** command, which removes all Syslog servers.

Syntax

```
[no]debug event
[no]debug all
```

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.

The **all** option sends debug logging to the configured debug destinations for all ACL, Event Log, IP-OSPF, and IP-RIP options.

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
```

Enables syslog logging to configured syslog servers so that the debug message types specified by the `debug debug-type` command (see "Debug Messages" on page A-47) are sent.

(Default: Logging disabled)

To configure a syslog server IP address, see "[Configuring a syslog server](#)" (page 428).

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” \(page 438\)](#).

<p>session</p>	<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. Event message types received on the selected CLI session are configured with the <code>debug debug-type</code> command.</p>
<p>buffer</p>	<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 `loggingcommand` 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" on page A-47.)
- `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 437) and "Configuring the system module used to select the Event Log messages sent to a syslog server" (page 437).)

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 414).)

Syntax

```
[ no ] logging syslog-ip-addr
```

Enables or disables syslog messaging to the specified IP address. You can configure up to six addresses. If you configure an address when none are already configured, this command enables destination logging (syslog) and the Event debug type.

Therefore, at a minimum, the switch begins sending Event Log messages to configured syslog servers. The ACL, IP-OSPF, and/or IP-RIP message types are also

sent to the syslog servers if they are currently enabled as debug types. (See "Debug Messages" on page A-47.)

<code>no logging</code>	Removes all currently configured syslog logging destinations from the running configuration. Using this form of the command to delete the only remaining syslog server address disables debug destination logging on the switch, but the default Event debug type does not change.
<code>no logging syslog-ip-address</code>	Removes only the specified syslog logging destination from the running configuration. Removing all configured syslog destinations with the <code>no logging</code> command (or a specified syslog server destination with the <code>no logging syslog-ip-address</code> command) does not delete the syslog server IP addresses stored in the startup configuration.

Deleting syslog addresses in the startup configuration

Enter a `no logging` command followed by the `write memory` command.

Verifying the deletion of a syslog server address

Display the startup configuration by entering the `show config` command.

Blocking the messages sent to configured syslog servers from the currently configured debug message type

Enter the `no debug debug-type` command. (See "Debug Messages" on page A-47.)

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.

Examples

Figure 239 Configuring UDP for logging message transmission using the default port

```
HP Switch(config)# logging 192.123.4.5 tcp
```

Default TCP port 1470 is used.

Figure 240 Configuring TCP for logging message transmission using a specified port

```
HP Switch(config)# logging 192.123.4.5 tcp 9514
```

TCP port 9514 is used.

Figure 241 Configuring UDP for logging message transmission using the default port

```
HPswitch(config)# logging 192.123.4.5 udp
```

Default UDP port 514 is used.

Figure 242 Configuring UDP for logging message transmission using a specified port

```
HPswitch(config)# logging 192.123.4.5 udp 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.

For a list of supported HP switches, see the Note on page A-41.

Disable LinkUp/Down Syslog Messages Based on Port

This feature provides a per-port basis filter that can restrict the logging of events that are associated with a link status change. Unimportant linkup/linkdown events can be filtered out, avoiding unwanted messages in the event log and reducing troubleshooting time.

The specific port-based events to be controlled are:

RMON_PMGR_PORT_UP—Indicates that the port has changed from an off-line to an on-line state. To be online the port must be both connected to the LAN and enabled through configuration.

RMON_PMGR_PORT_DOWN—Indicates that the port has changed from an on-line state to an offline state. For this state to occur, the port is physically disconnected from the LAN, disabled through the configuration, or both.

The following rules apply:

- Only one filter can be enabled at a time.
- The maximum number of configured filters is 10.
- A filter is identified by a unique name of up to 16 printable ASCII characters.
- Filters can be dynamically replaced; the newly enabled filter automatically disables the previous filter.
- A filter always contains a default sub-filter that functions as the filtering rules terminator.
- To apply filtering to an event logging process, the filter must be explicitly enabled from the CLI.
- Enabled filter modules can be dynamically modified; the changes will take effect immediately.

A filter module may include up to 19 option sub-filters and a default sub-filter. The sub-filter types are:

- **Severity**—checks the severity level of the event log message. The severity values are:
 - major
 - warning
 - error
 - info
 - debug
- **Event number**—Checks the event number of the event log message.
- **Regular expression**—Checks everything beyond the date/time portion of the event log message.

A sub-filter has a sequence number, criteria to be matched, and a resulting action when a match occurs. All of the parameters must be specified in order to create the sub-filter.

- **Sequence number**: Used for the ordering of sub-filters. Range 1-98.
- **Matching criteria**: Can be the severity level, event number, or a regular expression.
- **Action to execute**: When a match occurs, the resulting action is either **permit** the logging of the event, or **deny** the logging of the event.

The following sub-filter rules apply:

- Up to 19 optional sub-filters and a default sub-filter are allowed in a filter module.
- Sub-filters in the filter module can be of the same or different types.
- Sub-filter entries can be modified with new criteria and action definitions.
- Sub-filters are executed from the lowest sequence number to the highest. As soon as a match is found the log event is immediately accepted or rejected and no further matching operation is performed.

- The default sub-filter must always be the last entry in a filter module. It functions as the rules terminator when the criteria matching performed by the prior sub-filters in a filter does not produce an action.
- The default sub-filter cannot be deleted, re-ordered, or changed. The only parameter that can be modified is the action parameter of **permit** or **deny**. The default is **permit**.

Creating a Filter

Syntax

```
[no] logging filter <name> <sequence> {severity <severity> | event-num <num> | <regexp>}
{permit | deny}
```

Creates a logging filter to restrict which events are logged. The **no** form of the command removes the logging filter.

<name>: The name that identifies the filter.

severity <severity>: Specifies the severity of an event—major, warning, error, info, or debug.

event-num <num>: Specifies an event number to match.

deny: If the log entry matches the specified criteria, do not log the event message. No further criteria are evaluated for a match.

permit: If the log entry matches the specified criteria, log the event message. No further criteria are evaluated for a match.

Enabling a Filter after Creation

Syntax [no] logging filter <name> enable | disable

Enables a log filter. Only one filter can be enabled at a time. An enabled filter automatically disables a previously enabled filter.

<name>: The name that identifies the filter.

Clearing a Filter

Syntax clear logging filter {<name> | all}

Clears statistics counters for the named logging filter or for all filters.

Displaying Filter Configuration Information

Syntax show logging filter [<name>]

Displays the logging filter's configuration information. The Matches column indicates the number of times that criteria has matched.

Example

Figure 243 Specifying the Criteria for a Filter and then Enabling the Filter

```
HP Switch(config)# logging filter SevWarnFatal 10 severity warning permit
HP Switch(config)# logging filter SevWarnFatal 20 severity major permit
HP Switch(config)# logging filter SevWarnFatal default deny

HP Switch(config)# logging filter SevWarnFatal enable
```

1. The filter named SevWarnFatal adds a sub-filter of the severity type, with a sequence number of 10. The sub-filter specifies that a match for an event log message with a severity of "warning" will be logged.
2. The second sub-filter has a sequence number of 20 and a severity type of major. The sub-filter specifies that a match for an event log message with a severity of "major" will be logged.
3. The default sub-filter, which is created automatically at the time of filter creation, is always the last entry in the filter module. It matches "anything" and cannot be changed. You can change

the actions to either **permit** or **deny**. This example specifies that any message that did not meet the prior matching criteria will not be logged.

4. The last command enables the filter named SevWarnFatal. If there was another filter enabled, this filter automatically replaces it and the other filter is disabled.

Example

Figure 244 Specifying the Criteria for a Filter named noUpDownEvents and then Enabling the Filter

```
HP Switch(config)# logging filter noUpDownEvents 10 event-num 76 deny
HP Switch(config)# logging filter noUpDownEvents 20 event-num 77 deny
HP Switch(config)# logging filter noUpDownEvents default permit

HP Switch(config)# logging filter noUpDownEvents enable
```

1. The filter named noUpDownEvents adds a sub-filter with a type of event-num, and a sequence number of 10. The sub-filter specifies that a match for an event log message with an event number of "76" will not be logged.
2. The second sub-filter has a sequence number of 20 and a type of event-num. The sub-filter specifies that a match for an event log message with an event number of "77" will not be logged.
3. The default sub-filter, which is created automatically at the time of filter creation, is always the last entry in the filter module. It matches "anything" and cannot be changed. You can change the actions to either **permit** or **deny**. This example specifies that any message that did not meet the prior matching criteria will be logged.
4. The last command enables the filter named noUpDownEvents. If there was another filter enabled, this filter automatically replaces it and the other filter is disabled.

Example

Figure 245 Specifying the Criteria for a Match using a Regular Expression and then Enabling the Filter

```
HP Switch(config)# logging filter noUpPorts 10 "(A10|A22|B5) is now on-line" deny
HP Switch(config)# logging filter noUpPorts default permit

HP Switch(config)# logging filter noUpPorts enable
```

This example denies logging of the matching regular expression "port <port-num> is now on-line" for ports A10, A22, and B5.

1. The filter named noUpPorts adds a sub-filter with a type of regular expression for ports A10, A22, and B5. The sub-filter specifies the matching criteria for the regular expression and if there is a match, the event log message is not logged.
2. The default sub-filter specifies that any message that did not meet the prior matching criteria will be logged.
3. The last command enables the filter named noUpPorts.

Example

Figure 246 Specifying the Criteria for a Match using a Regular Expression for Specific Ports

```
HP Switch(config)# logging filter noStpBlockPorts 10 "(A[1-9]|A10|B[1-4])
.*Blocked by STP" permit
HP Switch(config)# logging filter noStpBlockPorts 20 ".*Blocked by STP" deny
HP Switch(config)# logging filter noStpBlockPorts default permit

HP Switch(config)# logging filter noStpBlockPorts enable
```

1. The filter named noStpBlockPorts adds a sub-filter with a type of regular expression with a sequence number of 10. This rule specifies that event messages from ports A1-A10, and B1-B4 with the ".*Blocked by STP" expression pattern in the message body are logged.
2. The second command adds a sub-filter with a type of regular expression and a sequence number of 20. This rule specifies that event messages generated from **any** ports with the ".*Blocked by STP" expression pattern in the message body are not logged.
3. The default sub-filter specifies that any message that did not meet the prior matching criteria will be logged.
4. The last command enables the filter named noStpBlockPorts.

Output Examples

The following example displays the **show** command output for displaying the configured logging filters.

Figure 247 The Configured Logging Filters

```
HP Switch# show logging filter

Status and Counters - Log Filters Information

Name                Enabled
-----
noUpPorts           No
SevWarnFatal        No
noUpDownEvents      No
noStpBlockPorts     Yes
```

The following example displays the output when specific filters are specified.

Figure 248 Output for Specified Logging Filters

```
HP Switch# show logging filter sevWarnFatal

Status and Counters - Log Filter Information

Name           : SevWarnFatal
Enabled        : Yes
Messages Dropped : 0

Seq Type      Value                               Action Matches
-----
10 Severity   warning                               Permit 2
20 Severity   major                               Permit 2
def           (any)                               Deny 0

HP Switch(config)# show logging filter noStpBlockPorts

Status and Counters - Log Filter Information

Name           : noStpBlockPorts
Enabled        : Yes
Messages Dropped : 0

Seq Type      Value                               Action Matches
-----
10 RegExp     (A[1-9]|A10|B[1-4]) .*Blocked by STP Permit 4
20 RegExp     .*Blocked by STP                Deny 2
def           (any)                               Permit 0
```

The following example displays the running-config output.

Figure 249 Output of running-config File

```
HP Switch# show running-config

Running configuration:

; J9470A Configuration Editor; Created on release #K.15.13.0000x
; Ver #04:0f.ff.3f.ef:24
hostname "HP Switch"
module 1 type j94dda
logging filter "noUpPorts" 10 "(A10|A22|B5) is now on-line" deny
logging filter "noUpPorts" default permit
logging filter "SevWarnFatal" 10 severity warning permit
logging filter "SevWarnFatal" 20 severity major permit
logging filter "SevWarnFatal" default deny
logging filter "noUpDownEvents" 10 event-num 76 deny
logging filter "noUpDownEvents" 20 event-num 77 deny
logging filter "noUpDownEvents" default permit
logging filter "noStpBlockPorts" 10 "(A[1-9]|A10|B[1-4]) .*Blocked by STP" permit
logging filter "noStpBlockPorts" 20 " .*Blocked by STP" deny
logging filter "noStpBlockPorts" default permit
logging filter "noStpBlockPorts" enable
snmp-server community "public" unrestricted
snmp-server host 15.255.133.156 community "public"
snmp-server host 15.255.133.146 community "public"
vlan 1
.
.
.
```

Adding a description for a Syslog server

You can associate a user-friendly description with each of the IP addresses (IPv4 only) configured for syslog using the CLI or SNMP.

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 75 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

Example 76 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” \(page 428\)](#).

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” \(page 428\)](#).)

See [Table C-2 \(page 402\)](#) 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.

- Debug commands do not affect normal message output to the Event Log. Using the `debug event` command, you can specify that Event Log messages are sent to the debug destinations you configure (CLI session, syslog servers, or both) in addition to the Event Log.
- Ensure that your syslog servers accept debug messages. All syslog messages resulting from a debug operation have a "debug" severity level. If you configure the switch to send debug messages to a syslog server, ensure that the server's syslog application is configured to accept the "debug" severity level. (The default configuration for some syslog applications ignores the "debug" severity level.)
- Duplicate IP addresses are not stored in the list of syslog servers.
- If the default severity value is in effect, all messages that have severities greater than the default value are passed to syslog. For example, if the default severity is "debug," all messages that have severities greater than debug are passed to syslog.
- There is a limit of six syslog servers. All syslog servers are sent the same messages using the same filter parameters. An error is generated for an attempt to add more than six syslog servers.
- The **debug destination logging** command allows the forwarding of debug logging activities to the configured Syslog servers.
- Switch events can be sent to the current console or the debug logging buffer using the `debug event` command followed by the **debug destination session** or **debug destination buffer** command respectively.
- The **no logging <ip-address>** and **no logging** commands disable the sending of event logs to Syslog servers. These commands will not affect the results of the `debug event` and `debug destination <session | buffer | logging>` commands.
- If the Syslog server is configured, the **no debug destination logging** command does not disable the sending of events to Syslog. It does disable the sending of other debug logs to the Syslog server.
- If Syslog is configured, the **no debug event** or **no debug all** commands do not disable the sending of events to Syslog unless the debug destination is set to **buffer** or **session**.
- When debugging is finished and the **no debug all** command is entered, the commands being sent to the previously configured Syslog server(s) are not affected, and no additional commands

are needed. If both **debug event** and **debug destination logging** are configured, duplicate events are not sent to the configured Syslog servers.

- The show debug command won't display debug event or debug destination logging unless the debug **event** command is executed for **buffer** or **session** destinations.
- You can check on the configured Syslog settings with the **show syslog config** command. The Syslog settings also are visible in the output of **show running-config** or **show startup-config**.

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 Chapter 10, "Port Status and Configuration".

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 457).

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 250](#) (page 440).

Figure 250 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] [ip-option [
record-route | loose-source-route | strict-source-route |
include-timestamp | include-timestamp-and-address | include
timestamp-from ]] [tos 0-255]
```

```
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
<code>ip-option</code>	Specify an IP option, such as loose or strict source routing, or an include-timestamp option: <code>include-timestamp</code> : Adds the timestamp option to the IP header. The timestamp displays the amount of travel time to and from a host. Default: 9 <code>include-timestamp-and-address</code> : Records the intermediate router's timestamp and IP address. Default: 4 <code>include-timestamp-from</code> : Records the timestamp of the specified router addresses. <code>loose-source-route IP-addr</code> : The <code>loose-source-route</code> option prompts for the IP address of each source IP on the path. It allows you to specify the IP addresses that you want the ping packet to go through; the packet may go through other IP addresses as well. <code>record-route 1-9</code> : Displays the IP addresses of the interfaces that the ping packet goes through on its way to the destination and on the way back. When specified without loose or strict recording, the source route is not recorded. The source route is automatically recorded when loose or strict source routing is enabled. Default: 9 <code>strict-source-route IP-addr</code> : Restricts the ping packet to only those IP addresses that have been specified and no other addresses.
<code>tos 0-255</code>	Specifies the type of service to be entered in the header packet. Default: 0 (zero)

NOTE: For information about `ping6`, see *IPv6 Configuration Guide*.

Example

Figure 251 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 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 457).

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 252 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

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 ]] [dstport 1-34000] [srcport 1-34000] [ip-option [ record-route | loose-source-route | strict-source-route | include-timestamp | include-timestamp-and-address | include-timestamp-from ]] [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]	<p>For the current instance of <code>traceroute</code>, changes the minimum number of hops allowed for each probe packet sent along the route.</p> <ul style="list-style-type: none">• If <code>minttl</code> is greater than the actual number of hops, the output includes only the hops at and above the <code>minttl</code> threshold. (The hops below the threshold are not listed.)• If <code>minttl</code> matches the actual number of hops, only that hop is shown in the output.• If <code>minttl</code> is less than the actual number of hops, all hops are listed. <p>For any instance of <code>traceroute</code>, if you want a <code>minttl</code> value other than the default, you must specify that value.</p> <p>(Default: 1)</p>
[maxttl 1-255]	<p>For the current instance of <code>traceroute</code>, changes the maximum number of hops allowed for each probe packet sent along the route.</p> <p>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.</p> <p>For any instance of <code>traceroute</code>, if you want a <code>maxttl</code> value other than the default, you must specify that value.</p> <p>(Default: 30)</p>
[probes 1-5]	<p>For the current instance of <code>traceroute</code>, changes the number of queries the switch sends for each hop in the route.</p> <p>For any instance of <code>traceroute</code>, if you want a <code>probes</code> value other than the default, you must specify that value.</p> <p>(Default: 3)</p>
[source ip-addr vid loopback 0-7]	The source IPv4 address, VLAN ID, or Loopback address.
[dstport 1-34000]	Destination port.

[srcport 1-34000]	Source port.
[ip-option]	<p>Specify an IP option, such as loose or strict source routing, or an include-timestamp option:</p> <p>[include-timestamp]: Adds the timestamp option to the IP header. The timestamp displays the amount of travel time to and from a host. Default: 9</p> <p>[include-timestamp-and-address]: Records the intermediate router's timestamp and IP address. Default: 4</p> <p>[loose-source-route IP-addr]: Prompts for the IP address of each source IP on the path. It allows you to specify the IP addresses that you want the ping packet to go through; the packet may go through other IP addresses as well.</p> <p>[record-route 1-9]: Displays the IP addresses of the interfaces that the ping packet goes through on its way to the destination and on the way back.</p> <p>When specified without loose or strict recording, the source route is not recorded. The source route is automatically recorded when loose or strict source routing is enabled. Default: 9</p> <p>[strict-source-route IP-addr]: Restricts the ping packet to only those IP addresses that have been specified and no other addresses.</p> <p>[timeout 1-120]: For the current instance of traceroute, changes the timeout period the switch waits for each probe of a hop in the route. For any instance of traceroute, if you want a timeout value other than the default, you must specify that value. (Default: 5 seconds)</p>

NOTE: For information about traceroute6, see *IPv6 Configuration Guide*.

Halting an ongoing traceroute search

Press the **[Ctrl] [C]** keys.

Examples

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 253 A completed traceroute enquiry

```

HP Switch# traceroute 125.25.24.35
traceroute to 125.25.24.35 ,
  1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2      0 ms      0 ms      0 ms
 2 10.71.217.2       7 ms      3 ms      0 ms
 3 10.243.170.1     0 ms      1 ms      0 ms
 4 125.25.24.35    3 ms      3 ms      0 ms

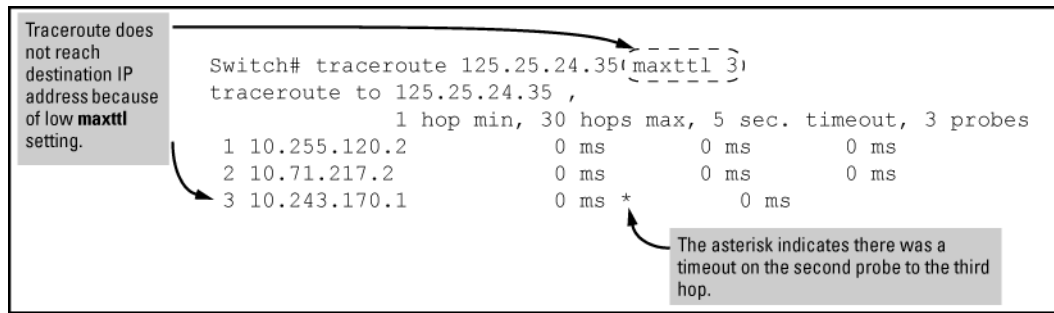
```

Intermediate router hops with the time taken for the switch to receive acknowledgement of each probe reaching each router.

Destination IP Address

Continuing from the previous example (Figure 253 (page 444)), executing traceroute with an insufficient `maxttl` for the actual hop count produces an output similar to this:

Figure 254 Incomplete traceroute because of low `maxttl` setting



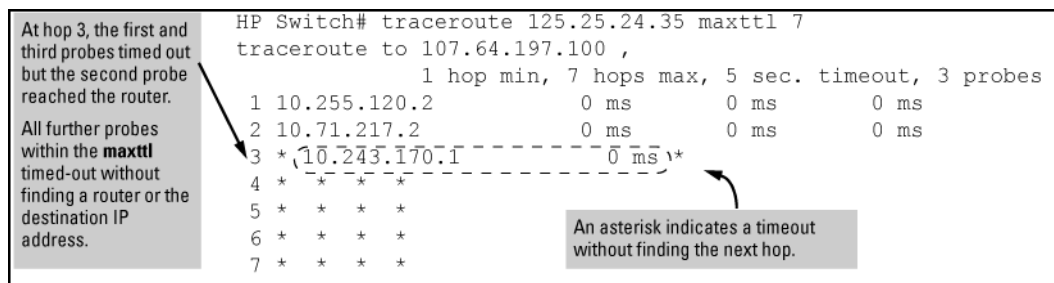
If a network condition prevents traceroute from reaching the destination

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 255 Traceroute failing to reach the destination address



Viewing Switch Configuration and Operation

In some troubleshooting scenarios, you may need to view the switch configuration to diagnose a problem. The complete switch configuration is contained in a file that you can browse from the CLI using the commands described in this section.

Viewing the startup or running configuration file

Syntax

`write terminal`

Displays the running configuration.

<code>show config</code>	Displays the startup configuration.
<code>show running-config</code>	Displays the running-config file.

For more information and examples of how to use these commands, see Chapter 6, "Switch Memory and Configuration".

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

Figure 256 (page 447) shows sample output from the `show tech` command.

Figure 256 show tech command

```
HP Switch# show tech

show system

Status and Counters - General System Information

System Name       : 5400_1
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 448).

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 `TransferCapture Text...`(see [Figure 257](#) (page 447).)

Figure 257 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 258](#) (page 448).

Figure 258 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 the following, press the Space bar to display and copy more information.

```
-- MORE --
```

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 "<i>command</i> "</pre>	Includes the output of a specified command in <code>show-tech</code> command output. Enter the command name between double-quotation marks, for example, <code>copy "show system" show-tech</code> .
<pre>crash-data [<i>slot-id</i> master] :</pre>	Includes the crash data from all management and interface modules in <code>show tech</code> command output.

	<p>To limit the amount of crash data displayed, specify an installed module or management modules, where:</p> <ul style="list-style-type: none"> • <code>slot-id</code>: Includes the crash data from an installed module. Valid slot IDs are the letters a through h. • <code>master</code>: Includes the crash data from both management modules.
<pre>crash-log [slot-id master] :</pre>	<p>Includes the crash logs from all management and interface modules in <code>show tech</code> command output.</p> <p>To limit the amount of crash-log data displayed, specify an installed module or management modules, where:</p> <p><code>slot-id</code>: Includes the crash log from an installed module. Valid slot IDs are the letters a through h.</p> <p><code>master</code>: Includes the crash log from both management modules.</p>
<pre>event-log</pre>	<p>Copies the contents of the Event Log to <code>show tech</code> command output.</p>
<pre>running-config</pre>	<p>Includes the contents of the running configuration file in <code>show tech</code> command output</p>
<pre>startup-config</pre>	<p>Includes the contents of the startup configuration file in <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><code>ip-addr</code>: Specifies the IP address of the remote host device.</p> <p><code>remote-file</code>: Specifies the pathname on the remote host for the configuration file whose contents you want to include in the command output.</p> <p><code>pcunix</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 "File Transfers" appendix.</p>
<pre>usb config startup-config filename command-file acl-filename.txt</pre>	<p>Copies the contents of a configuration file or ACL command file from a USB flash drive to <code>show tech</code> command output, where:</p> <p><code>startup-config filename</code> : Specifies the name of a startup configuration file on the USB drive.</p> <p><code>command-file acl-filename.txt</code> : Specifies the name of an ACL command file on the USB drive.</p> <p>For more information on using <code>copy usb</code> commands, see the "File Transfers" appendix.</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>pcunix</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 "File Transfers" appendix.</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 [see "Displaying Saved Crash Information" in the "Redundancy (Switch 8212zl)" chapter].

```
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 453).)

```
show system-information
```

Displays globally configured parameters and information on switch operation (see "CLI: Viewing and Configuring System Information" in the "Interface Access and System Information" chapter.)

```
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.

<code>include</code>	Only the lines that contain the matching pattern are displayed in the output.
<code>exclude</code>	Only the lines that contain the matching pattern are <i>not</i> displayed in the output.
<code>begin</code>	The display of the output begins with the line that contains the matching pattern.

NOTE: Pattern matching is case-sensitive.

Examples

Below are examples of what portions of the running config file display depending on the option chosen.

Figure 259 Pattern matching with `include` option

```
HP Switch(config)# show run | include ipv6
  ipv6 enable
  ipv6 enable
ipv6 access-list "EH-01"
HP Switch(config)#
```

Displays only lines that contain "ipv6".

Figure 260 Pattern matching with `exclude` option

```
HP Switch(config)# show run | exclude ipv6

Running configuration:

; J8697A Configuration Editor; Created on release #K.15.01

hostname "HP Switch Switch 5406z1"
module 1 type J8702A
module 2 type J8705A
snmp-server community "notpublic" Unrestricted
vlan 1
    name "DEFAULT_VLAN"
    untagged A1-A24,B1-B24
    ip address dhcp-bootp
    no untagged B21-B24
    exit
vlan 20
    name "VLAN20"
    untagged B21-B24
    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

HP Switch(config)#
```

Displays all lines that don't contain "ipv6".

Figure 261 Pattern matching with `begin` option

```
HP Switch(config)# show run | begin ipv6

ipv6 enable
no untagged B21-B24
exit
vlan 20
    name "VLAN20"
    untagged B21-B24
    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

HP Switch(config)#
```

Displays the running config beginning at the first line that contains "ipv6".

Figure 262 (page 453) 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.

Figure 262 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   B1
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. For more information on other CLI practices, see chapter 4, "Using the Command Line Interface (CLI)."

Syntax

`alias`

Creates a shortcut alias name for commonly used commands and command options.

For more information, see "Using a Command Alias" in the "Using the Command Line Interface (CLI)" chapter.

Syntax

`kill`

Terminates a currently running, remote troubleshooting session. Use the `show ip ssh` command to list the current management sessions.

For more information, see "Denying Interface Access by Terminating Remote Management Sessions" in the "Interface Access and System Information" chapter.

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.

For more information, see "Repeating a Command" in the "Using the Command Line Interface (CLI)" chapter.

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 *Access Security Guide*.

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.

Zeroization

The zeroization feature is used to restore the switch to its default, factory state. This is accomplished by sanitizing the persistent storage media (compact, SD, flash, etc..) Zeroization removes all customer installed/alterd files, including the following:

- switch configurations (including manager/operator credentials)
- system generated public keys
- system generated private keys

- user installed private keys and certificates
- legacy manager/operator password files
- crypto-key files
- fdr logs
- core dumps

Zeroization can be initiated in one of three ways:

- CLI command
- During Secure Mode transition, initiated through the `secure-mode` CLI command executed in a serial session
- ROM console command

The zeroization process can be time-consuming; for this reason, it is performed during the initial process of a switch reboot. After zeroization, the configuration file is rebuilt from the default config file, which is similar to the config rebuilding process performed by the `erase startup-config` command.

Syntax

```
erase all [zeroize]
```

Erases all management module files, including configuration files, core dumps, password files, crypto-key files, etc. Software images are not erased.

When executed without the `zeroize` option, files are removed, but the flash storage is not zeroized. The data is still physically present in the flash. The flash can be removed from the switch and the data recovered with file recovery tools.

[zeroize]: Zeroizes the file storage of the management modules.

NOTE: It is recommended that zeroization be performed from the serial console so that the status information can be viewed during the zeroization process.

Example 77 Zeroizing the management module files

```
HP Switch(config)# erase all zeroize
```

The system will be rebooted and all management module files except software images will be erased and zeroized. This will take up to 60 minutes and the switch will not be usable during that time. Continue (y/n)? y

Zeroizing from the ROM console

It is also possible to zeroize the file storage from the ROM console of the switch, using the `erase-all zeroize` command at the prompt. This most likely occurs during a switch recovery process. The warning messages are the same as for the CLI command.

Example 78 Zeroizing the management module files from the ROM console

```
=> erase-all zeroize
```

The system will be rebooted and all management module files except software images will be erased and zeroized. This will take up to 60 minutes and the switch will not be usable during that time. Continue (y/n)? y

Zeroizing with HA

When zeroization is triggered by a secure mode transition, HA handles zeroization on the AMM and SMM automatically.

When the CLI command (`erase all zeroize`) is used to start zeroization, the AMM syncs with the SMM and ensures that the SMM performs the same level of zeroization before the AMM starts the zeroization process on itself. The AMM before the zeroization process occurs remains the AMM, unless it takes over a minute for the AMM to boot up, in which case the prior SMM becomes the AMM.

When zeroization is started from the ROM console, there is no synchronization performed between the AMM and SMM, as zeroization from the ROM console is treated as a recovery facility. Each MM has to be zeroized individually.

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:

<ul style="list-style-type: none">• Baud rate: 9600• No parity• 8 Bits	<ul style="list-style-type: none">• 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 `CallDisconnect`.
- ii. Select `FileProperties`.
- 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 `CallConnect`.
 - 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 `TransferFile` 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 263 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. (At software release K.13.01, the DNS-compatible commands include `ping` and `traceroute`.)

Beginning with software release K.13.01, DNS operation supports both IPv4 and IPv6 DNS resolution and multiple, prioritized DNS servers. (For information on IPv6 DNS resolution, see *IPv6 Configuration Guide*.)

Basic operation

- When the switch is configured with only the IP address of a DNS server available to the switch, a DNS-compatible command, executed with a fully qualified domain name, can reach a device found in any domain accessible through the configured DNS server.
- When the switch is configured with both of the following:
 - The IP address of a DNS server available to the switch
 - The domain suffix of a domain available to the configured DNS server
 then:
 - A DNS-compatible command that includes the host name of a device in the same domain as the configured domain suffix can reach that device.
 - A DNS-compatible command that includes a fully qualified domain name can reach a device in any domain that is available to the configured DNS server.

Example

Suppose the switch is configured with the domain suffix `mygroup.HP Switch.net` and the IP address for an accessible DNS server. If an operator wants to use the switch to ping a target host in this domain by using the DNS name "leader" (assigned by a DNS server to an IP address used in that domain), the operator can use either of the following commands:

Figure 264 Example of using either a host name or a fully qualified domain name

```
HP Switch# ping leader
10.28.229.220 is alive, time = 1 ms

HP Switch# ping leader.mygroup.HP Switch.net
10.28.229.220 is alive, time = 1 ms
```

In the preceding example, if the DNS server's IP address is configured on the switch, but a domain suffix is either not configured or is configured for a different domain than the target host, the fully qualified domain name *must* be used.

Note that if the target host is in a domain *other than* the domain configured on the switch:

- The host's domain must be reachable from the switch. This requires that the DNS server for the switch must be able to communicate with the DNS servers in the path to the domain in which the target host operates.
- The fully qualified domain name must be used, and the domain suffix must correspond to the domain in which the target host operates, regardless of the domain suffix configured in the switch.

Example

Suppose the switch is configured with the domain suffix `mygroup.HP Switch.net` and the IP address for an accessible DNS server in this same domain. This time, the operator wants to use the switch to trace the route to a host named "remote-01" in a different domain named `common.group.net`. Assuming this second domain is accessible to the DNS server already configured on the switch, a `traceroute` command using the target's fully qualified DNS name should succeed.

Figure 265 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
```

Configuring and using DNS resolution with DNS-compatible commands

(At software release K.13.01, 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” \(page 457\)](#).) 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 264 \(page 458\)](#).) 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.

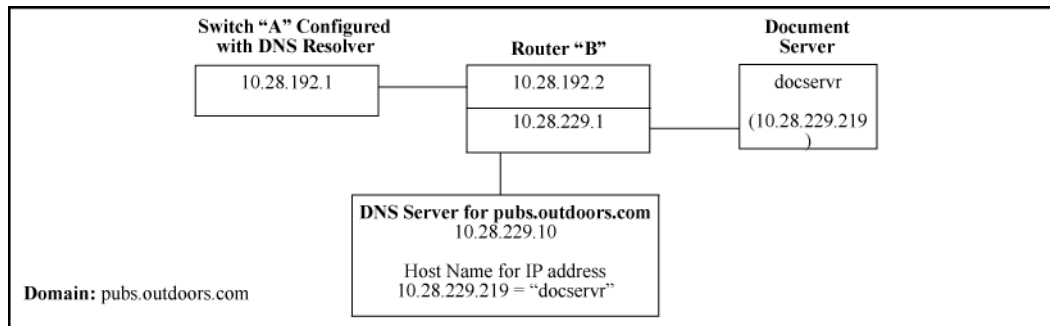
The switch supports one domain suffix entry and three DNS server IP address entries. (See the preceding command description.)

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 266 \(page 460\)](#), 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 266 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.

Figure 267 Configuring switch "A" in Figure 266 (page 460) to support DNS resolution

```
HP Switch(config)# ip dns server-address 10.28.229.10
HP Switch(config)# ip dns domain-name pubs.outdoors.com
```

Figure 268 Example of ping and traceroute execution for the network in Figure 266 (page 460)

```
HP Switch# 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 ms      0 ms      0 ms
 2 10.28.229.219   0 ms      0 ms      0 ms
```

As mentioned under “Basic operation” (page 457), 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 266 (page 460) as a target:

Figure 269 Example of ping and traceroute execution when only the DNS server IP address is configured

```
HP Switch# ping [docservr.pubs.outdoors.com]
10.28.229.219 is alive, time = 1 ms

HP Switch# traceroute [docservr.pubs.outdoors.com]
traceroute to 10.28.229.219
          1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.192.2      1 ms      0 ms      0 ms
 2 10.28.229.219   0 ms      0 ms      0 ms
```

Viewing the current DNS configuration

The `show ip` command displays the current domain suffix and the IP address of the highest priority DNS server configured on the switch, along with other IP configuration information. If the switch configuration currently includes a non-default (non-null) DNS entry, it will also appear in the `show run` command output.

Figure 270 Example of viewing the current DNS configuration

```
HP Switch# show ip

Internet (IP) Service

IP Routing : Disabled

Default Gateway : 10.28.192.2
Default TTL     : 64
Arp Age        : 20
[Domain Suffix  : pubs.outdoors.com]
[DNS server     : 10.28.229.10]

VLAN          | IP Config | IP Address   | Subnet Mask
-----+-----
DEFAULT_VLAN | Manual    | 10.28.192.1 | 255.255.255.0
```

Operating notes

- Configuring another IP address for a priority that has already been assigned to an IP address is not allowed. To replace one IP address at a given priority level with another address having the same priority, you must first use the `no` form of the command to remove the unwanted address. Also, only one instance of a given server address is allowed in the server list. Attempting to enter a duplicate of an existing entry at a different priority level is not allowed.

To change the priority of an existing server address, use the `no` form of the command to remove the entry, then re-enter the address with the new priority.

- To change the position of an address already configured with priority `x`, you must first use `no ip dns server-address priority x ip-addr` to remove the address from the configuration, then use `ip dns server-address priority ip-addr` to reconfigure the address with the new priority. Also, if the priority to which you want to move an address is already used in the configuration for another address, you must first use the `no` form of the command to remove the current address from the target priority.
- The DNS servers and domain configured on the switch must be accessible to the switch, but it is not necessary for any intermediate devices between the switch and the DNS server to be configured to support DNS operation.
- When multiple DNS servers are configured on the switch, they can reside in the same domain or different domains.
- A DNS configuration must include the IP address for a DNS server that is able to resolve host names for the desired domain. If a DNS server has limited knowledge of other domains, its ability to resolve DNS-compatible command requests is also limited.
- If the DNS configuration includes a DNS server IP address but does not also include a domain suffix, then any DNS-compatible commands should include the target host's fully qualified domain name.
- Switch-Initiated DNS packets go out through the VLAN having the best route to the DNS server, even if a Management VLAN has been configured.
- The DNS server address must be manually input. It is not automatically determined via DHCP.

Event Log messages

Please see the *Event Log Message Reference Guide* for information about Event Log messages.

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

Figure 271 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 464).

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 VLAN1 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: MAC address <i>mac-addr</i> not found.
[<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.

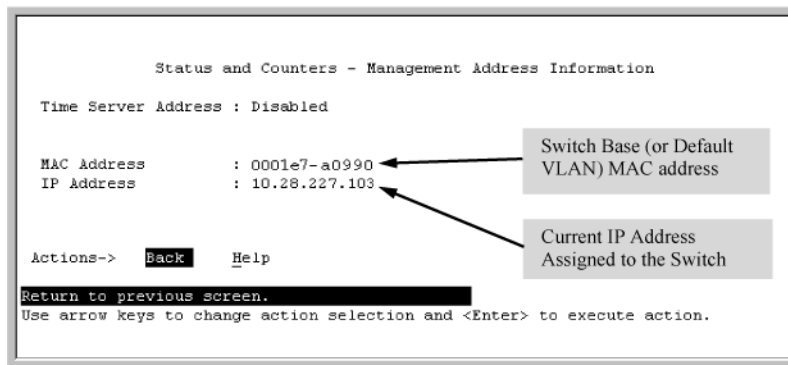
Also, the Base MAC address appears on a label on the back of the switch.

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 272 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.

Table 33 Switch series' and their MAC address allocations

Switch series	MAC address allocation
8212zl	The switch allots 24 MAC addresses per slot. For a given slot, if a four-port module is installed, the switch uses the first four MAC addresses in the allotment for that slot, and the remaining 18 MAC addresses are unused. If a 24-port module is installed, the switch uses the first 24 MAC addresses in the allotment, and so on.
All Models	The switch's base MAC address is assigned to VLAN (VID) 1 and appears in the <code>walkmib</code> listing after the MAC addresses for the ports. (All VLANs in the switch have the same MAC address.)

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 273](#) (page 465):

- 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 273 Example of Port MAC address assignments on a switch

HP Switch# walkmib ifphysaddress	
ifPhysAddress.1 = 00 12 79 88 b1 ff	ifPhysAddress.1 - 4: Ports A1 - A4 in Slot A (Addresses 5 - 24 in slot A are unused.)
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.49 - 72:Ports C1 - C24 in Slot C (In this example, there is no module in slot B.)
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.362 Base MAC Address (MAC Address for default VLAN; VID = 1)
ifPhysAddress.461 = 00 12 79 88 a1 00	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
ifPhysAddress.488 = 00 12 79 88 a1 00	
ifPhysAddress.4456 =	Virtual LANs (VLANs)" chapter of the <i>Advanced Traffic Management Guide</i> for your switch.

E Monitoring Resources

Displaying current resource usage

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.

show resources	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.
qos access-list openflow policy	Display the same command output and provide different ways to access task-specific information. NOTE: See <i>OpenFlow Administrators Guide</i> .

Example

Figure A-1 shows the resource usage on a 3500yl 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.
- The switch is also configured for VT and is either blocking or throttling routed traffic with a high rate-of-connection requests.
- Varying ICMP rate-limiting configurations on ports 1 to 24, on ports 25 to 48, and on slot A, have resulted in different meter usage and different rule usage listed under QoS. Global QoS settings would otherwise result in identical resource consumption on each port range in the switch.
- There is authenticated client usage of IDM resources on ports 25 to 48.

Figure 274 Displaying current QoS resource usage on a series 3500yl switch

```

HP Switch# show qos resources

Resource usage in Policy Enforcement Engine

      |      Rules      |      Rules Used
Ports | Available | ACL | QoS | IDM | VT | Mirror | Other |
-----+-----+-----+-----+-----+-----+-----+-----+
1-24 |      3014 |  15 |  11 |   0 |   1 |   0 |   3 |
25-48 |      3005 |  15 |  10 |  10 |   1 |   0 |   3 |
A     |      3017 |  15 |   8 |   0 |   1 |   0 |   3 |

      |      Meters      |      Meters Used
Ports | Available | ACL | QoS | IDM | VT | Mirror | Other |
-----+-----+-----+-----+-----+-----+-----+
1-24 |      250 |   |   5 |   0 |   |   |   0 |
25-48 |      251 |   |   4 |   0 |   |   |   0 |
A     |      253 |   |   2 |   0 |   |   |   0 |

      | Application |
      | Port Ranges | Application Port Ranges Used
Ports | Available | ACL | QoS | IDM | VT | Mirror | Other |
-----+-----+-----+-----+-----+-----+-----+
1-24 |      3014 |  2 |  0 |  0 |   |   0 |  0 |
25-48 |      3005 |  2 |  0 |  0 |   |   0 |  0 |
A     |      3017 |  2 |  0 |  0 |   |   0 |  0 |

0 of 8 Policy Engine management resources used.
Key:
ACL = Access Control Lists
QoS = Device & Application Port Priority, QoS Policies, ICMP rate limits
IDM = Identity Driven Management
VT  = Virus Throttling blocks
Mirror = Mirror Policies, Remote Intelligent Mirror endpoints
Other = Management VLAN, DHCP Snooping, ARP Protection, Jumbo IP-MTU.

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 and VT 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.
 - A VT configuration (connection-rate filtering) on the switch does not affect switch resources unless traffic behavior has triggered either a throttling or blocking action on the traffic from one or more clients. When the throttling action ceases or a blocked client is unblocked, the resources used for that action are released.
- 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
 - VT applied to any port (when a high-connection-rate client is being throttled or blocked)

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 ingress and egress broadcast rate limiting through the CLI using `rate-limit bcast/mcast`
 - 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`
- For chassis products (for example, the 5400zl or 8212zl switches), 'slots' are listed instead of 'ports,' with resources shown for all installed modules on the chassis.
- 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 IDMapapplication)
- VT and blocking on individual clients.

NOTE: Virus throttling does not operate on IPv6 traffic.

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-base 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 Daylight Saving Time on HP Switches

This information applies to the following HP switches:

<ul style="list-style-type: none">• 212M• 224M• 1600M• 2400M• 2424M• 4000M• 8000M	<ul style="list-style-type: none">• Series 2500• Series 2510• Series 2600• Series 2610• Series 2800• Switch 2910• Series 3400cl• Series 3500• Series 3500yl• Series 4100gl• Series 4200vl	<ul style="list-style-type: none">• Series 5300xl• Series 5400zl• Switch 6108• Switch 6200yl• Series 6400cl• Switch 6600• Series 8200zl• ProCurve AdvanceStack Switches• ProCurve 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:

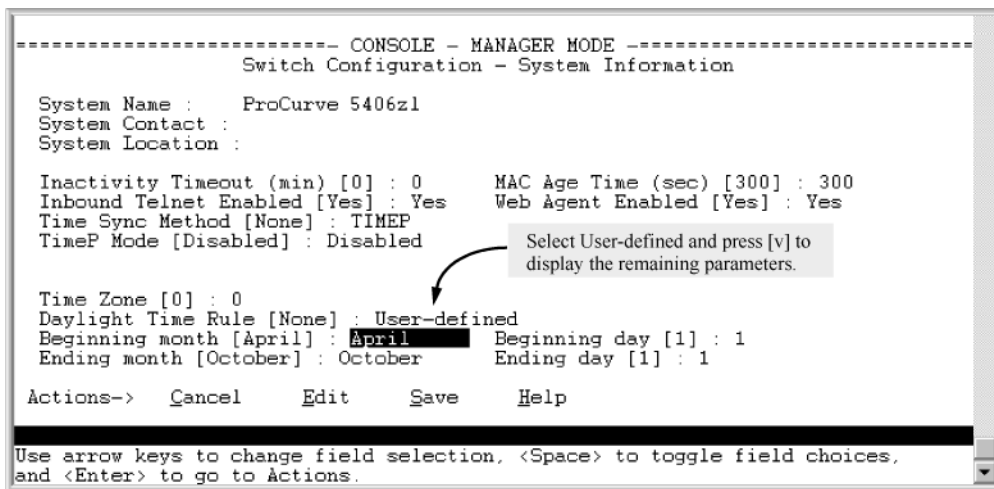
- Begin DST at 2 am the first Sunday on or after October 25th.
- End DST at 2 am the first Sunday on or after March 1st.

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 275 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.

G 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
IPv6 ACLs	
total IPv6 ACLs	Up to 2048 ¹
total ACEs in all IPv6 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
OSPFv2	
Interfaces/subnets	512 (128 active)
Max. areas supported	16
ECMP next hops	4
IPv6 Routing Protocol	
DHCPv6 Helper Addresses	32 unique addresses; multiple instances of same address counts as 1 towards maximum
OSPFv3	
Interfaces/subnets	512 (128 active)
Max. areas supported	16
ECMP next hops	4

¹ Actual availability depends on combined resource usage on the switch. See Appendix E, [“Monitoring Resources”](#) (page 466).

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.

H Switch Licensing

Switch software licensing enables advanced features in certain HP switches. The following table shows the software licenses available for the switches.

License type	Premium (includes OSPF, PIM – sparse mode, PIM – dense mode, VRRP, QinQ)
Switch family	License product
3500 and 3500yl	J8993A
5400zl	J8994A
6600	J9305A
8200zl	J9474A

General procedure

The general procedure for installing a software license involves several different numbers:

- **Registration ID** — This number comes with the license you purchase, and represents your right to install the particular type of license on a particular type of switch.
- **Hardware ID** — This number is provided by the switch that you are licensing, and includes the switch's serial number and an identifier for the feature that you are licensing.
- **License key** — This number is generated by the My HP Switch portal, based on the registration ID and the hardware ID that you provide. When you install this number into the switch, it enables the feature that you are licensing.

The procedure for installing a licensed feature into a switch is:

1. **Locate the registration ID.** When you purchase a software license, you receive a folded license registration card. The registration ID is located on the inside of the card, in the upper left corner.
2. **Get the switch's hardware ID.** Establish a console connection to the switch CLI and enter Manager level, using the `enable` command if necessary and the switch password if required. For example:

```
HP Switch enable
HP Switch#
```

From the Manager level, issue the `licenses hardware-id license_type` command. For example:

```
HP Switch# licenses hardware-id premium
```

The CLI returns a hardware ID number. Copy the hardware ID number from the screen (using Ctrl-C) or write it down. (Copying the number is easier and more accurate.) You will enter the number on the My HP Switch portal in the next step.

3. **Get the license key.** Point your Web browser at the My HP Switch portal (<http://my.HPSwitch.com>) and sign in. Click the My Licenses tab, enter the registration ID, and then enter the hardware ID. At the end of the procedure a license key is displayed. (It is also e-mailed to you.) Copy the license key from the screen (using Ctrl-C) or write it down.
4. **Enter the license key into the switch.** On the CLI console, save the configuration of the switch (**write memory.**) Then, from a Manager-level prompt, issue a `licenses install premium license-key` command. (The license key number is not case sensitive.) For example:

```
HP Switch# licenses install premium AA000GG000-A-0123ABC-ABCD123-0A2B3C4-0123ABC
```

5. Reboot the switch. For example:

```
HP Switch# boot
```

or:

```
HP Switch# reload
```

The licensed features should now be active on the switch.

I Power-Saving Features

Overview

There are several power-saving features that can be configured for the indicated switches and modules. The power-saving features include the ability to:

- Turn slot power on or off
- Turn LED power on or off
- Turn slot auto low power mode on or off
- Use LLDP for Energy Efficient Ethernet

The modules support the power-saving features as indicated in the table below.

Product number	Description	LED power on/off	Slot auto low power mode	Slot power on/off
J8702A	HP Switch zl 24 10/100/1000 PoE Module	Yes	Yes	Yes
J8705A	HP Switch zl 20 Gig-T + 4 mGBIC Module	Yes	Yes	Yes
J8706A	HP Switch zl 24-Port Mini-GBIC Module	Yes	No	Yes
J8707A	HP Switch zl 4-Port 10GbE X2 Module	Yes	No	Yes
J8708A	HP Switch zl 4-Port 10GbE CX4 Module	Yes	No	Yes
J9307A	HP Switch 24-Port 10/100/1000 PoE+ zl Module	Yes	Yes	Yes
J9308A	HP Switch 20-Port 10/100/1000 PoE+/4-Port MiniGBIC zl Module	Yes	Yes	Yes
J9309A	HP Switch 4-Port 10Gbe SFP+ zl Module	Yes	No	Yes
J9478A	HP Switch 24-Port 10/100 PoE+ zl Module	Yes	Yes	Yes

Configuring the power-saving options

Syntax

```
[ no ] savepower [module slot-list | all ] [led slot-list | all ] [port-low-pwr [ slot-list | all ]]
```

Configures power-saving features.

<code>module [slot-list all]</code>	Turns power-saving options on or off for modules.
<code>led [slot-list all]</code>	Turns power-saving options on or off for the LEDs for a module, list of modules or all modules.
<code>port-low-pwr [slot-list all]</code>	Enables or disables auto power down for slots.

Configuring the savepower module option

The `module` option provides the ability to turn the slot power on or off. If no module is specified, all slots are powered off. You can also specify `all` to turn off the power for all slots. If the command is preceded by `no`, all the slots are powered on, if off already. (See [Figure 276 \(page 477\)](#).)

Figure 276 Example of savepower module Command

```
HP Switch(config)# savepower module c
HP Switch(config)# show savepower module

Module Save Power Information

Slot | Status
----+-----
A   | Disabled
B   | Disabled
C   | Enabled
D   | Disabled
E   | Disabled
```

The `savepower module` command shuts down the specified modules in the order specified in the command. The ports on these modules no longer pass traffic. Any management traffic (SNMP, SSH, Telnet) that passes through these modules is interrupted. It can take up to two minutes to power down all the specified modules. Check the Event Log to see the current status of the module power down. This command applies to PoE/PoE+ modules as well as non-PoE/PoE+ modules.

Figure 277 Example of savepower module all Command

```
HP Switch(config)# savepower led timer 06/01/2009 12:01
duration 12:00 recur

HP Switch(config)# show savepower led

Led Save Power Information

Alarm Start Time      : 06/01/09 12:01:07
Alarm Duration (HH:MM) : 12:00
Recurrent Status      : Enabled

Led Save Power Information

Slot | Status
----+-----
A   | Disabled
B   | Disabled
C   | Disabled
D   | Disabled
E   | Disabled
```

You can verify the status of the `savepower` command by using the `show modules` command or by checking the log messages (for 8200zl and 5400zl switches.)

NOTE: If a `savepower module slot-list` or `savepower module all` command is immediately followed by a `no savepower module slot-list` or `no savepower module all` command, the first slot in the list is powered down and then brought up.

Configuring the savepower LED option

The `savepower LED` option provides the ability to turn off specified slot LEDs or all LEDs. You can also configure a timer for turning off the chassis LEDs or the specified slot LEDs. There is one system-wide timer; all the selected slots will have the chassis LEDs turned off for the same amount of time.

Syntax

```
[ no ] savepower led [ slot-list | all ] [ timer [ MM/DD [ / [YY]YY ]  
HH:MM | now | duration [HH:MM[recur] ] ]
```

If a `slot-list` is specified, the LEDs for that `slot-list` are turned off.

The `all` option can be specified for the `slot-list`. All the chassis LEDs are turned off.

When the timer option is configured, schedules a timer for turning off the chassis LEDs or configured slot LEDs. The LEDs are turned off for the configured time period and duration. (See [Figure 278 \(page 479\)](#).)

<code>MM/DD [/ [YY]YY] HH:MM</code>	Specifies the date and time to start the timer.
<code>now</code>	Instantaneously turns off the LEDs. The configured timer is canceled and all the configured modules go into power-saving mode immediately.
<code>duration [HH:]MM</code>	<p>The amount of time the LEDs remain turned off. (Optional)</p> <p>If the duration value is zero, when the timer starts, the LEDs are turned off indefinitely until the timer is canceled or the command is overridden with another command.</p> <p>Default: 0 (zero)</p>
<code>recur</code>	<p>(Optional) If specified, the LEDs are turned off on a daily basis at the configured time. The <code>recur</code> option is ignored if the duration is configured as zero.</p> <p>Default: disabled.</p>

A new command overrides the previous command, regardless of the current state. For example, if a timer is active and new command is given, the currently running timer is canceled and the new timer is scheduled.

The `no` form of the `savepower led` command cancels any scheduled or running timer and the LEDs are returned to their original state. The `no savepower led all` command turns on all the switch LEDs.

Figure 278 Example of setting a time and duration for savepower led command

```
HP Switch(config)# savepower port-low-pwr c
HP Switch(config)# show savepower port-low-pwr

Port Save Power Information

Slot | Status
----+-----
A   | Disabled
B   | Disabled
C   | Enabled
D   | Disabled
E   | Disabled
```

Configuring the savepower port-low-pwr option

The `port-low-pwr` option puts the slots into auto low power mode if they are not linked. If a particular slot is specified, only that slot goes into auto low-power mode. Specifying `savepower port-low-pwr all` puts all the slots into auto low power mode. (See [Figure 279 \(page 479\)](#).)

The ports in low-power mode periodically monitor to determine if the link has become active. If a LAN cable is connected to one of the ports, that port will come out of the low-power mode state after approximately 2 seconds (the monitor period) and enter into normal power mode. The remaining ports continue to be in low-power mode.

The `no` form of the command puts the specified slot into normal power mode. Entering `no savepower port-low-pwr all` puts all the slots into normal power mode.

Figure 279 Example of savepower port-low-power command for slot C

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

Module Save Power Information

Slot | Status
----+-----
A   | Disabled
B   | Disabled
C   | Enabled
D   | Disabled
E   | Disabled
```

`show savepower commands`

The settings for the `savepower` commands can be viewed using the appropriate `show` command.

`show savepower module`

Displays the settings for the `savepower module` command (see [Figure 280 \(page 480\)](#).)

Figure 280 Example of output for show savepower module command

```
HP Switch(config)# show savepower port-low-pwr

Port Save Power Information

Slot | Status
----+-----
A    | Enabled
B    | Enabled
C    | Enabled
D    | Enabled
E    | Enabled
```

`show savepower port-low-pwr`

Displays the status of the power-down feature for the slots (see [Figure 281 \(page 480\)](#).) For the stackable switches, the output shows if the feature is enabled or not enabled.

Figure 281 Example of output for show savepower port-low-pwr command

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

Led Save Power Information

Alarm Start Time      : 06/01/09 12:01:07
Alarm Duration (HH:MM) : 12:00
Recurrent Status      : Enabled

Led Save Power Information

Slot | Status
----+-----
A    | Enabled
B    | Enabled
C    | Enabled
D    | Enabled
E    | Enabled
```

`show savepower led`

Displays the configured status of the LED power-saving option (see [Figure 282 \(page 480\)](#).)

Figure 282 Example of output for show savepower led command

```
HP Switch(config)# int B5-B7 energy-efficient-ethernet
HP Switch(config)# show energy-efficient-ethernet

Port      | EEE Config | Current Status | txWake (µS)
-----+-----
B1        | Enabled    | Active          | 30
B2        | Enabled    | Inactive        | -
B3        | Disabled   | Inactive        | -
B4        | Enabled    | Unsupported     | -
B5        | Enabled    | Active          | 30
B6        | Enabled    | Active          | 30
B7        | Enabled    | Inactive        | -
```


Enabling energy efficient ethernet

Energy efficient ethernet (EEE) follows the 802.3az standard, which provides support for a system to operate in low-power idle mode during low-link use. This allows both sides of a link to disable or turn off a portion of the system's transmit/receive circuitry, saving power. When traffic is ready for transmission, the interface sends a "wake-up" message to the link partner to prepare to receive the traffic. The circuitry is returned to "normal" mode. Both sides of the link must be EEE-capable to support the power-saving idle mode.

Syntax

```
[ no ] int port-list energy-efficient-ethernet
```

Enables EEE for a given port or range of ports.

The no form of the command disables EEE for a port or range of ports.

(Default: Enabled)

Example

Figure 283 EEE enabled on ports B5 - B7

```
HP Switch(config)# lldp config B5 dot3TlvEnable eee_config
HP Switch(config)# show lldp config B5

LLDP Port Configuration Detail

Port : B5
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
*poe_config
*eee_config
```

The parameters are explained in the following table.

Parameter	Description
EEE Config	The EEE configuration status, read from the configuration database.
<ul style="list-style-type: none">• Enabled	EEE mode is enabled.
<ul style="list-style-type: none">• Disabled	EEE mode is disabled.
Current Status	Current EEE operational status.
<ul style="list-style-type: none">• Active	The port is advertised and auto-negotiated EEE with link partner (an EEE-capable partner.) EEE mode is enabled.
<ul style="list-style-type: none">• Inactive	Set to one of the following conditions: <ul style="list-style-type: none">• EEE configuration is disabled on the local port.• Local port advertises EEE capabilities with "EEE disabled" link partner or non-EEE link partner.• Auto-negotiation is mandatory for EEE to work. EEE configuration will not be applied if the port is in forced/manual (speed-duplex) mode. The current status will be 'inactive' for forced/manual mode port configuration.• EEE is not supported for 10Base-T. The current status will be 'inactive' if the link is operating in 10Base-T mode.
<ul style="list-style-type: none">• Unsupported	The local physical interface does not have EEE capability.
txWake	Current value of transmit wake-up time (in microseconds.)

NOTE: The interface modules do not support adjustment of both transmit and receive wake-up times. Therefore, txWake is constant.

LLDP support for EEE

Layer 2 (data link layer) EEE capability is a feature that allows fine-tuning for EEE that uses LLDP TLVs for the negotiation of physical link partners' wake up time values. An EEE-capable port notifies its link partner about the EEE capabilities supported. The ports then negotiate how to best optimize energy efficiency.

Syntax

```
[ no ] lldp config port-list dot3TlvEnable eee_config
```

Enables the advertisement of Layer 2 EEE TLVs for a given port or range of ports.

The `no` form of the command disables the advertisement of EEE TLVs.

(Default: Enabled)

Examples

Figure 284 Configuring Layer 2 TLVs on a port

```
HP Switch(config)# show lldp info local-device B5

LLDP Local Port Information Detail

Port      : B5
PortType  : local
PortID    : 5
PortDesc  : B5
Pvid      : 1

Energy Efficient Ethernet (EEE) Wake Times (microseconds)

Transmit      : 10
Receive       : 10
Echo Transmit : 10
Echo Receive  : 10
Fallback Receive : 10
```

To display the EEE TLV information for the local port, enter the `show lldp info local-device port-list` command, as shown in [Figure 285 \(page 483\)](#).

Figure 285 Output for LLDP information for a local port

```
HP Switch(config)# show lldp info remote-device B6

LLDP Remote Device Information Detail

Local Port   : B6
ChassisType  : mac-address
ChassisID    : 00 15 23 ff 2d 49
PortType     : Local
PortID       : 3
SysName      : HP Switch
System Desc  : HP Switch
PortDesc     : 3
Pvid         : 22
.
.
.
Energy Efficient Ethernet (EEE) Wake Times (microseconds)

Transmit      : 10
Receive       : 10
Echo Transmit : 10
Echo Receive  : 10
Fallback Receive : 10
```

To display the EEE TLV information for the link partner, enter the `show lldp info remote-device port-list` command, as shown in [Figure 33 \(page 86\)](#).

J Network Out-of-Band Management (OOBM) for the 6600 Switch

OOBM Configuration

OOBM configuration commands can be issued from the global configuration context (config) or from a specific OOBM configuration context (oobm.)

Entering the OOBM configuration context from the general configuration context

Syntax

```
oobm
```

Enters the OOBM context from the general configuration context.

Example

```
HP Switch (config)# oobm
HP Switch (oobm)#
```

Enabling and disabling OOBM

From the OOBM context:

Syntax

```
enable
disable
```

From the general configuration context:

Syntax

```
oobm enable
oobm disable
```

Enables or disables networked OOBM on the switch.

OOBM is not compatible with either a management VLAN or stacking. If you attempt to enable OOBM when a management VLAN is enabled or when stacking is enabled, the command will be rejected and you will receive an error message.

If an OOBM IP address exists and you disable OOBM, the OOBM IP address configuration is maintained. If you enable OOBM and there is a pre-existing OOBM IP address, it will be reinstated.

Network OOBM is enabled by default.

Examples

```
HP Switch (oobm)# enable
HP Switch (oobm)# disable
HP Switch (config)# oobm enable
HP Switch (config)# oobm disable
```

Enabling and disabling the OOBM port

The `OOBM interface` command enables or disables the OOBM interface (that is, the OOBM port, as opposed to the OOBM function.)

From the OOBM context:

Syntax

```
interface [ enable | disable ]
```

From the general configuration context:

Syntax

```
oobm interface [ enable | disable ]
```

Enables or disables the networked OOBM interface (port.)

Examples

```
HP Switch (oobm)# interface enable
HP Switch (config)# oobm interface disable
```

Setting the OOBM port speed

The OOBM port operates at 10 Mbps or 100 Mbps, half or full duplex. These can be set explicitly or they can be automatically negotiated using the `auto` setting.

From the OOBM context:

Syntax

```
interface speed-duplex [ 10-half | 10-full | 100-half |
100-full | auto ]
```

From the general configuration context:

Syntax

```
oobm interface speed-duplex [ 10-half | 10-full | 100-half |
100-full | auto ]
```

Enables or disables the networked OOBM interface (port.) Available settings are:

10-half	10 Mbps, half-duplex
10-full	10-Mbps, full-duplex
100-half	100-Mbps, half-duplex
100-full	100-Mbps, full-duplex
auto	auto negotiate for speed and duplex

Example

```
HP Switch (oobm)# interface speed-duplex auto
```

Configuring an OOBM IPv4 address

Configuring an IPv4 address for the OOBM interface is similar to VLAN IP address configuration, but it is accomplished within the OOBM context.

From the OOBM context:

Syntax

```
[ no ] ip address [ dhcp-bootp | ip-address/mask-length ]
```

From the general configuration context:

Syntax

```
[ no ] oobm ip address [ dhcp-bootp | ip-address/mask-length ]
```

Configures an IPv4 address for the switch's OOBM interface.

You can configure an IPv4 address even when global OOBM is disabled; that address will become effective when OOBM is enabled.

Example

```
HP Switch (oobm)# ip address 10.1.1.17/24
```

Configuring an OOBM IPv4 default gateway

Configuring an IPv4 default gateway for the OOBM interface is similar to VLAN default gateway configuration, but it is accomplished within the OOBM context.

From the OOBM context:

Syntax

```
[ no ] ip default-gateway ip-address
```

From the general configuration context:

Syntax

```
[ no ] oobm ip default-gateway ip-address
```

Configures an IPv4 default gateway for the switch's OOBM interface.

Example

```
HP Switch (oobm)# ip default-gateway 10.1.1.1
```

OOBM show commands

The `show` commands for OOBM are similar to the analogous commands for the data plane. Note that you must always include the `oobm` parameter to see the information for the OOBM interface, regardless of the context. For instance, even from the OOBM context, the `show ip` command displays the IP configuration for the data plane; to see the IP configuration of the OOBM interface, you need to use `show oobm ip`.

Showing the global OOBM and OOBM port configuration

Syntax

```
show oobm
```

Summarizes OOBM configuration information. This command displays the global OOBM configuration (enabled or disabled), the OOBM interface status (up or down), and the port status (enabled/disabled, duplex, and speed.)

You can issue this command from any context

Example

```
HP Switch# show oobm

Global Configuration
OOBM Enabled      : Yes
OOBM Port Type    : 10/100TX
```

```
OoBM Interface Status : Up
OoBM Port      : Enabled
OoBM Port Speed  : Auto
```

Showing OoBM IP configuration

Syntax

```
show oobm ip
```

Summarizes the IP configuration of the OoBM interface. This command displays the status of IPv4 (enabled/disabled), the IPv4 default gateway, and the IPv4 address configured for the interface.

You can issue this command from any context.

Example

```
HP Switch# show oobm ip
```

Showing OoBM ARP information

Syntax

```
show oobm arp
```

Summarizes the ARP table entries for the OoBM interface.

You can issue this command from any context.

Example

```
HP Switch# show oobm arp
```

Application server commands

Application servers (as described in OoBM and server applications in [“Concepts” \(page 490\)](#)) have added a `listen` keyword with `oobm|data|both` options to specify which interfaces are active.

Default value is `both` for all servers.

Telnet:

```
telnet-server [listen oobm | data | both ]
Management and Configuration Guide
```

SSH:

```
ip ssh [listen oobm | data | both ]
Access Security Guide
```

SNMP:

```
snmp-server [listen oobm | data | both ]
Management and Configuration Guide
```

TFTP:

```
tftp server [listen oobm | data | both ]
Management and Configuration Guide
```

HTTP:

```
web-management [listen oobm | data | both ]
Management and Configuration Guide
```

In all cases, `show running-config` displays the server configurations.

Use the `no` form of the command to prevent the server from running on either interface.

Examples

Telnet: no telnet-server

SSH: no ip ssh ...

SNMP: no snmp-server ...

TFTP: no tftp server

HTTP: no web-management ...

The show servers command shows the listen mode of the servers:

```
HP Switch# show servers
Server listen mode
```

```
Server      Listen mode
-----
Telnet      | both
Ssh         | both
Tftp        | both
Web-management | both
Snmp        | both
```

Application client commands

CLI commands for client applications have added the `oobm` keyword to allow you to specify that the outgoing request be issued from the OOBM interface. If you do not specify the `oobm` keyword, the request will be issued from the appropriate in-band data interface. Command syntax is:

Telnet:

```
telnet ip-address [oobm]
Management and Configuration Guide
```

TFTP:

```
copy tftp ... ip-address filename... [oobm]
Management and Configuration Guide
```

SNTP:

```
[ no ] sntp server priority priority ip-address [oobm] [version]
Management and Configuration Guide
```

TIMEP:

```
[ no ] ip timep [ dhcp | manual ip-address | [oobm] ]
Management and Configuration Guide
```

RADIUS:

```
[ no ] radius-server host ip-address [oobm]
Access Security Guide
```

TACACS+:

```
[ no ] tacacs-server host ip-address [oobm]
Access Security Guide
```

DNS:

```
[ no ] ip dns server-address priority priority ip-address [oobm]
Management and Configuration Guide
```

Syslog:

```
[ no ] logging ip-address [[control-descr] | [oobm] ]
Management and Configuration Guide
```

Ping:

```
ping [ ... ] [source [ ip-address | vlan-id | oobm ] ]
```


Traceroute:


```
traceroute [...] [source [ ip-address | vlan-id | oobm ]]
```

Example

Figure 286 (page 489) shows setup and use of network OOBM using the commands described above.

Assume that the figure below describes how you want to set up your data center.

Figure 286 Example data center



Assume that you are configuring the switch in the left-hand rack to communicate on both the data and management networks. You might do the following:

- Configure an IP address on the data network.
- Verify that out-of-band management is enabled. (It is enabled by default.)
- Configure an IP address on the management network.
- Verify that the switch can communicate on both networks.

The CLI commands that follow would accomplish those tasks. (The first time through the process you might easily make the omission shown near the end of the example.)

```
Switch 41# config
Switch 41(config)# vlan 1
Switch 41(vlan-1)# ip address 10.1.129.7/20          Set up IP address on data network.
Switch 41(vlan-1)# end                             Exit back to manager context.
Switch 41# show oobm                               Look at default OOBM configuration.
```

```
Global Configuration
OOBM Enabled      : Yes
OOBM Port Type   : 10/100TX
OOBM Interface Status : Up
OOBM Port       : Enabled
OOBM Port Speed  : Auto
Defaults look appropriate.
```

```
Switch 41# config
Switch 41(config)# oobm
Switch 41(oobm)# ip address 10.255.255.41/24      Go to OOBM context and
Switch 41(oobm)# ip default-gateway 10.255.255.1 add IP address and
Switch 41(oobm)# end                             default gateway.
Switch 41# ping 10.1.131.44                       Exit back to manager context.
10.1.131.44 is alive, time = 19 ms                Ping server in this rack (on data network.)
Switch 41# ping 10.1.131.51                       Ping server in adjacent rack.
```

```

10.1.131.51 is alive, time = 15 ms
Switch 41# ping 10.255.255.42
The destination address is unreachable.
Switch 41# ping source oobm 10.255.255.42
10.255.255.42 is alive, time = 2 ms
Switch 41#

```

```

Ping switch in adjacent rack.
Oops! It's on the management network.
Go through the management port
and it works fine.

```

Concepts


Management communications with a managed switch can be:

- In band—through the networked data ports of the switch
- Out of band—through a dedicated management port (or ports) separate from the data ports

Out-of-band ports have typically been serial console ports using DB-9 or specially wired 8-pin modular (RJ-style) connectors. Some recent HP switches have added networked OOBM ports.

Figure 287 (page 490) shows management connections for a typical switch.

Figure 287 Management ports



OOBM operates on a "management plane" that is separate from the "data plane" used by data traffic on the switch and by in-band management traffic. That separation means that OOBM can continue to function even during periods of traffic congestion, equipment malfunction, or attacks on the network. In addition, it can provide improved switch security: a properly configured switch can limit management access to the management port only, preventing malicious attempts to gain access via the data ports.

Network OOBM typically occurs on a management network that connects multiple switches. It has the added advantage that it can be done from a central location and does not require an individual physical cable from the management station to each switch's console port.

Of the switches covered by this manual, network OOBM is available on:

- HP Switch 6600-24XG switch (J9265A)
- HP Switch 6600-48G switch (J9451A)
- HP Switch 6600-48G-4XG switch (J9452A)

Table 34 (page 491) summarizes the switch management ports.

Table 34 Switch management ports

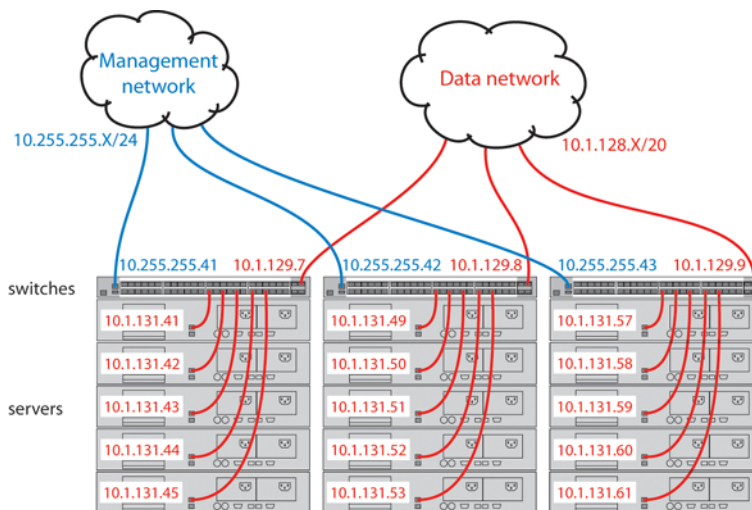
	In band		Out of band	
	Networked	Directly connected	Directly connected	Networked
Management interface	Command line (CLI), menu, Web	Command line (CLI), menu	Command line (CLI), menu	Command line (CLI), menu
Communication plane	Data plane	Management plane	Management plane	Management plane
Connection port	Any data port	Dedicated serial or USB console port	Dedicated serial or USB console port	Dedicated networked management port
Connector type	Usually RJ-45; also CX4, SFP, SFP+, and XFP	DB9 serial, serial-wired 8-pin RJ	DB9 serial, serial-wired 8-pin RJ	RJ-45
Advantages	Allows centralized management	Not affected by events on data network, shows boot sequence	Not affected by events on data network, shows boot sequence	Not affected by events on data network, allows centralized management, allows improved security
Disadvantages	Can be affected by events on data network; does not show boot sequence	Requires direct connection to console port (can be done via networked terminal server)	Requires direct connection to console port (can be done via networked terminal server)	Does not show boot sequence

Example

In a typical data center installation, top-of-rack switches connect servers to the data network, while the management ports of those switches connect to a physically and logically separate management network. This allows network administrators to manage the switches even if operation on the data network is disrupted.

In [Figure 288 \(page 491\)](#), the switches face the hot aisle of the data center, allowing easy connection to the network ports on the backs of the servers.

Figure 288 Network OOBM in a data center



For even more control, the serial console ports of the switches can be connected to the management network through a serial console server (essentially, a networked serial switch), allowing the network administrators to view the CLI activity of each switch at boot time and to control the switches through the console ports (as well as through the management ports.)

OOBM and switch applications

The table below shows the switch applications that are supported on the OOBM interface as well as on the data interfaces. In this list, some applications are client-only, some are server-only, and some are both.

Application	Inbound OOBM (server)	Outbound OOBM (client)	Inbound data plane (server)	Outbound data plane (client)
Telnet	yes	yes	yes	yes
SSH	yes	1	yes	1
SNMP	yes	yes ²	yes	yes
TFTP	yes	yes	yes	yes
HTTP	yes	1	yes	1
SNTP	1	yes	1	yes
TIMEP	1	yes	1	yes
RADIUS	1	yes	1	yes
TACACS	1	yes	1	yes
DNS ³	1	yes	1	yes
Syslog	1	yes	1	yes
Ping	yes ⁴	yes	yes ⁴	yes
Traceroute	yes ⁴	yes	yes ⁴	yes

¹ N/A = not applicable

² *=SNMP client refers to SNMP traps as they originate from the switch.

³ **=DNS has a limit of two servers—primary and secondary. Either can be configured to use the OOBM interface.

⁴ ***=Ping and Traceroute do not have explicit servers. Ping and Traceroute responses are sent by the host stack.

For applications that have servers, `oobm/data/both` options have been added to listen mode. There is now a `listen` keyword in the CLI commands to allow selection of those options. Default **value** is `both` for all servers. For details of the new command syntax, see [Support and Other Resources \(page 493\)](#).

K 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 Switch 620 Redundant and External Power Supply Installation and Getting Started Guide*
- *HP Switch 630 Redundant and/or External Power Supply Installation and Getting Started 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/support>.

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 35 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 8212zl#
```

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 289 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*, available at www.hp.com/support:

- 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.

Product warranties

For information about HP Networking product warranties, see the warranty information website:

<http://www.hp.com/networking/support>

Table 36 (page 496) lists related products and their part numbers.

Table 36 Applicable Products

Product	Part Number
HP Switch E3500yl-24G-PoE+ Switch	J9310A
HP Switch E3500yl-48G-PoE+ Switch	J9311A
HP Switch 10GbE 2-Port SFP+/2-Port CX4 yl Module	J9312A
630 Redundant and/or External Power Supply	J9443A
E3500-24 Switch	J9470A
E3500-48 Switch	J9472A
E3500-24-PoE Switch	J9471A
E3500-48-PoE Switch	J9473A
E3500yl-24G-PWR Intelligent Edge	J8692A
E3500yl-48G-PWR Intelligent Edge	J8693A
Switch E6200yl-24G mGBIC Premium Edge	J8992A
Switch E3500yl 2p 10GbE X2 + 2p CX4 Module	J8694A
620 Redundant and External Power Supply	J8696A
Switch E3500yl/E6200yl Fan Tray	5069-8598
Switch E3500yl/E6200yl Rack Mounting Kit	5069-5705
Switch E3500yl/E6200yl 10K Rack Rail Kit	356578-B21
Switch zl and yl RPS/EPS Cable	5070-0102

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/support.

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

Index

Symbols

- 802.1X
 - effect, LLDP, 234
 - LLDP blocked, 225
- 802.1X access control
 - authentication failure, SNMP notification, 217
 - SNMP notification of authentication failure, 217
- =prompt, 456

A

- access
 - manager, 213
 - operator, 213
- ACL
 - debug messages, 415
 - dynamic port ACL, 468
 - gateway fails, 383
 - mirroring
 - converted to classifier-based policies, 356
 - replaced by classifier-based match criteria, 356, 367
 - resources, 466
 - transferring command files, 310
 - troubleshooting, 382
- ACL, IPv4
 - limit, 472
 - RADIUS-assigned, limit, 472
 - scalability, 472
- ACL, IPv6
 - limit, 472
 - RADIUS-assigned, limit, 472
 - scalability, 472
- active management module
 - how determined, 274
- address
 - network manager, 211
- address table, port, 323
- address, network manager, 211
- advertise location, 228
- AES encryption, 164
- ARP
 - maximums, 472
- ARP protection
 - SNMP notification, 214, 217
- authentication
 - notification messages, 214, 217
 - SNTP, 31
 - SNTP client, 44
- authentication trap, 215
- authorized IP managers
 - SNMP, blocking, 210
- auto MDI/MDI-X
 - configuration, display, 61
 - operation, 60, 61
 - port mode, display, 61
- Auto-10, 111, 114, 119

- auto-TFTP, 283
 - disable, 283, 285, 305
 - disabled, 283
 - download to a redundant management system, 282
 - downloading software images, 283
- autonegotiate, 228

B

- bandwidth
 - displaying port utilization, 54
 - guaranteed minimum, 157
- blue locator LED, 462
- boot
 - with OSPF, 271
- boot active, 263
- boot ROM console, 303
- boot-history command, 255
- Bootp
 - effect of no reply, 381
- Bootp/DHCP, LLDP, 192
- bootset-default, redundancy, 247
- broadcast
 - limit, 73, 142
- broadcast mode
 - SNTP, 31
- broadcast storm, 110, 388
- broadcast storm loop, network port trunk, 99
- broadcast traffic
 - IPX, 73, 143
 - RIP, 73, 143

C

- CDP, 206, 235, 236
- chassislocate
 - LED, 462
- Classifier
 - ACL match criteria deprecated, 356, 367
 - benefits, 356, 367
 - class configuration
 - mirroring, 368
 - IPv6 match criteria, 367
 - mirroring configuration, 356, 357, 363, 367
 - default class, 369
 - examples, 371
 - restrictions, 367, 369
 - viewing, 343
 - policy configuration
 - mirroring, 368
 - resource usage, displaying, 350
 - restrictions
 - mirroring configuration, 367, 369
- clear
 - statistics
 - global, 321
 - ports, 321
- CLI

- context level, 57
- command line interface, 57
- Command Syntax
 - alias, 453
 - auto-tftp, 283
 - boot, 282
 - set-default, 251
 - system, 288
 - system flash, 286, 293
 - chassislocate, 462
 - class, 335
 - clear logging, 411
 - copy
 - command-output, 297
 - config, 293
 - core dump, 300
 - crash, 254
 - crash log, 254, 299
 - crash-data, 298
 - event-log, 297
 - fdr-log, 301
 - flash trtp, 290
 - flash usb, 290
 - flash xmodem, 290
 - running-config, 294
 - source, 448
 - startup-config, 291, 293, 294
 - tftp, 281, 288, 291, 294
 - tftp config, 291
 - trtp, 289, 292
 - usb, 296
 - usb flash, 287
 - usb startup-config, 294
 - xmodem, 285, 293, 296
 - xmodem config, 293
 - enable, 484
 - erase, 454, 455
 - fault-finder, 393
 - interface, 331, 343, 485
 - ip
 - default-gateway, 494
 - kill, 453
 - link, 442
 - logging, 436
 - logging filter, 209
 - mac-notify traps, 216
 - mirror, 331, 335, 337
 - mirror endpoint, 334, 336
 - mirror endpoint ip, 331
 - monitor all, 331, 332
 - monitor ip, 332
 - monitor mac, 332
 - no
 - autorun, 302
 - cdp, 204, 206, 207
 - class, 342
 - debug, 421, 425, 426
 - default-class, 343
 - ignore-untagged-mac, 209
 - int, 481
 - interface, 335, 339
 - ip address, 485
 - ip default-gateway, 486
 - ip dns, 459
 - lldp, 482
 - log-numbers, 412
 - logging, 428, 429, 430, 432, 436, 437
 - mac-notify traps, 215
 - mirror, 337
 - monitor mac, 331, 336, 341
 - oobm, 486
 - page, 453
 - policy mirror, 342
 - redundancy, 237
 - rmon alarm, 219
 - savepower, 476, 478
 - seq-number, 342
 - trtp, 282
 - oobm, 484
 - disable, 484
 - enable, 484
 - oobm interface, 485
 - ospf, 261
 - ospf3, 262
 - ping, 440
 - ping6, 440
 - policy mirror, 336
 - redundancy, 241, 242, 243
 - fabric-module, 255
 - reload, 252, 282, 286, 288, 293
 - repeat, 453
 - service-policy, 332
 - setup, 453
 - show
 - boot-history, 255, 450
 - cdp, 205
 - class, 348
 - clear statistics, 349
 - command option, 450
 - interfaces, 201, 394, 495
 - lldp, 200, 201, 202
 - logging, 253, 409
 - mac-address, 463
 - mac-notify traps, 216
 - modules, 246
 - monitor, 332, 344, 345, 346
 - oobm, 486, 487
 - policy, 349
 - policy resources, 350
 - qos, 466
 - redundancy, 237
 - tech, 292, 446
 - usb-port, 301
 - test, 399
 - traceroute, 443
 - usb-port, 301
 - vlan, 332, 340, 343
 - vrrp, 258

- write terminal, 445
- Command syntax
 - chassislocate, 315, 316
 - class, 330
 - clear
 - link-keepalive, 69
 - clear statistics, 321
 - int
 - poel, 83, 84
 - interface, 101
 - poel, 80
 - interfaces, 56, 61, 63
 - link-keepalive, 67, 68
 - lldp config, 85
 - module, 67
 - no
 - interface, 80
 - interfaces, 59, 63, 67, 68
 - lacp, 101, 102
 - mirror, 330
 - module, 67
 - modules, 319
 - poel, 80, 81, 83
 - port/trunk/mesh, 330
 - SNTP, 23, 25, 28
 - sntp authentication, 29, 31
 - sntp server priority, 30
 - switch-interconnect, 106
 - task-monitor, 316
 - timep, 36, 37, 38, 40, 47
 - timesync, 27, 36, 39, 47
 - trunk, 101
 - uplink-failure-detection, 71
 - poel, 83
 - policy mirror, 330
 - Show
 - chassislocate, 314
 - interfaces, 320
 - management, 317
 - modules, 318
 - show, 24
 - chassislocate, 316
 - config, 65, 66
 - interfaces, 48, 53, 61, 64, 65, 320
 - lacp, 100
 - link-keepalive, 69
 - lldp config, 85
 - mac-address, 322
 - management, 22
 - name, 64, 65
 - poel, 87, 88, 90
 - SNTP, 21
 - spanning-tree, 325
 - timep, 35
 - trunks, 99
 - trunks load-balance, 105
 - usb-port, 58
 - sntp, 23, 24, 27, 43
 - authentication, 29
 - broadcast, 31
 - server priority, 27
 - unicast, 31
 - switch-interconnect, 106
 - timep, 35, 36, 37, 38, 39, 47
 - timesync, 36, 37, 38, 43, 47
 - SNTP, 23, 24, 25
 - timep, 38
 - trunk, 100
 - trunk-load-balance, 104
 - usb-port, 57
- communities, SNMP
 - viewing and configuring with the menu, 168
- configuration, 129
 - clearing module, 67
 - copying, 310
 - impacts of software download on, 303
 - port, 48
 - port trunk group, 99
 - port, duplex, 56
 - port, speed, 56
 - restoring factory defaults, 454
 - SNMP, 211, 213
 - SNMP communities, 168
 - traffic mirroring, 352
 - transferring, 310
 - trap receivers, 215
 - usb autorun, 311
- configuration file
 - browsing for troubleshooting, 445
- configuration file, multiple
 - copy from a USB device, 294
 - copy to a USB device, 293
 - copy via tftp, 291
 - copy via Xmodem, 292
- connection-rate filtering
 - affect on switch resources, 468
- console
 - measuring network activity, 381
 - status and counters menu, 314
 - troubleshooting access problems, 380
- copy
 - command output, 297
 - config
 - oobm, 291
 - crash data, 298
 - crash log, 299
 - event log output, 297
 - event-log smm, 297
 - show tech, 448
 - tftp
 - show-tech, 292
 - tftp show-tech, 292
- core-dump, 272
- CPU utilization, 317
- cpu utilization data, 316
- crash log
 - oobm, 299
- crash-log

redundancy, 254
customizing, show command output, 53

D

date format, events, 401

debug

acl messages, 415
compared to event log, 413
destination, logging, 415
displaying debug configuration, 418
filtering messages, 425
forwarding IPv4 messages, 416
lldp messages, 417
message filtering, 425
overview, 413
packet messages, 416
sending event log messages, 414
standard event log messages, 415
using CLI session, 414
wireless services, 417

debug command

all, 415, 421
cdp, 415, 421
configuring debug/Syslog operation, 417
destination, 421
destinations, 414, 426
event log, 438
event log as default, 414
event types supported, 414
ip, 416
ip fib, 416
ip ospfv3, 416
ip pim, 416
ipv6 dhcpv6-client, 416
ipv6 dhcpv6-relay, 416
ipv6 forwarding, 416
ipv6 nd, 416
ipv6 ospf3, 416
ipv6 packet, 417
lldp, 424
operating notes, 438
OSPF messages, 422
pim, 422
rip, 423
security, 417
services, 417, 425
show debug, 418
snmp, 417, 425
support for "debug" severity on Syslog servers, 427, 438
syntax, 415, 421
using CLI session, 425
vrrp, 417, 425
wireless-services, 425

default settings, 21, 43, 57, 59, 61, 67, 88, 193, 282, 283, 303
ping, 440, 441
PoE, 95
redundancy, 242, 255

security, 311

traceroute, 444

default trunk type, 104

DES

encryption, 164

DHCP

address problems, 381

effect of no reply, 381

DHCP snooping

SNMP notification, 214, 217

DHCP/Bootp, LLDP, 192

DHCPv6

client, 416

debug messages, 416

diagnostics tools, 439

browsing the configuration file, 445

displaying switch operation, 446, 448

ping and link tests, 439

traceroute, 443

viewing switch operation, 445

distributed trunking, 106, 129

broadcast forwarding, 132

display info, 107

ISC port config, 106

max trunk and links, 130

multicast forwarding, 132

peer-keepalive, 107

peer-keepalive settings, 109

port config, 106

restrictions, 135

show switch interconnect, 109

spanning-tree, 131

unicast traffic forwarding, 131

DNS

configuration, 458, 460

configuration, viewing, 461

DNS-compatible commands, 457, 458

domain name, fully qualified, 458, 461

event log messages, 462

example, 460

IPv6 DNS resolution, 457

operating notes, 461

ping, 457, 458

resolver, 457

resolver operation, 457

secure management VLAN, 462

server address, DHCP not used, 462

server IP address, 458, 462

three entries supported, 459

traceroute, 457, 458

VLAN, best route selection, 462

documentation

latest versions, 280

release notes, 280

dot1TlvEnable, 193

download, 288

software, 266

software using TFTP, 281

switch-to-switch, 288

- TFTP, 284
- troubleshooting, 303
- Xmodem, 285
- duplex advertisements, 227
- duplex information, displaying, 233
- duplicate MAC address, 392
- Dyn1, 114

E

- edge ports, 138, 151
- EEE, 481
- Emergency Location Id Number, 231
- Energy Efficient Ethernet
 - low power support, 481
 - show TLV information, 483
 - with LLDP, 482
- erase all zeroize, 455
- event log, 280
 - compared to debug/Syslog operation, 413
 - debugging by severity level, 414, 415, 427
 - debugging by system module, 414, 415, 427
 - generated by system module, 401
 - how to read entries, 400
 - listing entries, 410
 - losing messages, 401
 - navigation, 410
 - not affected by debug configuration, 438
 - security levels, 170
 - sending event log messages as traps, 169
 - sending messages to Syslog server, 415
 - severity level, 401, 437
 - system module, 437
 - time format, 401
 - used for debugging, 414
 - used for troubleshooting, 400
- excessive frames, 162

F

- fabric modules
 - disabling, 255
 - enabling, 255
- facility
 - logging, 415
- factory default configuration
 - restoring, 454
- failover, locator LED, 462
- failover, management module, locator LED, 462
- failure, switch software download, 304
- fans, show status, 314
- fault-finder
 - transceiver link-flap, 392
 - transceiver sensitivities, 393
 - warn and disable, 393
 - Web interface, 393
- fault-tolerance, 111
- features, 276
- fiber optics, monitoring links, 76
- filter, source-port
 - jumbo VLANs, 161

- firmware version, 317
- flow control
 - constraints, 58, 73
 - effect on rate-limiting, 152, 156
 - global, 58
 - global requirement, 73
 - jumbo frames, 160
 - per-port, 58, 73
- flow sampling, 210
- friendly port names, 74

G

- gateway
 - routing fails, 383
- giant frames, 162
- guaranteed minimum bandwidth
 - apportioning unallocated bandwidth, 158
 - configuration, 144
 - displaying current configuration, 147
 - impacts of QoS queue configuration, 158
 - operation, 157
 - outbound queue priority, 158
 - starving queues, 158

H

- Help
 - for CLI, 497
 - for menu interface, 496
- hotswapping mgmt module, 245, 264
- HP
 - Auto-MDIX feature, 60

I

- ICMP
 - resources, 466
- ICMP rate-limiting
 - caution, 154
 - configuring, 140
 - current rate-limit configuration, 141
 - effect of flow control, 156
 - effect on port trunks, 156
 - event log messages, 141, 157
 - interface support, 156
 - monitoring/mirroring, 156
 - network application, 154
 - no meshing, 156
 - note on testing, 156
 - operating notes, 156
 - operation, 140, 153
 - optimum packet size, 156

IDM

- resources, 466, 469

- IEEE 802.1d, 388

- IEEE P802.1AB/D9, 225

IGMP

- host not receiving, 384
- not working, 384
- statistics, 326

- Inbound Telnet Enabled parameter, 380

- include-credentials, SNMP, 33, 45
- informs
 - sending to trap receiver, 169
 - SNMP, 217
- intelligent mirroring, 353
- IP
 - address maximums, 472
 - duplicate address, 381
 - duplicate address, DHCP network, 381
 - time server address, 22, 35
- IP address
 - for SNMP management, 210
- IP routing
 - debug messages, 415
- ip-option with ping, 441
- IPv4
 - static route, maximum, 472
- IPv6
 - debug dhcpv6 messages, 416
 - match criteria
 - classifier, 357, 363, 367
 - static route, maximum, 472
- IPX
 - broadcast traffic, 73, 143
 - network number, 318
- ISC
 - port config, 106
- J**
- jumbo frames
 - configuration, 148
 - excessive inbound, 160
 - flow control, 160
 - GVRP operation, 159
 - management VLAN, 160
 - maximum size, 150, 159
 - meshing, 159
 - mirroring, 375
 - MTU, 159, 375, 377
 - port adds and moves, 159
 - port speed, 159
 - security concerns, 161
 - switch mesh domain, 161
 - through non-jumbo ports, 161
 - traffic sources, 160
 - troubleshooting, 162
 - VLAN tag, 159, 376
 - voice VLAN, 160
- L**
- LACP
 - 802.1X not allowed, 122
 - active, 116
 - blocked ports, 123
 - default port operation, 121
 - described, 113, 119
 - Dyn1, 114
 - dynamic, 119
 - full-duplex required, 111, 119
 - IGMP, 123
 - mirroring static trunk, 358
 - no half-duplex, 124
 - operation not allowed, 385
 - overview of port mode settings, 111
 - passive, 116
 - removing port from active trunk, 101
 - restrictions, 122
 - standby link, 120
 - status, terms, 122
 - STP, 123
 - trunk limit, 119
 - VLANs, 123
 - with 802.1X, 122
 - with port security, 123
- Layer-3
 - scalability, 472
- licensing, 474
- limit, broadcast, 142
- line rate, 358
- link failures
 - detecting, 76
- link speed, port trunk, 110
- link test, 439
- link-change traps, 175, 214
- listening
 - snmp-server, 184, 218
- LLDP
 - 802.1D-compliant switch, 234
 - 802.1X blocking, 225
 - 802.1X effect, 234
 - advertisement content, 226
 - advertisement, mandatory data, 226
 - advertisement, optional data, 227
 - advertisements, delay interval, 190, 226
 - CDP neighbor data, 235
 - chassis ID, 226
 - chassis type, 226
 - clear statistics counters, 233
 - comparison with CDP data fields, 235
 - configuration options, 222
 - configuring optional data, 227
 - data options, 223
 - data read options, 224
 - debug messages, 415, 417
 - default configuration, 222
 - DHCP/Bootp operation, 225
 - display neighbor data, 201
 - enable/disable, global, 189
 - Energy Efficient Ethernet, 482
 - general operation, 221
 - global counters, 203
 - holdtime multiplier, 226
 - hub, packet-forwarding, 222
 - IEEE P802.1AB/D9, 225
 - inconsistent value, 190
 - information options, 223
 - invalid frames, 203
 - IP address advertisement, 225, 234

- IP address subelement, 227
- IP address, DHCP/Bootp, 192
- IP address, options, 192
- IP address, version advertised, 192
- mandatory TLVs, 234
- MIB, 222, 225
- neighbor data remaining, 234
- neighbor data, displaying, 201
- neighbor statistics, 203
- neighbor, maximum, 233
- operating rules, 225
- operation, 221
- outbound packet options, 223
- packet boundaries, 222
- packet dropped, 222
- packet time-to-live, 224
- packet-forwarding, 222
- per-port counters, 203
- port description, 227
- port ID, 227
- port speed, 227
- port trunks, 225
- port type, 227
- refresh interval, 189
- reinitialization delay, 190
- remote management address, 224
- remote manager address, 227
- reset counters, 233
- setmib, delay interval, 226
- setmib, reinit delay, 190
- show commands, 187, 188
- show outbound advertisement, 200
- SNMP notification, 223
- SNMP traps, 223
- spanning-tree blocking, 225
- standards compatibility, 225
- statistics, 233
- statistics, displaying, 202, 233
- system capabilities, 227
- system description, 227
- system name, 227
- Time-to-Live, 222
- time-to-live, 189, 222
- transmission frequency, 222
- transmission interval, change, 189
- transmit and receive, 222
- transmit/receive modes, 222
- trap notice interval, 226
- trap notification, 226
- trap receiver, data change notice, 191
- TTL, 222, 224
- VLAN, untagged, 234
- walkmib, 224
- with PoE, 97

lldp

- port vlan ID support, 227

LLDP-MED

- displaying speed, 233
- ELIN, 231
- enable or disable, 222
- endpoint support, 228
- fast start control, 195
- location data, 197, 231
- medTlvenable, 196
- Neighbors MIB, 201
- topology change notification, 195
- Voice over IP, 228

load balancing, 109, 126

logging

- facility, 415
- neighbor-adjacency, 415
- priority-desc, 415
- udp, 429

logging command, 421

- syntax, 414, 415, 427

logical port, 116

loop, network, 110

loopback address

- remote mirroring, 361

M

MAC address, 317

- displaying detected devices, 463
- duplicate, 388, 392
- per-slot or per-switch, 464
- port, 463
- traffic selection in mirroring, 357
- VLAN, 464
- walkmib, 464

Management Information Base, 211

management module

- resetting, 246

management module failover, locator LED, 462

management modules

- hotswapping, 245, 264

management port, 490

management VLAN, 210

- DNS, 462

manager access, 213

manager password

- SNMP notification, 214, 217

match criteria

- mirroring, classifier-based, 356

max frame size, jumbo, 150

maximums, 473

MDI/MDI-X

- configuration, display, 61
- operation, 60
- port mode, display, 61

media type, port trunk, 110

mesh

- jumbo frames, 161
- mirroring, 353

meshing

- with redundancy, 264

MIB

- HP proprietary, 210
- listing, 210

- standard, 210
- mini-GBICs, displaying info, 269
- mirroring
 - 802.1Q tag, 377
 - ACL criteria (deprecated), 359
 - ACL criteria (deprecated), 356, 357, 363, 365
 - ACLs converted to classifier-based policies in K.14.xx and later, 356, 365
 - ACLs replaced by classifier-based criteria, 356, 367
 - ARP request, 378
 - booting pre-K.12.xx OS, 359
 - caution
 - configure destination first, 334, 356, 363, 369
 - endpoint removal, 337
 - exit port connection, 329, 378
 - classifier, first release, 356
 - CLI option, 355, 358
 - configuration options, 357
 - configuration override, 360
 - configuration, endpoint switch, 362
 - configuration, Menu, 332
 - configuration, session identity, 361
 - configuration, source switch, 362
 - configuration, traffic selection, 363
 - destination
 - classifier-based configuration, 369
 - local session, 334, 356
 - remote session, 334, 356
 - display configuration, 344
 - distributed traffic, 358
 - dropped traffic, 358
 - dynamic LACP trunk not supported, 358
 - effect of STP state, 376
 - encapsulation, 359
 - endpoint configuration, 336
 - endpoint switch, 358, 361
 - examples
 - classifier-based, 371
 - local session, 354, 371
 - remote session, 354, 373
 - exit port
 - caution, 329
 - local session, 329
 - oversubscribe, 355, 358
 - remote session, 335
 - requirements, 358
 - VLAN subnet, 355, 361, 362, 374, 378
 - frame fragment, 359
 - frame truncation, not supported, 359, 375
 - in show running configuration, 351
 - intelligent mirroring, 353
 - intermediate switches, 358
 - intrusion detection system (IDS), 353
 - IPv4 encapsulation, 334, 356, 359, 363, 369, 375, 377
 - jumbo frames, 375
 - K.12.xx, earlier software, 359
 - legacy configuration, 359
 - local session
 - configuration steps, 329
 - maximum sessions
 - destination, 354
 - source, 354
 - maximum sources on a destination, 354
 - Menu interface limit, 358, 360
 - Menu interface, local-only, 358
 - Menu option, 355
 - migration to K.12.xx, 359
 - migration to release K.14.xx, 359
 - mirror command, 337
 - monitor, auto config session 1, 364
 - monitor, autoconfig session 1, 340, 377
 - monitored interface
 - mesh, 333, 357
 - of source traffic, 357
 - port, 333, 357
 - trunk, 333, 357
 - VLAN, 333, 357
 - MTU, 375, 377
 - no-tag-added, 364, 377
 - operating notes, 376
 - overload on destination, 358
 - oversized frames, 359
 - port screen, 353
 - rate, 358
 - remote session
 - configuration steps, 360
 - first release supported, 353
 - quick reference, 361
 - supported switches, 356
 - restrictions
 - classifier-based, 367, 369
 - local sessions, 354
 - remote sessions, 354
 - source switch, 363
 - session 1, legacy configuration, 359
 - session limits, 355
 - show commands, 344, 345, 351
 - simultaneous source/destination, 355
 - SNMP, 358, 360
 - SNMP for no-tag-added mirroring, 364
 - static VLAN, 357
 - terminology, 354
 - traffic overload, 358
 - traffic selection
 - classifier-based criteria, 356, 357, 363, 367, 368
 - direction-based criteria, 355, 357, 364
 - MAC-based criteria, 344, 356, 357, 363, 366
 - overview, 355, 357
 - troubleshooting, 378
 - UDP destination address, 344, 346
 - UDP port, 344, 345, 374
 - untagged mirror packets, 364
 - VLAN
 - maximum frame size, 376
 - no-tag-added, 339
 - Web interface, 358
 - Web limits, 360

- MM Reset button, [246](#)
- MM shutdown button, [245](#)
- module
 - clearing the config, [67](#)
 - CLI command, [67](#)
 - configuring when not inserted, [75](#)
 - pre-configuring, [75](#)
- modules
 - compatibility for v2 zl and zl, [319](#)
 - mini-GBIC information, [319](#)
 - mixed v2 zl and zl, [352](#)
 - show command, [269](#), [318](#)
 - show details, [318](#)
 - zl unsupported with nonstop switching, [276](#)
- monitoring
 - links between ports, [76](#)
 - locator LED, [462](#)
- monitoring, traffic, [353](#)
- multicast, [144](#)
- multiple forwarding database, [318](#), [322](#)
- multiple VLAN, [210](#)

N

- navigation, event log, [411](#)
- network management functions, [211](#), [214](#)
- network manager address, [211](#)
- network slow, [381](#)
- no-tag-added, [364](#), [377](#)
- nonstop switching, [256](#)
 - advantage, [256](#)
 - unsupported modules, [276](#)
- notifications
 - authentication messages, [214](#), [217](#)
 - configuring trap receivers, [215](#)
 - link-change traps, [214](#)
 - network security, [217](#)

O

- oobm
 - address config, [485](#)
 - client commands, [488](#)
 - copy config to remote host, [291](#)
 - copy crash-data, [299](#)
 - copy crash-log, [299](#), [300](#)
 - default gateway config, [486](#)
 - enable/disable, [484](#)
 - server commands, [487](#)
 - show arp, [487](#)
 - show commands, [486](#)
 - show config, [487](#)
 - SNTP, [25](#)
 - timep, [37](#)
- operating system, [303](#)
- operation not allowed, LACP, [385](#)
- operator access, [213](#)
- OS, [303](#)
 - version, [290](#)
- OSPF
 - boot, [271](#)

- debug command, [422](#)
- debug messages, [416](#)
- reload, [271](#)

P

- packet
 - debug messages, [416](#)
- password
 - disables usb autorun, [313](#)
 - SNMP notification, [217](#)
 - SNMP notification for invalid login, [214](#)
- pattern matching, show command output, [450](#)
- pbr
 - debug messages, [416](#)
- peer-keepalive
 - configuring links, [107](#)
 - display settings, [109](#)
- ping, [440](#), [457](#), [458](#)
 - ip-option, [441](#)
- ping test, [439](#)
- PoE
 - advertisement of power, [85](#)
 - advertisements, [231](#)
 - allocate-by, [93](#)
 - benefit of LLDP-MED, [228](#)
 - changing priority level, [95](#)
 - changing the threshold, [83](#), [95](#)
 - configuration options, [93](#)
 - configuring operation, [95](#)
 - data link layer, [84](#)
 - detection status, [89](#)
 - displaying power status, [87](#)
 - DLC, [98](#)
 - enable or disable operation, [80](#)
 - enabling, disabling redundancy, [82](#)
 - EPS, defined, [80](#)
 - fault, [82](#)
 - IEEE 802.3at std, [98](#)
 - LLDP detection, enabling or disabling, [83](#), [84](#)
 - lldp negotiation, [98](#)
 - manually configuring power levels, [82](#)
 - max module power, [93](#), [94](#)
 - MPS
 - absent cnt, [90](#)
 - needed power for PoE+, [94](#)
 - negotiation, [84](#)
 - other fault, [90](#)
 - over current cnt, [90](#)
 - overview of status, [88](#)
 - PD support, [93](#)
 - poe-lldp-detect command, [97](#)
 - power denied cnt, [90](#)
 - power negotiation, [84](#)
 - priority critical, [95](#)
 - priority high, [96](#)
 - priority low, [96](#)
 - priority, port, [93](#), [94](#)
 - RPS, defined, [80](#)
 - setting allocation, [82](#)

- short cnt, 91
- slot-id-range option, 97
- status, 230
- status on specific ports, 90
- terminology, 80
- threshold, power, 83
- TLVs, 85, 98
- usage, 93
- using LLDP, 97
- policy enforcement engine
 - described, 468
 - displaying resource usage, 468
- poll interval, 21
- port
 - address table, 323
 - blocked by UDLD, 76
 - broadcast limit, 142
 - configuration, 48
 - configuring UDLD, 76
 - context level, 57
 - counters, 353
 - counters, reset, 353
 - enabling UDLD, 68
 - fiber-optic, 72
 - MAC address, 463, 464
 - management, 490
 - menu access, 62
 - mirroring, static LACP trunk, 358
 - monitoring, 353
 - traffic patterns, 353
 - transceiver status, 55
 - trunk, 113
 - utilization, 54
 - CLI, 54
- port configuration, 99
- port mirroring, 353
- port names, friendly
 - configuring, 63
 - displaying, 64
 - summary, 74
- port security
 - port trunk restriction, 110
 - trunk restriction, 116
- port trunk, 109, 111
 - bandwidth capacity, 110
 - caution, 101, 102, 110
 - default trunk type, 104
 - enabling UDLD, 68
 - IGMP, 115
 - limit, 109
 - limit, combined, 119
 - link requirements, 110
 - logical port, 116
 - media requirements, 114
 - media type, 110
 - mirroring, 353
 - monitor port restrictions, 116
 - nonconsecutive ports, 109
 - port security restriction, 116
 - removing port from static trunk, 101
 - requirements, 114
 - spanning tree protocol, 115
 - static trunk, 114
 - static trunk, overview, 111
 - static/dynamic limit, 119
 - STP, 115
 - STP operation, 114
 - traffic distribution, 115
 - Trk1, 114
 - trunk (non-protocol) option, 113
 - trunk option described, 125
 - types, 113
 - UDLD configuration, 76
 - VLAN, 115
 - VLAN operation, 114
- port trunk group
 - interface access, 99
- port-based access control
 - event log, 385
 - LACP not allowed, 122
 - troubleshooting, 385
- port-utilization and status displays, 54
- power levels, configuring, 82
- power supply
 - show settings, 314
- power-over-ethernet, 80
- Premium License
 - installing, 474
- ProCurve
 - Auto-MDIX feature, 60
 - HP, URL, 211
 - switch documentation, 280
- ProCurve Manager
 - reading USB autorun files, 312
 - required for USB autorun, 311
 - security concerns when deleting public community, 211
 - SNMP and network management, 210
 - updating switch software, 310
- prompt, =, 456
- public SNMP community, 211, 214

Q

- QoS, 468
- Quality of Service
 - queue configuration, 158
 - resources, 466

R

- RADIUS-assigned ACLs
 - resources, 468
- rapid switchover stale timer
 - redundancy
 - rapid switchover stale timer, 257
- rapid-switchover, 242
- rate display for ports, 54
- rate-limiting, 144
 - bcast command, 143
 - broadcast traffic, 143

- caution, [138](#), [151](#)
- configuration, [138](#), [140](#)
- disabling multicast, [144](#)
- displaying configuration, [138](#), [141](#)
- edge ports, [138](#), [151](#)
- effect of flow control, [152](#), [156](#)
- effect on port trunks, [152](#), [156](#)
- how measured, [152](#)
- ICMP, [151](#)
- ICMP operation, [153](#)
- intended use, [138](#), [151](#)
- mcast command, [143](#)
- multicast traffic, [143](#)
- note on testing, [153](#), [156](#)
- operating notes, [152](#)
- optimum packet size, [153](#), [156](#)
- per-port only, [151](#)
- purpose, [151](#)
- traffic filters, [152](#)
- redundancy, [82](#)
 - active-management, [243](#)
 - boot command, [270](#)
 - boot-history, [255](#)
 - bootset-default, [247](#)
 - causes of switchover, [263](#)
 - core-dump, [272](#)
 - crash files, [254](#)
 - disabling, [268](#)
 - downloading software, [266](#)
 - enabling redundant management, [237](#)
 - hotswapping module, [245](#), [264](#)
 - how active module determined, [274](#)
 - IGMP note
 - IGMP, [264](#)
 - locator LED, [462](#)
 - management module interaction, [256](#)
 - meshing note, [264](#)
 - MM Reset button, [246](#)
 - MM shutdown button, [245](#)
 - module in failed mode, [244](#)
 - nonstop switching, [256](#)
 - nonstop switching commands, [237](#)
 - nonstop switching features
 - nonstop switching, [276](#)
 - other commands affected, [269](#)
 - rapid switchover, [242](#)
 - reload command, [252](#)
 - resetting mgmt module, [246](#)
 - setting active module, [243](#)
 - setting default flash for boot, [251](#)
 - show detail, [248](#)
 - show flash, [248](#)
 - show log, [249](#)
 - show module, [269](#)
 - show modules, [269](#)
 - show redundancy, [247](#)
 - show version, [249](#)
 - software version mismatch, [246](#), [265](#), [266](#)
 - switchover, [257](#)

- switchover operations, [263](#)
- syncing CLI commands, [275](#)
- transitioning to nonstop switching, [257](#)
- warm-standby, [256](#)
- reload, [252](#)
 - with OSPF, [271](#)
- remote intelligent mirroring, [353](#)
- remote mirroring
 - loopback address, [361](#)
- reset
 - port counters, [353](#)
- resetting the switch
 - factory default reset, [454](#)
- resource monitor
 - event log, [469](#)
- resource usage
 - displaying, [466](#)
 - insufficient resources, [469](#)
- restricted write access, [213](#)
- RFCs, [210](#)
 - RFC 1493, [210](#)
 - RFC 1515, [210](#)
 - RFC 2737, [225](#)
 - RFC 2863, [225](#)
 - RFC 2922, [225](#)
- RIP
 - broadcast traffic, [73](#), [143](#)
 - debug messages, [416](#)
- RMON, [210](#)
- rmon alarm, [219](#)
- RMON groups supported, [219](#)
- router
 - maximum routes, [472](#)
 - OSPF area maximum, [472](#)
 - OSPF interface maximum, [472](#)
 - RIP interface maximum, [472](#)
 - supported routes, [472](#)
- routing
 - gateway fails, [383](#)
 - traceroute, [443](#)
- running-config
 - history of changes, [176](#)
 - traps for changes, [218](#)

S

- savepower
 - led option, [478](#)
 - port-low-pwr, [479](#)
 - show led, [480](#)
 - show module, [479](#)
 - show port-low-pwr, [480](#)
- scalability, [472](#)
- SCP/SFTP
 - enabling, [285](#)
 - session limit, [307](#), [309](#)
 - troubleshooting, [308](#)
- secure copy, [285](#)
- secure FTP, [285](#)
- secure management VLAN, DNS, [462](#)

- security
 - enabling network security notifications, 217
 - USB autorun, 302
- Self Test LED
 - behavior during factory default reset, 454
- serial number, 317
- setmib
 - delay interval, 226
 - reinit delay, 190
- severity level
 - event log, 401
 - selecting Event Log messages for debugging, 437
- sFlow, 210
 - CLI-owned versus SNMP-owned configurations, 185
 - configuring via the CLI, 184
 - sampling-polling information, 187
 - show commands, 186
- show
 - custom option, 53
 - displaying specific output, 450
 - pattern matching with, 450
- show cpu, 316
- show debug, 418
- show flash, 248
- show interfaces
 - dynamic display, 53
- show interfaces display, 450
- show log, 249
- show management, 35
- show module, 269
- show modules command, 269, 318
- show policy resources command, 350
- show redundancy, 247
- show redundancy, detail, 248
- show tech, 292, 446
- show version, 249
- slow network, 381
- SNMP, 210
 - ARP protection events, 214
 - authentication notification, 214, 217
 - CLI commands, 213
 - communities, 211, 213
 - configuring with the menu, 168
 - mapping, 213
 - config changes traps, 218
 - configure, 211
 - configuring security groups, 171
 - configuring SNMPv3 notification, 171
 - configuring SNMPv3 users, 171
 - configuring trap receivers, 215
 - configuring trap receivers, 215
 - DHCP snooping events, 214
 - different versions, 214
 - enabling informs, 217
 - enabling SNMPv3, 171
 - enabling traps, 175
 - fixed traps, 215
 - invalid password in login, 214
 - IP, 210
 - link-change traps, 175, 214
 - manager password change, 214
 - mirroring, 358
 - network security notification, 217
 - no-tag-added mirroring, 364
 - notification, LLDP
 - SNMP notification, 223
 - public community, 211, 214
 - running config changes, 218
 - running-config changes, 218
 - supported notifications, 214
 - system thresholds, 215
 - traps, 76, 210, 214
 - traps for config changes, 218
 - walkmib, 464, 465
 - well-known traps, 215
- SNMP trap, LLDP, 226
- snmp-server
 - listening mode, 184, 218
- SNMPv3
 - "public" community access caution, 163
 - access, 211
 - assigning users to groups, 212
 - communities, 213
 - enable command, 163
 - enabling, 211
 - encryption, configuring, 164
 - group access levels, 212, 213
 - groups, 165
 - network management problems with snmpv3 only, 163
 - restricted-access option, 163
 - set up, 211
 - users, 211
- SNTP
 - authentication command, 31
 - authentication mode, 29
 - broadcast
 - mode, 31
 - broadcast mode, 23, 42
 - broadcast mode, requirement, 42
 - client authentication, 44
 - configuration, 42
 - disabling, 26
 - display config information, 32
 - display statistics, 33
 - event log messages, 46
 - include-credentials, 33, 45
 - key-id, 29, 30, 44
 - key-value, 29, 44
 - menu interface operation, 44
 - oobm, 25
 - operating modes, 42
 - poll interval, 27
 - priority, 27, 30
 - server priority, 27
 - show authentication, 32
 - trusted key, 29, 45
 - unicast mode, 25, 31, 42
 - unicast time polling, 43

- unicast, replacing servers, 44
- viewing, 21, 42
- software, 303
- software image, 303
- software licensing, 474
- software version, 317
- source port filters
 - jumbo VLANs, 161
- spanning tree
 - fast-uplink, troubleshooting, 388
 - mirroring blocked traffic, 376
 - problems related to, 388
 - show tech, copy output, 447
 - using with port trunking, 115
- SSH
 - file transfer, 283
 - TACACS exclusion, 307
 - troubleshooting, 308, 389
- standard MIB, 210
- static route, maximum, 472
- statistics
 - clearing, 321
 - SNTP, 33
- subnet
 - VLAN, mirroring exit port, 355, 361, 362, 374, 378
- switch software, 303
 - download using TFTP, 281
 - download, failure indication, 304
 - download, troubleshooting, 303
 - download, using TFTP, 281
 - installing a license, 474
 - software image, 303
 - version, 285, 287
- switchover
 - causes, 263
 - operations during, 263
 - when it doesn't occur, 264
- Syslog, 413
 - "debug" severity level as default, 438
 - adding priority description, 436
 - compared to event log, 413
 - config friendly descriptions, 436
 - configuring for debugging, 417
 - configuring server address, 414
 - configuring server IP address, 421
 - configuring Syslog servers and debug destinations, 414
 - control-desc, 436
 - displaying Syslog configuration, 418
 - logging command, 421, 426
 - operating notes, 438
 - overview, 413
 - priority-descr, 436
 - sending event log messages, 414
 - server configuration, 428
 - severity, "debug", 427
 - specifying severity level events for debugging, 437
 - specifying system module events for debugging, 437
 - user facility as default, 438
 - using event log for debugging, 415, 427

- system module
 - selecting event log messages for debugging, 437
- System Support Modules, information about, 269

T

- TACACS
 - SSH exclusion, 307
- task monitor, 316
- taskusage -d, 316
- taskUsageShow, 316
- Telnet
 - troubleshooting access, 380
- temperature, show settings, 314
- TFTP
 - auto-TFTP, 283
 - auto-TFTP feature, 283
 - auto-TFTP, disable, 283, 285, 305
 - copy command output, 297
 - copy crash data, 298
 - copy crash log, 299
 - copy event log output, 297
 - copying a configuration file, 291
 - copying software image, 290
 - disable, 285, 305
 - disabled, 283
 - download software using CLI, 281
 - downloading software using console, 284
 - enable client or server, 282
 - enabling client functionality, 282
 - enabling server functionality, 282
 - switch-to-switch transfer, 288
 - troubleshooting download failures, 303
 - uploading an ACL command file, 294
 - using to download switch software, 281
- threshold setting, 211, 214
- thresholds, SNMP, 215
- throttling, broadcast/multicast traffic, 143
- time format, events, 401
- time protocol
 - selecting, 21
- Time-to-Live
 - LLDP, 222
- TimeP
 - assignment methods, 46
 - disabling, 39
 - poll interval, 39
 - server address listing, 22, 35
 - show management, 35
 - viewing and configuring, menu, 35
 - viewing, CLI, 35
- timesync, disabling, 39
- TLV advertisement, 193
- TLVs, mandatory, 234
- traceroute, 457, 458
 - blocked route, 445
 - fails, 444
- traffic
 - broadcast rate-limiting, 143
 - multicast rate-limiting, 143

- port, 353
- traffic mirroring, 353
- traffic monitoring, 210, 211, 214, 353
- transceiver
 - error messages, 56
 - fault sensitivities, 393
 - fault-finder, 392
 - fiber-optic, 72
 - flapping, 392
 - link-flap, 392
 - view status, 55
- transceivers
 - configuring when not inserted, 75
 - not inserted, 75
- trap
 - CLI access, 215
 - configuring trap receivers, 215
 - security levels, 170
- trap notification, 226
- trap receiver, 211
 - configuring, 215
 - sending event log messages, 169
 - sending SNMPv2 informs, 169
 - SNMP, 215
 - up to ten supported, 215
- traps, 215
 - arp-protect, 173
 - auth-server-fail, 173
 - dhcp-snooping, 174
 - dynamic-ip-lockdown, 174
 - fixed, 215
 - for config changes, 218
 - link-change, 174, 175
 - login-failure-mgr, 174
 - password-change-mgr, 174
 - port-security, 174
 - running-config changes, 218
 - snmp-authentication, 174
 - threshold, 215
- troubleshooting
 - ACL, 382
 - approaches, 379
 - browsing the configuration file, 445
 - configuring debug destinations, 414
 - console access problems, 380
 - diagnosing unusual network activity, 381
 - diagnostics tools, 439
 - displaying switch operation, 446, 448
 - DNS, 457
 - fast-uplink, 388
 - ping and link tests, 439
 - resource usage, 467
 - restoring factory default configuration, 454
 - spanning tree, 388
 - SSH, 389
 - SSH, SFTP, and SCP Operations, 308
 - switch software download, 303
 - switch won't reboot, shows = prompt, 456
 - traceroute, 457, 458
 - unusual network activity, 381
 - using CLI session, 414
 - using debug and Syslog messaging, 413
 - using the event log, 400
 - viewing switch operation, 445
 - web browser access problems, 380
- trunk, 109
 - distributed, 106
 - L4 load balancing, 126
 - load balancing, 126
 - number supported, 110
- trunking
 - distributed trunking config, 129
 - ISC port config, 106
- TTL
 - LLDP, 222
- U
- UDLD
 - changing the keepalive interval, 68
 - changing the keepalive retries, 69
 - configuration, 77
 - configuring for tagged ports, 69
 - enabling on a port, 68
 - operation, 76
 - overview, 76
 - supported switches, 77
 - viewing configuration, 69
- UDP
 - logging messages, 429
- undersize frames, 162
- Uni-directional Link Detection, 76
- unicast mode
 - SNTP, 31
- unrestricted write access, 213
- unusual network activity, 381
- up time, 317
- URL
 - ProCurve, 211
- USB, 303
 - autorun, 311
 - AutoRun file, 311
 - command file, 311
 - configuring passwords, 313
 - creating a command file, 302
 - enabling or disabling, 313
 - LED indications, 312
 - report outputs, 312
 - required software versions, 311
 - secure-mode, 313
 - security, 302
 - troubleshooting, 312
 - auxiliary port, 309, 311
 - auxiliary port LEDs, 312
 - configuration, 57
 - copy command output, 297
 - copy configuration file to/from a USB device, 293
 - copy crash data, 298
 - copy crash log, 299

- copy event log output, [297](#)
- copy software image to a USB device, [290](#)
- devices with secure partitions not supported, [309](#)
- enable port, [57](#)
- flash drives must be formatted, [309](#)
- supported capabilities, [309](#)
- uploading an ACL command file, [296](#)
- viewing flash drive contents, [309](#), [310](#)
- usb-port, [57](#)
- users, SNMPv3, [212](#)
- utilization, port, [54](#)

V

- v2 zl modules, compatibility with zl modules, [319](#)
- version, OS, [290](#)
- version, switch software, [285](#), [287](#)
- view
 - transceiver status, [55](#)
- VLAN
 - address, [210](#)
 - configuring UDLD for tagged ports, [69](#)
 - device not seen, [391](#)
 - event log entries, [401](#)
 - IP address maximum, [472](#)
 - jumbo max frame size, [150](#)
 - link blocked, [388](#)
 - MAC address, [464](#)
 - management and jumbo frames, [160](#)
 - management VLAN, SNMP block, [210](#)
 - maximums, [472](#)
 - mirroring, [352](#), [353](#)
 - multiple, [210](#)
 - port configuration, [391](#)
 - secure management VLAN, with DNS, [462](#)
 - switch software download, [281](#)
 - tagging broadcast, multicast, and unicast traffic, [391](#)
- VoIP
 - LLDP-MED support, [228](#)
- VRRP
 - debug messages, [417](#)

W

- walkmib, [224](#), [464](#), [465](#)
- warm-standby
 - booting, [256](#)
- warm-standby redundant management, [256](#)
- web browser interface
 - troubleshooting access problems, [380](#)
- web site, HP, [211](#)
- wireless services
 - debug messages, [417](#)
- write access, [213](#)

X

- Xmodem
 - copy command output, [297](#)
 - copy crash data, [298](#)
 - copy crash log, [299](#)
 - copy event log output, [297](#)

- copying a configuration file, [292](#)
- copying a software image, [290](#)
- uploading an ACL command file, [296](#)
- using to download switch software, [285](#)

Z

- zeroize
 - file storage, [454](#)
 - from ROM console, [455](#)
 - management modules, [454](#)
 - redundant MM, [455](#)
 - with HA, [455](#)