

HP Switch Software

Management and Configuration Guide for WB.15.16

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.

Applicable Products

HP Switch 2920-series:

J9726A
J9727A
J9728A
J9729A



© Copyright 2015 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® and Windows® are trademarks of the Microsoft group of companies.

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

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

Warranty

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

Contents

| | |
|--|-----------|
| 1 Time Protocols | 17 |
| General steps for running a time protocol on the switch | 17 |
| TimeP time synchronization | 17 |
| SNTP time synchronization | 17 |
| Selecting a time synchronization protocol | 17 |
| Disabling time synchronization | 18 |
| SNTP: Selecting and configuring | 18 |
| Viewing and configuring SNTP (Menu) | 19 |
| Viewing and configuring SNTP (CLI) | 20 |
| Configuring (enabling or disabling) the SNTP mode | 21 |
| Enabling SNTP in Broadcast Mode | 22 |
| Enabling SNTP in unicast mode (CLI) | 23 |
| Changing the SNTP poll interval (CLI) | 25 |
| Changing the SNTP server priority (CLI) | 25 |
| Disabling time synchronization without changing the SNTP configuration (CLI) | 25 |
| Disabling the SNTP Mode | 26 |
| SNTP client authentication | 26 |
| Requirements | 27 |
| Configuring the key-identifier, authentication mode, and key-value (CLI) | 27 |
| Configuring a trusted key | 28 |
| Configuring a key-id as trusted (CLI) | 28 |
| Associating a key with an SNTP server (CLI) | 28 |
| Enabling SNTP client authentication | 29 |
| Configuring unicast and broadcast mode for authentication | 29 |
| Viewing SNTP authentication configuration information (CLI) | 30 |
| Viewing all SNTP authentication keys that have been configured on the switch (CLI) | 30 |
| Viewing statistical information for each SNTP server (CLI) | 30 |
| Saving configuration files and the include-credentials command | 31 |
| TimeP: Selecting and configuring | 32 |
| Viewing, enabling, and modifying the TimeP protocol (Menu) | 33 |
| Viewing the current TimeP configuration (CLI) | 34 |
| Configuring (enabling or disabling) the TimeP mode | 35 |
| Enabling TimeP in manual mode (CLI) | 35 |
| SNTP unicast time polling with multiple SNTP servers | 38 |
| Displaying all SNTP server addresses configured on the switch (CLI) | 38 |
| Adding and deleting SNTP server addresses | 39 |
| Adding addresses | 39 |
| Deleting addresses | 39 |
| Operating with multiple SNTP server addresses configured (Menu) | 39 |
| SNTP messages in the Event Log | 40 |
| Monitoring resources | 40 |
| Displaying current resource usage | 40 |
| Viewing information on resource usage | 41 |
| Policy enforcement engine | 42 |
| Usage notes for show resources output | 43 |
| When insufficient resources are available | 43 |
| 2 Port Status and Configuration | 44 |
| Viewing port status and configuring port parameters | 44 |
| Connecting transceivers to fixed-configuration devices | 44 |
| Viewing port configuration (Menu) | 45 |
| Configuring ports (Menu) | 46 |

| | |
|---|-----------|
| Viewing port status and configuration (CLI)..... | 47 |
| Dynamically updating the show interfaces command (CLI/Menu)..... | 47 |
| Customizing the show interfaces command (CLI)..... | 48 |
| Error messages associated with the show interfaces command..... | 49 |
| Note on using pattern matching with the show interfaces custom command..... | 49 |
| Viewing port utilization statistics (CLI)..... | 50 |
| Operating notes for viewing port utilization statistics..... | 50 |
| Viewing transceiver status (CLI)..... | 50 |
| Operating notes..... | 51 |
| Enabling or disabling ports and configuring port mode (CLI)..... | 52 |
| Enabling or disabling flow control (CLI)..... | 52 |
| Port shutdown with broadcast storm..... | 54 |
| Viewing broadcast storm..... | 55 |
| SNMP MIB..... | 56 |
| Configuring auto-MDIX..... | 58 |
| Manual override..... | 59 |
| Configuring auto-MDIX (CLI)..... | 59 |
| Using friendly (optional) port names..... | 60 |
| Configuring and operating rules for friendly port names..... | 60 |
| Configuring friendly port names (CLI)..... | 61 |
| Configuring a single port name (CLI)..... | 61 |
| Configuring the same name for multiple ports (CLI)..... | 61 |
| Displaying friendly port names with other port data (CLI)..... | 62 |
| Listing all ports or selected ports with their friendly port names (CLI)..... | 62 |
| Including friendly port names in per-port statistics listings (CLI)..... | 63 |
| Searching the configuration for ports with friendly port names (CLI)..... | 64 |
| Uni-directional link detection (UDLD)..... | 65 |
| Configuring UDLD..... | 66 |
| Configuring uni-directional link detection (UDLD) (CLI)..... | 66 |
| Enabling UDLD (CLI)..... | 67 |
| Changing the keepalive interval (CLI)..... | 67 |
| Changing the keepalive retries (CLI)..... | 67 |
| Configuring UDLD for tagged ports..... | 67 |
| Viewing UDLD information (CLI)..... | 68 |
| Viewing summary information on all UDLD-enabled ports (CLI)..... | 68 |
| Viewing detailed UDLD information for specific ports (CLI)..... | 68 |
| Clearing UDLD statistics (CLI)..... | 69 |
| 3 Power Over Ethernet (PoE/PoE+) Operation..... | 70 |
| Introduction to PoE..... | 70 |
| PoE terminology..... | 70 |
| PoE operation..... | 70 |
| Configuration options..... | 70 |
| PD support..... | 71 |
| Power priority operation..... | 71 |
| When is power allocation prioritized?..... | 71 |
| How is power allocation prioritized?..... | 71 |
| Configuring PoE operation..... | 72 |
| Disabling or re-enabling PoE port operation..... | 72 |
| Enabling support for pre-standard devices..... | 72 |
| Configuring the PoE port priority..... | 72 |
| Controlling PoE allocation..... | 73 |
| Manually configuring PoE power levels..... | 74 |
| Configuring PoE redundancy..... | 75 |
| Changing the threshold for generating a power notice..... | 75 |

| | |
|--|------------|
| PoE/PoE+ allocation using LLDP information..... | 77 |
| LLDP with PoE..... | 77 |
| Enabling or disabling ports for allocating power using LLDP..... | 77 |
| Enabling PoE detection via LLDP TLV advertisement..... | 77 |
| LLDP with PoE+..... | 77 |
| Overview..... | 77 |
| PoE allocation..... | 78 |
| Viewing PoE when using LLDP information..... | 79 |
| Viewing LLDP port configuration..... | 79 |
| Operating note..... | 80 |
| Viewing the global PoE power status of the switch..... | 81 |
| Viewing PoE status on all ports..... | 82 |
| Viewing the PoE status on specific ports..... | 84 |
| Using the HP 2920 Switch with an external power supply..... | 85 |
| Overview..... | 85 |
| Supported PSUs..... | 86 |
| Using the XPS for additional PoE power..... | 86 |
| Determining the maximum available PoE power..... | 86 |
| Operating rules..... | 88 |
| Using redundant (N+1) power..... | 89 |
| Providing non-PoE redundant power..... | 89 |
| Configuring the HP 2920 PoE switches to use the XPS..... | 89 |
| Enabling and disabling power from the XPS..... | 89 |
| Configuring auto-recovery..... | 90 |
| Restoring the default external power supply settings..... | 91 |
| Distributing power to specified ports..... | 91 |
| Example: of the power-share option..... | 92 |
| Example: of adding a switch..... | 92 |
| Example: of using the force option..... | 92 |
| Reducing allocated external power..... | 93 |
| Example: configurations..... | 93 |
| Non-PoE configuration..... | 93 |
| PoE configuration for full PoE power to one XPS port..... | 94 |
| PoE configuration for multiple switches..... | 96 |
| Viewing power information..... | 98 |
| Examples for show external-power-supply..... | 100 |
| Examples for show power-over-ethernet commands..... | 102 |
| Example: for show running-config command..... | 104 |
| Planning and implementing a PoE configuration..... | 105 |
| Power requirements..... | 106 |
| Assigning PoE ports to VLANs..... | 106 |
| Applying security features to PoE configurations..... | 106 |
| Assigning priority policies to PoE traffic..... | 106 |
| PoE Event Log messages..... | 106 |
| 4 Port Trunking..... | 107 |
| Overview of port trunking..... | 107 |
| Port connections and configuration..... | 107 |
| Port trunk features and operation..... | 108 |
| Fault tolerance | 108 |
| Trunk configuration methods..... | 108 |
| Dynamic LACP trunk..... | 108 |
| Using keys to control dynamic LACP trunk configuration..... | 109 |
| Static trunk..... | 109 |
| Viewing and configuring a static trunk group (Menu)..... | 112 |

| | |
|--|------------|
| Viewing and configuring port trunk groups (CLI)..... | 114 |
| Viewing static trunk type and group for all ports or for selected ports..... | 114 |
| Viewing static LACP and dynamic LACP trunk data..... | 115 |
| Dynamic LACP Standby Links..... | 115 |
| Configuring a static trunk or static LACP trunk group..... | 116 |
| Removing ports from a static trunk group..... | 116 |
| Enabling a dynamic LACP trunk group..... | 117 |
| Removing ports from a dynamic LACP trunk group..... | 117 |
| Viewing existing port trunk groups (WebAgent)..... | 118 |
| Trunk group operation using LACP..... | 118 |
| Default port operation..... | 120 |
| LACP notes and restrictions..... | 120 |
| 802.1X (Port-based access control) configured on a port..... | 120 |
| Port security configured on a port..... | 121 |
| Changing trunking methods..... | 121 |
| Static LACP trunks..... | 121 |
| Dynamic LACP trunks..... | 121 |
| VLANs and dynamic LACP..... | 121 |
| Blocked ports with older devices..... | 122 |
| Spanning Tree and IGMP..... | 122 |
| Half-duplex, different port speeds, or both not allowed in LACP trunks..... | 122 |
| Dynamic/static LACP interoperation..... | 123 |
| Trunk group operation using the "trunk" option..... | 123 |
| How the switch lists trunk data..... | 123 |
| Outbound traffic distribution across trunked links..... | 123 |
| Trunk load balancing using port layers..... | 125 |
| Enabling trunk load balancing..... | 125 |
| 5 Port Traffic Controls..... | 127 |
| Rate-limiting..... | 127 |
| All traffic rate-limiting..... | 127 |
| Configuring rate-limiting..... | 127 |
| Displaying the current rate-limit configuration..... | 128 |
| Operating notes for rate-limiting..... | 129 |
| Guaranteed minimum bandwidth (GMB)..... | 131 |
| GMB operation..... | 131 |
| Impacts of QoS queue configuration on GMB operation..... | 132 |
| Configuring GMB for outbound traffic..... | 133 |
| Viewing the current GMB configuration..... | 135 |
| GMB operating notes..... | 136 |
| Impact of QoS queue configuration on GMB commands..... | 136 |
| Jumbo frames..... | 137 |
| Operating rules..... | 137 |
| Configuring jumbo frame operation..... | 137 |
| Overview..... | 137 |
| Viewing the current jumbo configuration..... | 137 |
| Enabling or disabling jumbo traffic on a VLAN..... | 139 |
| Configuring a maximum frame size..... | 139 |
| Configuring IP MTU..... | 139 |
| SNMP implementation..... | 140 |
| Jumbo maximum frame size..... | 140 |
| Jumbo IP MTU..... | 140 |
| Displaying the maximum frame size..... | 140 |
| Operating notes for maximum frame size..... | 140 |
| Operating notes for jumbo traffic-handling..... | 141 |

| | |
|---|------------|
| Troubleshooting..... | 142 |
| A VLAN is configured to allow jumbo frames, but one or more ports drops all inbound jumbo frames..... | 142 |
| A non-jumbo port is generating "Excessive undersize/giant frames" messages in the Event Log..... | 142 |
| 6 Configuring for Network Management Applications..... | 143 |
| Using SNMP tools to manage the switch..... | 143 |
| SNMP management features..... | 143 |
| SNMPv1 and v2c access to the switch..... | 144 |
| SNMPv3 access to the switch..... | 144 |
| Enabling and disabling switch for access from SNMPv3 agents..... | 144 |
| Enabling or disabling restrictions to access from only SNMPv3 agents..... | 145 |
| Enabling or disabling restrictions from all non-SNMPv3 agents to read-only access..... | 145 |
| Viewing the operating status of SNMPv3..... | 145 |
| Viewing status of message reception of non-SNMPv3 messages..... | 145 |
| Viewing status of write messages of non-SNMPv3 messages..... | 145 |
| Enabling SNMPv3..... | 145 |
| SNMPv3 users..... | 146 |
| Adding users..... | 146 |
| SNMPv3 user commands..... | 147 |
| Listing Users..... | 147 |
| Assigning users to groups (CLI)..... | 148 |
| Group access levels..... | 148 |
| SNMPv3 communities..... | 149 |
| Mapping SNMPv3 communities (CLI)..... | 149 |
| SNMP community features..... | 150 |
| Viewing and configuring non-version-3 SNMP communities (Menu)..... | 151 |
| Listing community names and values (CLI)..... | 151 |
| Configuring community names and values (CLI)..... | 152 |
| SNMP notifications..... | 153 |
| Supported Notifications..... | 153 |
| General steps for configuring SNMP notifications..... | 154 |
| SNMPv1 and SNMPv2c Traps..... | 154 |
| SNMP trap receivers..... | 154 |
| Configuring an SNMP trap receiver (CLI)..... | 155 |
| SNMP trap when MAC address table changes..... | 156 |
| show command..... | 156 |
| SNMPv2c informs..... | 156 |
| Enabling SNMPv2c informs (CLI)..... | 157 |
| Configuring SNMPv3 notifications (CLI)..... | 157 |
| Network security notifications..... | 160 |
| Enabling or disabling notification/traps for network security failures and other security events (CLI)..... | 160 |
| Viewing the current configuration for network security notifications (CLI)..... | 161 |
| Enabling Link-Change Traps (CLI)..... | 162 |
| Readable interface names in traps..... | 162 |
| Source IP address for SNMP notifications..... | 162 |
| Configuring the source IP address for SNMP notifications (CLI)..... | 162 |
| Viewing SNMP notification configuration (CLI)..... | 164 |
| Configuring the MAC address count option..... | 165 |
| Displaying information about the mac-count-notify option..... | 166 |
| Advanced management: RMON..... | 167 |
| CLI-configured sFlow with multiple instances..... | 167 |
| Configuring sFlow (CLI)..... | 168 |

| | |
|---|-----|
| Configuring sFlow when destination is OOBM port..... | 168 |
| Viewing sFlow Configuration and Status (CLI)..... | 169 |
| Configuring UDLD Verify before forwarding..... | 171 |
| UDLD time delay..... | 171 |
| Restrictions..... | 171 |
| UDLD configuration commands..... | 172 |
| Show commands..... | 172 |
| RMON generated when user changes UDLD mode..... | 173 |
| LLDP..... | 173 |
| General LLDP operation..... | 173 |
| LLDP-MED..... | 174 |
| Packet boundaries in a network topology..... | 174 |
| LLDP operation configuration options..... | 174 |
| Enable or disable LLDP on the switch..... | 174 |
| Enable or disable LLDP-MED..... | 174 |
| Change the frequency of LLDP packet transmission to neighbor devices..... | 174 |
| Change the Time-To-Live for LLDP packets sent to neighbors..... | 174 |
| Transmit and receive mode..... | 174 |
| SNMP notification..... | 175 |
| Per-port (outbound) data options..... | 175 |
| Remote management address..... | 176 |
| Debug logging..... | 176 |
| Options for reading LLDP information collected by the switch..... | 176 |
| LLDP and LLDP-MED standards compatibility..... | 176 |
| LLDP operating rules..... | 177 |
| Port trunking..... | 177 |
| IP address advertisements..... | 177 |
| Spanning-tree blocking..... | 177 |
| 802.1X blocking..... | 177 |
| Configuring LLDP operation..... | 177 |
| Displaying the global LLDP, port admin, and SNMP notification status (CLI)..... | 177 |
| Viewing port configuration details (CLI)..... | 178 |
| Configuring Global LLDP Packet Controls..... | 179 |
| LLDP operation on the switch..... | 179 |
| Enabling or disabling LLDP operation on the switch (CLI)..... | 179 |
| Changing the packet transmission interval (CLI)..... | 179 |
| Time-to-Live for transmitted advertisements..... | 180 |
| Delay interval between advertisements generated by value or status changes to the LLDP MIB..... | 180 |
| Reinitialization delay interval..... | 181 |
| Configuring SNMP notification support..... | 182 |
| Enabling LLDP data change notification for SNMP trap receivers (CLI)..... | 182 |
| Changing the minimum interval for successive data change notifications for the same neighbor..... | 182 |
| Configuring per-port transmit and receive modes (CLI)..... | 182 |
| Basic LLDP per-port advertisement content..... | 183 |
| Mandatory Data..... | 183 |
| Configuring a remote management address for outbound LLDP advertisements (CLI)..... | 183 |
| Optional Data..... | 184 |
| Support for port speed and duplex advertisements..... | 184 |
| Configuring support for port speed and duplex advertisements (CLI)..... | 185 |
| Port VLAN ID TLV support on LLDP..... | 185 |
| Configuring the VLAN ID TLV..... | 185 |
| Viewing the TLVs advertised..... | 186 |
| SNMP support..... | 188 |

| | |
|---|-----|
| LLDP-MED (media-endpoint-discovery)..... | 188 |
| LLDP-MED endpoint support..... | 189 |
| LLDP-MED endpoint device classes..... | 190 |
| LLDP-MED operational support..... | 190 |
| LLDP-MED fast start control..... | 190 |
| Advertising device capability, network policy, PoE status and location data..... | 191 |
| Network policy advertisements..... | 191 |
| VLAN operating rules..... | 191 |
| Policy elements..... | 191 |
| Enabling or Disabling medTlvEnable..... | 192 |
| PoE advertisements..... | 193 |
| Location data for LLDP-MED devices..... | 193 |
| Configuring location data for LLDP-MED devices..... | 194 |
| Configuring coordinate-based locations..... | 195 |
| Viewing switch information available for outbound advertisements..... | 197 |
| Displaying the current port speed and duplex configuration on a switch port..... | 198 |
| Viewing the current port speed and duplex configuration on a switch port..... | 199 |
| Viewing advertisements currently in the neighbors MIB..... | 199 |
| Displaying LLDP statistics..... | 200 |
| Viewing LLDP statistics..... | 200 |
| LLDP Operating Notes..... | 202 |
| Neighbor maximum..... | 202 |
| LLDP packet forwarding..... | 202 |
| One IP address advertisement per port..... | 202 |
| 802.1Q VLAN Information..... | 203 |
| Effect of 802.1X Operation..... | 203 |
| Neighbor data can remain in the neighbor database after the neighbor is disconnected..... | 203 |
| Mandatory TLVs..... | 203 |
| LLDP and CDP data management..... | 203 |
| LLDP and CDP neighbor data..... | 203 |
| CDP operation and commands..... | 204 |
| Viewing the current CDP configuration of the switch..... | 205 |
| 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..... | 206 |
| Filtering CDP information..... | 208 |
| Configuring the switch to filter untagged traffic..... | 209 |
| Displaying the configuration..... | 209 |
| Filtering PVID mismatch log messages..... | 209 |
| DHCPv4 server..... | 210 |
| Introduction to DHCPv4..... | 210 |
| IP pools..... | 210 |
| DHCP options..... | 210 |
| BootP support..... | 210 |
| Authoritative server and support for DHCP inform packets..... | 211 |
| Authoritative pools..... | 211 |
| Authoritative dummy pools..... | 211 |
| Change in server behavior..... | 211 |
| DHCPv4 configuration commands..... | 212 |
| Enable/disable the DHCPv4 server..... | 212 |
| Configuring the DHCP address pool name..... | 212 |
| Authoritative..... | 213 |
| Specify a boot file for the DHCP client | 213 |
| Configure a default router for a DHCP client..... | 213 |

| | |
|---|------------|
| Configure the DNS IP servers | 214 |
| Configure a domain name..... | 214 |
| Configure lease time..... | 214 |
| Configure the NetBIOS WINS servers..... | 214 |
| Configure the NetBIOS node type..... | 214 |
| Configure subnet and mask | 215 |
| Configure DHCP server options..... | 215 |
| Configure the range of IP address..... | 215 |
| Configure the static binding information..... | 215 |
| Configure the TFTP server domain name..... | 216 |
| Configure the TFTP server address..... | 216 |
| Change the number of ping packets..... | 216 |
| Change the amount of time..... | 216 |
| Configure DHCP Server to save automatic bindings..... | 216 |
| Configure a DHCP server to send SNMP notifications..... | 217 |
| Enable conflict logging on a DHCP server..... | 217 |
| Enable the DHCP server on a VLAN..... | 217 |
| Clear commands..... | 217 |
| Reset all DHCP server and BOOTP counters..... | 217 |
| Delete an automatic address binding..... | 217 |
| Show commands..... | 218 |
| Display the DHCPv4 server address bindings..... | 218 |
| Display address conflicts..... | 218 |
| Display DHCPv4 server database agent..... | 218 |
| Display DHCPv4 server statistics..... | 218 |
| Display the DHCPv4 server IP pool information..... | 218 |
| Display DHCPv4 server global configuration information..... | 218 |
| Event log..... | 219 |
| Event Log Messages..... | 219 |
| 7 Link Aggregation Control Protocol—Multi-Active Detection (LACP-MAD)..... | 221 |
| LACP-MAD commands..... | 221 |
| Configuration command..... | 221 |
| show commands..... | 221 |
| clear command..... | 221 |
| LACP-MAD overview..... | 221 |
| 8 Scalability IP Address VLAN and Routing Maximum Values..... | 222 |
| 9 File Transfers..... | 224 |
| Overview..... | 224 |
| Downloading switch software..... | 224 |
| General software download rules..... | 224 |
| Using TFTP to download software from a server..... | 224 |
| Downloading from a server to primary flash using TFTP (Menu)..... | 225 |
| Troubleshooting TFTP download failures..... | 226 |
| Downloading from a server to flash using TFTP (CLI)..... | 227 |
| Enabling TFTP (CLI)..... | 228 |
| Configuring the switch to download software automatically from a TFTP server using auto-TFTP (CLI)..... | 228 |
| Using SCP and SFTP..... | 229 |
| Enabling SCP and SFTP..... | 230 |
| Disabling TFTP and auto-TFTP for enhanced security..... | 230 |
| Enabling SSH V2 (required for SFTP)..... | 232 |
| Confirming that SSH is enabled..... | 232 |
| Disabling secure file transfer..... | 232 |

| | |
|--|-----|
| Authentication..... | 232 |
| SCP/SFTP operating notes..... | 232 |
| Troubleshooting SSH, SFTP, and SCP operations..... | 234 |
| Broken SSH connection..... | 234 |
| Attempt to start a session during a flash write..... | 234 |
| Failure to exit from a previous session..... | 234 |
| Attempt to start a second session..... | 234 |
| Using Xmodem to download switch software from a PC or UNIX workstation..... | 235 |
| Downloading to primary flash using Xmodem (Menu)..... | 235 |
| Downloading to primary or secondary flash using Xmodem and a terminal emulator (CLI).... | 236 |
| Using USB to transfer files to and from the switch..... | 236 |
| Downloading switch software using USB (CLI)..... | 237 |
| Switch-to-switch download..... | 238 |
| Switch-to-switch download to primary flash (Menu)..... | 238 |
| Downloading the OS from another switch (CLI)..... | 239 |
| Downloading from primary only (CLI)..... | 239 |
| Downloading from either flash in the source switch to either flash in the destination switch (CLI)..... | 239 |
| Using IMC to update switch software..... | 240 |
| Copying software images..... | 240 |
| TFTP: Copying a software image to a remote host (CLI)..... | 240 |
| Xmodem: Copying a software image from the switch to a serially connected PC or UNIX workstation (CLI)..... | 240 |
| USB: Copying a software image to a USB device (CLI)..... | 240 |
| Transferring switch configurations..... | 241 |
| TFTP: Copying a configuration file to a remote host (CLI)..... | 241 |
| TFTP: Copying a configuration file from a remote host (CLI)..... | 241 |
| TFTP: Copying a customized command file to a switch (CLI)..... | 242 |
| Xmodem: Copying a configuration file to a serially connected PC or UNIX workstation (CLI).... | 242 |
| Xmodem: Copying a configuration file from a serially connected PC or UNIX workstation (CLI).... | 243 |
| USB: Copying a configuration file to a USB device (CLI)..... | 243 |
| USB: Copying a configuration file from a USB device (CLI)..... | 244 |
| Transferring ACL command files..... | 244 |
| TFTP: Uploading an ACL command file from a TFTP server (CLI)..... | 244 |
| Xmodem: Uploading an ACL command file from a serially connected PC or UNIX workstation (CLI)..... | 246 |
| USB: Uploading an ACL command file from a USB device (CLI)..... | 247 |
| Copying diagnostic data to a remote host, USB device, PC or UNIX workstation..... | 247 |
| Copying command output to a destination device (CLI)..... | 248 |
| Copying Event Log output to a destination device (CLI)..... | 248 |
| Copying crash data content to a destination device (CLI)..... | 249 |
| Flight Data Recorder (FDR)..... | 250 |
| Using USB autorun..... | 250 |
| Security considerations..... | 251 |
| Troubleshooting autorun operations..... | 252 |
| USB auxiliary port LEDs..... | 252 |
| AutoRun status files..... | 252 |
| Event log or syslog..... | 252 |
| Configuring autorun on the switch (CLI)..... | 253 |
| Autorun secure mode..... | 253 |
| Operating notes and restrictions..... | 253 |
| Autorun and configuring passwords..... | 254 |
| Viewing autorun configuration information..... | 254 |

| | |
|--|------------|
| 10 Monitoring and Analyzing Switch Operation..... | 255 |
| Overview..... | 255 |
| Status and counters data..... | 255 |
| Accessing status and counters (Menu)..... | 255 |
| General system information..... | 256 |
| Accessing system information (Menu)..... | 256 |
| Accessing system information (CLI)..... | 256 |
| Collecting processor data with the task monitor (CLI)..... | 257 |
| Task usage reporting..... | 258 |
| Switch management address information..... | 260 |
| Accessing switch management address information (Menu)..... | 260 |
| Accessing switch management address information (CLI)..... | 261 |
| Port Status..... | 261 |
| Viewing port status (CLI)..... | 261 |
| Viewing port status (Menu)..... | 261 |
| Viewing port and trunk group statistics (WebAgent)..... | 262 |
| Port and trunk group statistics and flow control status..... | 262 |
| Accessing port and trunk statistics (Menu)..... | 262 |
| Accessing port and trunk group statistics (CLI)..... | 263 |
| Viewing the port counter summary report..... | 263 |
| Viewing a detailed traffic summary for specific ports..... | 263 |
| Displaying trunk load balancing statistics..... | 264 |
| Clearing trunk load balancing statistics..... | 264 |
| Resetting the port counters..... | 264 |
| Viewing the switch's MAC address tables..... | 265 |
| Accessing MAC address views and searches (CLI)..... | 265 |
| Listing all learned MAC addresses on the switch, with the port number on which each MAC address was learned..... | 265 |
| Listing all learned MAC addresses on one or more ports, with their corresponding port numbers..... | 265 |
| Listing all learned MAC addresses on a VLAN, with their port numbers..... | 265 |
| Finding the port on which the switch learned a specific MAC address..... | 265 |
| Accessing MAC address views and searches (Menu)..... | 265 |
| Viewing and searching per-VLAN MAC-addresses..... | 265 |
| Finding the port connection for a specific device on a VLAN..... | 266 |
| Viewing and searching port-level MAC addresses..... | 267 |
| Determining whether a specific device is connected to the selected port..... | 267 |
| Accessing MSTP Data (CLI)..... | 267 |
| Viewing internet IGMP status (CLI)..... | 268 |
| Viewing VLAN information (CLI)..... | 269 |
| WebAgent status information..... | 271 |
| Interface monitoring features..... | 272 |
| Configuring port and static trunk monitoring (Menu)..... | 272 |
| Configuring port and static trunk monitoring (CLI)..... | 273 |
| Displaying the monitoring configuration..... | 273 |
| Configuring the monitor port..... | 274 |
| Selecting or removing monitoring source interfaces..... | 274 |
| 11 Troubleshooting..... | 276 |
| Overview..... | 276 |
| Troubleshooting approaches..... | 276 |
| Browser or Telnet access problems..... | 277 |
| Cannot access the WebAgent..... | 277 |
| Cannot Telnet into the switch console from a station on the network..... | 277 |
| Unusual network activity..... | 278 |

| | |
|--|-----|
| General problems..... | 278 |
| The network runs slow; processes fail; users cannot access servers or other devices..... | 278 |
| Duplicate IP addresses..... | 278 |
| Duplicate IP addresses in a DHCP network..... | 278 |
| The switch has been configured for DHCP/Bootp operation, but has not received a DHCP or Bootp reply..... | 279 |
| 802.1Q Prioritization problems..... | 279 |
| Ports configured for non-default prioritization (level 1 to 7) are not performing the specified action..... | 279 |
| Addressing ACL problems..... | 280 |
| ACLs are properly configured and assigned to VLANs, but the switch is not using the ACLs to filter IP layer 3 packets..... | 280 |
| The switch does not allow management access from a device on the same VLAN..... | 280 |
| Error (Invalid input) when entering an IP address..... | 281 |
| Apparent failure to log all "deny" matches..... | 281 |
| The switch does not allow any routed access from a specific host, group of hosts, or subnet... | 281 |
| The switch is not performing routing functions on a VLAN..... | 281 |
| Routing through a gateway on the switch fails..... | 281 |
| Remote gateway case..... | 281 |
| Local gateway case..... | 282 |
| IGMP-related problems..... | 282 |
| IP multicast (IGMP) traffic that is directed by IGMP does not reach IGMP hosts or a multicast router connected to a port..... | 282 |
| IP multicast traffic floods out all ports; IGMP does not appear to filter traffic..... | 283 |
| LACP-related problems..... | 283 |
| Unable to enable LACP on a port with the interface <i><port-number></i> lacp command..... | 283 |
| Port-based access control (802.1X)-related problems..... | 283 |
| The switch does not receive a response to RADIUS authentication requests..... | 283 |
| The switch does not authenticate a client even though the RADIUS server is properly configured and providing a response to the authentication request..... | 283 |
| During RADIUS-authenticated client sessions, access to a VLAN on the port used for the client sessions is lost..... | 284 |
| 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..... | 284 |
| The supplicant statistics listing shows multiple ports with the same authenticator MAC address..... | 284 |
| 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..... | 284 |
| RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch..... | 284 |
| 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..... | 285 |
| A trunked port configured for 802.1X is blocked..... | 285 |
| QoS-related problems..... | 285 |
| Loss of communication when using VLAN-tagged traffic..... | 285 |
| Radius-related problems..... | 285 |
| The switch does not receive a response to RADIUS authentication requests..... | 285 |
| RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch..... | 286 |
| MSTP and fast-uplink problems..... | 286 |
| Broadcast storms appearing in the network..... | 286 |
| STP blocks a link in a VLAN even though there are no redundant links in that VLAN..... | 286 |
| Fast-uplink troubleshooting..... | 287 |
| SSH-related problems..... | 287 |

| | |
|---|-----|
| Switch access refused to a client..... | 287 |
| Executing IP SSH does not enable SSH on the switch..... | 287 |
| Switch does not detect a client's public key that does appear in the switch's public key file (show ip client-publickey)..... | 287 |
| An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages..... | 287 |
| Client ceases to respond ("hangs") during connection phase..... | 288 |
| TACACS-related problems..... | 288 |
| Event Log..... | 288 |
| All users are locked out of access to the switch..... | 288 |
| No communication between the switch and the TACACS+ server application..... | 288 |
| Access is denied even though the username/password pair is correct..... | 289 |
| Unknown users allowed to login to the switch..... | 289 |
| System allows fewer login attempts than specified in the switch configuration..... | 289 |
| TimeP, SNTP, or Gateway problems..... | 289 |
| The switch cannot find the time server or the configured gateway..... | 289 |
| VLAN-related problems..... | 289 |
| Monitor port..... | 289 |
| None of the devices assigned to one or more VLANs on an 802.1Q-compliant switch are being recognized..... | 289 |
| Link configured for multiple VLANs does not support traffic for one or more VLANs..... | 289 |
| Duplicate MAC addresses across VLANs..... | 290 |
| Disabled overlapping subnet configuration..... | 290 |
| Fan failure..... | 291 |
| Mitigating flapping transceivers..... | 292 |
| Fault finder thresholds..... | 293 |
| Enabling fault finder using the CLI..... | 293 |
| Viewing transceiver information..... | 296 |
| Viewing information about transceivers (CLI)..... | 297 |
| MIB support..... | 297 |
| Viewing transceiver information..... | 297 |
| Information displayed with the detail parameter..... | 298 |
| Viewing transceiver information for copper transceivers with VCT support..... | 301 |
| Testing the Cable..... | 301 |
| Using the Event Log for troubleshooting switch problems..... | 303 |
| Event Log entries..... | 303 |
| Using the Menu..... | 311 |
| Using the CLI..... | 312 |
| Clearing Event Log entries..... | 313 |
| Turning event numbering on..... | 313 |
| Using log throttling to reduce duplicate Event Log and SNMP messages..... | 313 |
| Log throttle periods..... | 314 |
| Example: of event counter operation..... | 315 |
| Reporting information about changes to the running configuration..... | 316 |
| Debug/syslog operation..... | 316 |
| Debug/syslog messaging..... | 317 |
| Hostname in syslog messages..... | 317 |
| Logging origin-id..... | 317 |
| Viewing the identification of the syslog message sender..... | 319 |
| SNMP MIB..... | 320 |
| Debug/syslog destination devices..... | 321 |
| Debug/syslog configuration commands..... | 321 |
| Configuring debug/syslog operation..... | 324 |
| Viewing a debug/syslog configuration..... | 325 |
| Debug command..... | 327 |

| | |
|---|-----|
| Debug messages..... | 327 |
| Debug destinations..... | 329 |
| Logging command..... | 330 |
| Configuring a syslog server..... | 330 |
| Deleting syslog addresses in the startup configuration..... | 331 |
| Verifying the deletion of a syslog server address..... | 331 |
| Blocking the messages sent to configured syslog servers from the currently configured debug message type..... | 331 |
| Disabling syslog logging on the switch without deleting configured server addresses..... | 331 |
| Sending logging messages using TCP..... | 332 |
| Adding a description for a Syslog server..... | 333 |
| Adding a priority description..... | 334 |
| Configuring the severity level for Event Log messages sent to a syslog server..... | 334 |
| Configuring the system module used to select the Event Log messages sent to a syslog server..... | 335 |
| Operating notes for debug and Syslog..... | 335 |
| Diagnostic tools..... | 336 |
| Port auto-negotiation..... | 336 |
| Ping and link tests..... | 336 |
| Ping test..... | 336 |
| Link test..... | 336 |
| Executing ping or link tests (WebAgent)..... | 336 |
| Testing the path between the switch and another device on an IP network..... | 337 |
| Halting a ping test..... | 338 |
| Issuing single or multiple link tests..... | 339 |
| Tracing the route from the switch to a host address..... | 339 |
| Halting an ongoing traceroute search..... | 341 |
| A low maxttl causes traceroute to halt before reaching the destination address..... | 341 |
| If a network condition prevents traceroute from reaching the destination..... | 341 |
| Viewing switch configuration and operation..... | 342 |
| Viewing the startup or running configuration file..... | 342 |
| Viewing the configuration file (WebAgent)..... | 342 |
| Viewing a summary of switch operational data..... | 342 |
| Saving show tech command output to a text file..... | 343 |
| Customizing show tech command output..... | 344 |
| Viewing more information on switch operation..... | 346 |
| Searching for text using pattern matching with show command..... | 347 |
| Displaying the information you need to diagnose problems..... | 349 |
| Restoring the factory-default configuration..... | 350 |
| Resetting to the factory-default configuration..... | 350 |
| Using the CLI..... | 350 |
| Using Clear/Reset..... | 351 |
| Restoring a flash image..... | 351 |
| Recovering from an empty or corrupted flash state..... | 351 |
| DNS resolver..... | 352 |
| Basic operation..... | 353 |
| Configuring and using DNS resolution with DNS-compatible commands..... | 354 |
| Configuring a DNS entry..... | 354 |
| Using DNS names with ping and traceroute: Example:..... | 355 |
| Viewing the current DNS configuration..... | 356 |
| Operating notes..... | 357 |
| Event Log messages..... | 357 |
| Locating a switch (Locator LED)..... | 357 |

| | |
|--|------------|
| 12 MAC Address Management..... | 359 |
| Overview..... | 359 |
| Determining MAC addresses..... | 359 |
| Viewing the MAC addresses of connected devices..... | 359 |
| Viewing the switch's MAC address assignments for VLANs configured on the switch..... | 359 |
| Viewing the port and VLAN MAC addresses..... | 360 |
| A Network Out-of-Band Management (OOBM)..... | 362 |
| Concepts..... | 362 |
| Example:..... | 363 |
| OOBM and switch applications..... | 363 |
| OOBM configuration..... | 364 |
| Entering the OOBM configuration context from the general configuration context..... | 364 |
| Enabling and disabling OOBM..... | 364 |
| Enabling and disabling the OOBM port..... | 365 |
| Setting the OOBM port speed..... | 365 |
| Configuring an OOBM IPv4 address..... | 366 |
| Configuring an OOBM IPv4 default gateway..... | 366 |
| OOBM show commands..... | 366 |
| Showing the global OOBM and OOBM port configuration..... | 367 |
| Showing OOBM IP configuration..... | 367 |
| Showing OOBM ARP information..... | 367 |
| Application server commands..... | 367 |
| Application client commands..... | 368 |
| 13 Support and other resources..... | 371 |
| Contacting HP..... | 371 |
| Subscription service..... | 371 |
| Typographic conventions..... | 371 |
| 14 Documentation feedback..... | 373 |
| Index..... | 374 |

1 Time Protocols

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

General steps for running a time protocol on the switch

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

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

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

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

TimeP time synchronization

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

SNTP time synchronization

SNTP provides two operating modes:

- **Broadcast mode**

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

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

- **Unicast mode**

The switch requests a time update from the configured SNTP server. (You can configure one server using the menu interface, or up to three servers using the CLI `sntp server` command.) This option provides increased security over the Broadcast mode by specifying which time server to use instead of using the first one detected through a broadcast.

Selecting a time synchronization protocol

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

3. Configure the remaining parameters for the time protocol you selected.

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

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

Disabling time synchronization

You can use either of the following methods to disable time synchronization without changing the Timep or SNTP configuration:

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

SNTP: Selecting and configuring

Table 1 (page 18) shows the SNTP parameters and their operations.

Table 1 SNTP parameters

| SNTP parameter | Operation |
|--------------------------------|--|
| Time Sync Method | Used to select either SNTP, TIMEP, or None as the time synchronization method. |
| SNTP Mode | |
| Disabled | The Default. SNTP does not operate, even if specified by the Menu interface Time Sync Method parameter or the CLI <code>timesync</code> command. |
| Unicast | Directs the switch to poll a specific server for SNTP time synchronization. Requires at least one server address. |
| Broadcast | Directs the switch to acquire its time synchronization from data broadcast by any SNTP server to the network broadcast address. The switch uses the first server detected and ignores any others. However, if the Poll Interval expires three times without the switch detecting a time update from the original server, the switch accepts a broadcast time update from the next server it detects. |
| Poll Interval (seconds) | <p>In Unicast Mode: Specifies how often the switch polls the designated SNTP server for a time update.</p> <p>In Broadcast Mode: Specifies how often the switch polls the network broadcast address for a time update.</p> <p>Value is between 30 to 720 seconds.</p> |
| Server Address | Used only when the SNTP Mode is set to <code>Unicast</code> . Specifies the IP address of the SNTP server that the switch accesses for time synchronization updates. You can configure up to three servers; one using the menu or CLI, and two more using the CLI. |
| Server Version | Specifies the SNTP software version to use and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3. Default: 3; range: 1 to 7. |
| Priority | Specifies the order in which the configured servers are polled for getting the time. Value is between 1 and 3. |

Viewing and configuring SNTP (Menu)

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

Figure 1 System Information screen (default values)

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

System Name : HP Switch
System Contact :
System Location :

Inactivity Timeout (min) [0] : 0      MAC Age Time (sec) [300] : 300
Inbound Telnet Enabled [Yes] : Yes    Web Agent Enabled [Yes] : Yes
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled
Tftp-enable [Yes] : Yes                Server Address :
Time Zone [0] : 0                      Jumbo Max Frame Size [9216] : 9216
Daylight Time Rule [None] : None       Jumbo IP MTU [9198] : 9198

Actions->  Cancel      Edit      Save      Help

Time Protocol Selection Parameter
- TIMEP
- SNTP
- None
```

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

Option 1

- a. Use the **Space** bar to select the **Broadcast** mode.
- b. Move the cursor to the **Poll Interval** field.
- c. Go to step 6 (page 20). (For Broadcast mode details, see “SNTP time synchronization” (page 17))

Figure 2 Time configuration fields for SNTP with broadcast mode

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

Option 2

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

NOTE: This step replaces any previously configured server IP address. If you will be using backup SNTP servers (requires use of the CLI), see “SNTP unicast time polling with multiple SNTP servers” (page 38).

- g. Move the cursor to the **Server Version** field. Enter the value that matches the SNTP server version running on the device you specified in the preceding step .
If you are unsure which version to use, HP recommends leaving this value at the default setting of 3 and testing SNTP operation to determine whether any change is necessary.

NOTE: Using the menu to enter the IP address for an SNTP server when the switch already has one or more SNTP servers configured, the switch deletes the primary SNTP server from the server list. The switch then selects a new primary SNTP server from the IP addresses in the updated list. For more on this topic, see “SNTP unicast time polling with multiple SNTP servers” (page 38).

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

Figure 3 SNTP configuration fields for SNTP configured with unicast mode

```

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

```

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

6. In the **Poll Interval** field, enter the time in seconds that you want for a Poll Interval. (For Poll Interval operation, see [Table 1 \(page 18\)](#), on “SNTP parameters” (page 18).)
7. Press **Enter** to return to the Actions line, then **S** (for **Save**) to enter the new time protocol configuration in both the startup-config and running-config files.

Viewing and configuring SNTP (CLI)

Syntax:

```
show sntp
```

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

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

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

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

```
SNTP Configuration
```

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

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

In the factory-default configuration (where TimeP is the selected time synchronization method), `show sntp` still lists the SNTP configuration, even though it is not currently in use. In [Example 2](#), even though TimeP is the current time synchronous method, the switch maintains the SNTP configuration.

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

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

```
SNTP Configuration
```

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

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

Syntax:

```
show management
```

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

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

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

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

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

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

```
Default Gateway :10.0.9.80
```

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

Configuring (enabling or disabling) the SNTP mode

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

Syntax:

```
timesync sntp
```

Selects SNTP as the time protocol.

```
sntp <broadcast | unicast>
```

Enables the SNTP mode.

Syntax:

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

Required only for unicast mode.

Syntax:

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

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

Syntax:

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

Configures the amount of time between updates of the system clock via SNTP.
Default: 720 seconds

Enabling SNTP in Broadcast Mode

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

Syntax:

```
timesync sntp
```

Selects SNTP as the time synchronization method.

Syntax:

```
sntp broadcast
```

Configures broadcast as the SNTP mode.

Example:

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

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

The commands and output would appear as follows:

Example 4 Enabling SNTP operation in Broadcast Mode

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

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

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

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

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

Enabling SNTP in unicast mode (CLI)

Like broadcast mode, configuring SNTP for unicast mode enables SNTP. However, for unicast operation, you must also specify the IP address of at least one SNTP server. The switch allows up to three unicast servers. You can use the Menu interface or the CLI to configure one server or to replace an existing unicast server with another. To add a second or third server, you must use the CLI. For more on SNTP operation with multiple servers, see [“SNTP unicast time polling with multiple SNTP servers” \(page 38\)](#)

Syntax:

```
timesync sntp
Selects SNTP as the time synchronization method.
```

Syntax:

```
sntp unicast
Configures the SNTP mode for unicast operation.
```

Syntax:

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

Use the no version of the command to disable SNTP.

priority Specifies the order in which the configured SNTP servers are polled for the time.

ip-address An IPv4 or IPv6 address of an SNTP server.

version The protocol version of the SNTP server. Allowable values are 1 through 7; default is 3.

Syntax:

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

Deletes the specified SNTP server.

NOTE: Deleting an SNTP server when only one is configured disables SNTP unicast operation.

Example:

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

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

Selects SNTP.

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

Activates SNTP in unicast mode.

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

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

Example 5 Configuring SNTP for unicast operation

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

Sntp Configuration

Time Sync Mode: Sntp

Sntp Mode : Unicast

Poll Interval (sec) [720] : 720

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

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

Both IPv4 and IPv6 addresses are displayed.

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

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

Example 6 Specifying the SNTP protocol version number

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

SNTP Configuration

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

IP Address      Protocol Version
-----
10.28.227.141  4 3
```

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

Changing the SNTP poll interval (CLI)

Syntax:

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

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

Example:

To change the poll interval to 300 seconds:

```
HP Switch(config)# sntp 300
```

Changing the SNTP server priority (CLI)

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

Syntax:

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

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

NOTE: You can enter both IPv4 and IPv6 addresses. For more information about IPv6 addresses, see the *IPv6 Configuration Guide* for your switch.

Example:

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

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

```
HP Switch(config)# sntp server priority 2
                    2001:db8::215:60ff:fe79:8980
```

Disabling time synchronization without changing the SNTP configuration (CLI)

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

Syntax:

```
no timesync
```

Halts time synchronization without changing your SNTP configuration.

Example:

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

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

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

Example 7 SNTP with time synchronization disabled

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

Disabling the SNTP Mode

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

Syntax:

```
no sntp
```

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

Example:

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

Example 8 Disabling time synchronization by disabling the SNTP mode

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

SNTP Configuration

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

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

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

SNTP client authentication

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

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

Requirements

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

SNTP client authentication support

- Timesync mode must be SNTP. Use the `timesync sntp` command. (SNTP is disabled by default).
- SNTP must be in unicast or broadcast mode. See “[Configuring unicast and broadcast mode for authentication](#)” (page 29).
- The MD5 authentication mode must be selected.
- An SNTP authentication key-identifier (`key-id`) must be configured on the switch and a value (`key-value`) must be provided for the authentication key. A maximum of 8 sets of `key-id` and `key-value` can be configured on the switch.
- Among the keys that have been configured, one key or a set of keys must be configured as trusted. Only trusted keys are used for SNTP authentication.
- If the SNTP server requires authentication, one of the trusted keys has to be associated with the SNTP server.
- SNTP client authentication must be enabled on the HP Switch. If client authentication is disabled, packets are processed without authentication.

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

SNTP server authentication support

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

You must perform the following on the SNTP server:

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

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

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

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

Syntax:

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

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

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

Default: No default keys are configured on the switch.

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

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

Example 9 Setting parameters for SNTP authentication

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

Configuring a trusted key

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

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

In broadcast mode, the SNTP client switch checks the size of the received packet to determine if it is authenticated. If the broadcast packet is authenticated, the key-id value is checked to see if the same key-id value is configured on the SNTP client switch. If the switch is configured with the same key-id value, and the key-id value is configured as "trusted," the authentication succeeds. Only trusted key-id value information is used for SNTP authentication. For information about configuring these modes, see [“Configuring unicast and broadcast mode for authentication” \(page 29\)](#).

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

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

Configuring a key-id as trusted (CLI)

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

Syntax:

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

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

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

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

Default: No key is trusted by default.

For detailed information about trusted keys, see [“Configuring a trusted key” \(page 28\)](#)

Associating a key with an SNTP server (CLI)

Syntax:

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

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

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

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

| | |
|--------------------------|---|
| <code>priority</code> | Specifies the order in which the configured servers are polled for getting the time. |
| <code>version-num</code> | Specifies the SNTP software version to use and is assigned on a per-server basis. The version setting is backwards-compatible. For example, using version 3 means that the switch accepts versions 1 through 3. Default: 3; range: 1 - 7. |
| <code>key-id</code> | Optional command. The key identifier sent in the SNTP packet. This <code>key-id</code> is associated with the SNTP server specified in the command. |

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

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

Enabling SNTP client authentication

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

Syntax:

```
[no] sntp authentication
```

Enables the SNTP client authentication.

The `no` version of the command disables authentication.

Default: SNTP client authentication is disabled.

Configuring unicast and broadcast mode for authentication

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

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

Syntax:

```
sntp unicast  
sntp broadcast
```

Enables SNTP for either broadcast or unicast mode.

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

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

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

| | |
|-----------|---|
| Broadcast | Directs the switch to acquire its time synchronization from data broadcast by any SNTP server to the network broadcast address. The switch uses the first server detected and ignores any others. However, if the Poll Interval (configurable up to 720 seconds) expires three times without the switch detecting a time update from the original server, the switch accepts a broadcast time update from the next server it detects. |
|-----------|---|

Viewing SNTP authentication configuration information (CLI)

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

Example 11 SNTP configuration information

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

```
SNTP Configuration
```

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

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

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

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

Example 12 Show sntp authentication command output

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

```
SNTP Authentication Information
```

```
SNTP Authentication : Enabled
```

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

Viewing statistical information for each SNTP server (CLI)

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

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

Example 13 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
```

Saving configuration files and the include-credentials command

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

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

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

Example 14 Configuration file with SNMP authentication information

```
HP Switch (config) # show config
Startup configuration:
.
.
.
timesync snmp
snmp broadcast
snmp 50
snmp authentication
snmp server priority 1 10.10.10.2.3 key-id 55
snmp server priority 2 fe80::200:24ff:fec8:4ca8 4 key-id 55
```

NOTE: SNMP 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 SNMP authentication information is not saved and is not present in the retrieved configuration files, as shown in the following Example:

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

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

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

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

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

```
HP Switch(config)# show config

Startup configuration:

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

TimeP: Selecting and configuring

Table 2 (page 32) shows TimeP parameters and their operations.

Table 2 TimeP parameters

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

Table 2 TimeP parameters *(continued)*

| TimeP parameter | Operation |
|-----------------------|--|
| Disabled | The Default. Timep does not operate, even if specified by the Menu interface Time Sync Method parameter or the CLI <code>timesync</code> command. |
| DHCP | When Timep is selected as the time synchronization method, the switch attempts to acquire a Timep server IP address via DHCP. If the switch receives a server address, it polls the server for updates according to the Timep poll interval. If the switch does not receive a Timep server IP address, it cannot perform time synchronization updates. |
| Manual | When Timep is selected as the time synchronization method, the switch attempts to poll the specified server for updates according to the Timep poll interval. If the switch fails to receive updates from the server, time synchronization updates do not occur. |
| Server Address | Used only when the TimeP Mode is set to Manual . Specifies the IP address of the TimeP server that the switch accesses for time synchronization updates. You can configure one server. |

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

- From the Main Menu, select:

2. Switch Configuration

1. System Information

Figure 5 System Information screen (default values)

```

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

System Name : HP Switch
System Contact :
System Location :

Inactivity Timeout (min) [0] : 0      MAC Age Time (sec) [300] : 300
Inbound Telnet Enabled [Yes] : Yes   Web Agent Enabled [Yes] : Yes
Time Sync Method [None] : TIMEP
TimeP Mode [Disabled] : Disabled
Tftp-enable [Yes] : Yes               Server Address :
Time Zone [0] : 0                    Jumbo Max Frame Size [9216] : 9216
Daylight Time Rule [None] : None     Jumbo IP MTU [9198] : 9198

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

Actions->  Cancel  Edit  Save  Help
    
```

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

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

- Move the cursor to the **Time Sync Method** field.
- If **TIMEP** is not already selected, use the **Space** bar to select **TIMEP**, then move to the **TIMEP Mode** field.
- Do one of the following:
 - Use the **Space** bar to select the **DHCP** mode.
 - Move the cursor to the **Poll Interval** field.
 - Go to step 6.

Enabling TIMEP or DHCP

```

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

Daylight Time Rule [None] : None

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

NOTE: This step replaces any previously configured TimeP server IP address.

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

Viewing the current TimeP configuration (CLI)

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

Syntax:

```
show timep
```

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

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

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

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

```
Timep Configuration
```

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

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

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

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

```
Timep Configuration
```

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

Syntax:

```
show management
```

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

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

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

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

```
Time Server Address : 10.10.28.100
```

| Priority | SNTP Server Address | Protocol | Version |
|----------|---------------------|----------|---------|
| 1 | 10.10.28.101 | 3 | |
| 2 | 10.255.5.24 | 3 | |
| 3 | fe80::123%vlan10 | 3 | |

```
Default Gateway : 10.0.9.80
```

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

Configuring (enabling or disabling) the TimeP mode

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

Syntax:

```
timesync timep
```

Selects TimeP as the time synchronization method.

Syntax:

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

Enables the selected TimeP mode.

Syntax:

```
[no] ip timep
```

Disables the TimeP mode.

Syntax:

```
[no] timesync
```

Disables the time protocol.

Enabling TimeP in manual mode (CLI)

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

Syntax:

```
timesync timep
```

Selects TimeP.

Syntax:

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

Activates TimeP in manual mode with a specified TimeP server.

Syntax:

```
no ip timep
```

Disables TimeP.

Enabling TimeP in DHCP Mode

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

Syntax:

```
timesync timep
```

Selects TimeP as the time synchronization method.

Syntax:

```
ip timep dhcp
```

Configures DHCP as the TimeP mode.

For example, suppose:

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

Enabling TimeP in Manual Mode

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

Syntax:

```
timesync timep
```

Selects TimeP.

Syntax:

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

Activates TimeP in manual mode with a specified TimeP server.

Syntax:

```
[no]ip timep
```

Disables TimeP.

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

Example:

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

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

Selects TimeP.

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

Activates TimeP in Manual mode.

Example 19 Configuring TimeP for manual operation

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

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

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

```
Timep Configuration
```

```
Time Sync Mode: Timep
```

```
TimeP Mode : Manual
```

```
Server Address : 10.28.227.141
```

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

Changing from one TimeP server to another (CLI)

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

Changing the TimeP poll interval (CLI)

Syntax:

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

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

Example:

To change the poll interval to 60 minutes:

```
HP Switch(config)# ip timep interval 60
```

Disabling time synchronization without changing the TimeP configuration (CLI)

Syntax:

```
no timesync
```

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

Example:

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

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

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

Example 20 TimeP with time synchronization disabled

```
HP Switch(config)# show timep

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

Disabling the TimeP mode

Syntax:

```
no ip timep
```

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

Example:

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

Example 21 Disabling time synchronization by disabling the TimeP mode parameter

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

HP Switch(config)# show timep

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

SNTP unicast time polling with multiple SNTP servers

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

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

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

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

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

```
HP Switch(config)# show management

Status and Counters - Management Address Information

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

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

Default Gateway : 10.0.9.80

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

Adding and deleting SNTP server addresses

Adding addresses

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

Example 23 Creating additional SNTP server addresses with the CLI

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

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

Deleting addresses

Syntax:

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

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

Example:

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

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

Operating with multiple SNTP server addresses configured (Menu)

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

SNTP messages in the Event Log

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

Monitoring resources

Displaying current resource usage

To display current resource usage in the switch, enter the following command:

Syntax:

```
show <qos | access-list | policy> resources
```

Displays the resource usage of the policy enforcement engine on the switch by software feature. For each type of resource, the amount still available and the amount used by each software feature is shown.

| | |
|--|---|
| 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 "Viewing OpenFlow Resources" in the <i>OpenFlow Administrators Guide</i> for your switch. |

“Displaying current resource usage” (page 41) shows the resource usage on a switch configured for ACLs, QoS, RADIUS-based authentication, and other features:

- The "Rules Used" columns show that ACLs, VT, mirroring, and other features (For example, Management VLAN) have been configured globally or per-VLAN, because identical resource consumption is displayed for each port range in the switch. If ACLs were configured per-port, the number of rules used in each port range would be different.

Example 24 Displaying current resource usage

```
HP Switch(config)# show access-list resources
```

```
Resource usage in Policy Enforcement Engine
```

| Ports | Rules | Rules Used | | | |
|-------|-----------|------------|-----|-----|-------|
| | Available | ACL | QoS | IDM | Other |
| 1-48 | 2006 | 10 | 5 | 0 | 6 |

| Ports | Meters | Meters Used | | | |
|-------|-----------|-------------|-----|-----|-------|
| | Available | ACL | QoS | IDM | Other |
| 1-48 | 255 | | 5 | | 0 |

| Ports | Application Port Ranges | Application Port Ranges Used | | | |
|-------|-------------------------|------------------------------|-----|-----|-------|
| | Available | ACL | QoS | IDM | Other |
| 1-48 | 31 | 1 | 0 | 0 | 0 |

```
2 of 16 Policy Engine management resources used.
```

Key:

ACL = Access Control Lists

QoS = Device & Application Port Priority

IDM = Identity Driven Management

Other = Management VLAN, DHCP Snooping, ARP Protection, RA Guard.

Resource usage includes resources actually in use, or reserved for future use by the listed feature. Internal dedicated-purpose resources, such as port bandwidth limits or VLAN QoS priority, are not included.

Viewing information on resource usage

The switch allows you to view information about the current usage and availability of resources in the Policy Enforcement engine, including the following software features:

- Access control lists (ACL)
- Quality-of-service (QoS), including device and application port priority, ICMP rate-limiting, and QoS policies
- Dynamic assignment of per-port or per-user ACLs and QoS through RADIUS authentication designated as “IDM”, with or without the optional identity-driven management (IDM) application
- Virus throttling (VT) using connection-rate filtering

- Mirroring policies, including switch configuration as an endpoint for remote intelligent mirroring
- Other features, including:
 - Management VLAN
 - DHCP snooping
 - Dynamic ARP protection
 - Jumbo IP-MTU

Policy enforcement engine

The policy enforcement engine is the hardware element in the switch that manages QoS, mirroring, and ACL policies, as well as other software features, using the rules that you configure. Resource usage in the policy enforcement engine is based on how these features are configured on the switch:

- Resource usage by dynamic port ACLs is determined as follows:
 - Dynamic port ACLs configured by a RADIUS server (with or without the optional IDM application) for an authenticated client determine the current resource consumption for this feature on a specified slot. When a client session ends, the resources in use for that client become available for other uses.
- When the following features are configured globally or per-VLAN, resource usage is applied across all port groups or all slots with installed modules:
 - ACLs
 - QoS configurations that use the following commands:
 - QoS device priority (IP address) through the CLI using the `qos device-priority` command
 - QoS application port through the CLI using `qos tcp-port` or `qos udp-port`
 - VLAN QoS policies through the CLI using `service-policy`
 - Management VLAN configuration
 - DHCP snooping
 - Dynamic ARP protection
 - Remote mirroring endpoint configuration
 - Mirror policies per VLAN through the CLI using `monitor service`
 - Jumbo IP-MTU
- When the following features are configured per-port, resource usage is applied only to the slot or port group on which the feature is configured:
 - ACLs or QoS applied per-port or per-user through RADIUS authentication
 - ACLs applied per-port through the CLI using the `ip access-group` or `ipv6 traffic-filter` commands
 - QoS policies applied per port through the CLI using the `service-policy` command
 - Mirror policies applied per-port through the CLI using the `monitor all service` and `service-policy` commands
 - ICMP rate-limiting through the CLI using the `rate-limit icmp` command

Usage notes for show resources output

- A 1:1 mapping of internal rules to configured policies in the switch does not necessarily exist. As a result, displaying current resource usage is the most reliable method for keeping track of available resources. Also, because some internal resources are used by multiple features, deleting a feature configuration may not increase the amount of available resources.
- Resource usage includes resources actually in use or reserved for future use by the listed features.
- "Internal dedicated-purpose resources" include the following features:
 - Per-port ingress and egress rate limiting through the CLI using `rate-limit in/out`
 - Per-port or per-VLAN priority or DSCP through the CLI using `qos priority` or `qos dscp`
 - Per protocol priority through the CLI using `qos protocol`
- The "Available" columns display the resources available for additional feature use.
- The "IDM" column shows the resources used for RADIUS-based authentication with or without the IDM option.
- "Meters" are used when applying either ICMP rate-limiting or a QoS policy with a rate-limit class action.

When insufficient resources are available

The switch has ample resources for configuring features and supporting RADIUS-authenticated clients (with or without the optional IDM application).

If the resources supporting these features become fully subscribed:

- The current feature configuration, RADIUS-authenticated client sessions, and VT instances continue to operate normally.
- The switch generates an event log notice to say that current resources are fully subscribed.
- Currently engaged resources must be released before any of the following actions are supported:
 - Modifying currently configured ACLs, IDM, VT, and other software features, such as Management VLAN, DHCP snooping, and dynamic ARP protection.
You can modify currently configured classifier-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.

2 Port Status and Configuration

Viewing port status and configuring port parameters

Connecting transceivers to fixed-configuration devices

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

- To check the mode setting for a port on the switch, use either the Port Status screen in the menu interface or `show interfaces brief` in the CLI (see “Viewing port status and configuration (CLI)” (page 47)).
- To display information about the transceivers installed on a switch, enter the `show tech receivers` command in the CLI (Example 31 (page 51)).

Table 3 Status and parameters for each port type

| Status or parameter | Description |
|---------------------|--|
| Enabled | Yes (default): The port is ready for a network connection. No: The port will not operate, even if properly connected in a network. Use this setting, For example, if the port needs to be shut down for diagnostic purposes or while you are making topology changes. |
| Status (read-only) | Up: The port senses a link beat. Down: The port is not enabled, has no cables connected, or is experiencing a network error. For troubleshooting information, see the <i>Installation and Getting Started Guide</i> you received with the switch. See also to Appendix C, "Troubleshooting" (in this manual). |
| Mode | The port's speed and duplex (data transfer operation) setting. 10/100/1000Base-T Ports: <ul style="list-style-type: none">• Auto-MDIX (default): Senses speed and negotiates with the port at the other end of the link for port operation (MDI-X or MDI). To see what the switch negotiates for the auto setting, use the CLI <code>show interfaces brief</code> command or the 3. Port Status option under 1. Status and Counters in the menu interface.• MDI: Sets the port to connect with a PC using a crossover cable (manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables)• MDIX: Sets the port to connect with a PC using a straight-through cable (manual mode—applies only to copper port switches using twisted-pair copper Ethernet cables)• Auto-10: Allows the port to negotiate between half-duplex (HDx) and full-duplex (FDx) while keeping speed at 10 Mbps. Also negotiates flow control (enabled or disabled). HP recommends auto-10 for links between 10/100 auto-sensing ports connected with Cat 3 cabling. (Cat 5 cabling is required for 100 Mbps links.).• 10HDx: 10 Mbps, half-duplex• 10FDx: 10 Mbps, full-duplex• Auto-100: Uses 100 Mbps and negotiates with the port at the other end of the link for other port operation features.• Auto-10-100: Allows the port to establish a link with the port at the other end at either 10 Mbps or 100 Mbps, using the highest mutual speed and duplex mode available. Only these speeds are allowed with this setting.• Auto-1000: Uses 1000 Mbps and negotiates with the port at the other end of the link for other port operation features. |

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

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

Viewing port configuration (Menu)

The menu interface displays the configuration for ports and (if configured) any trunk groups. From the Main Menu, select:

1. Status and Counters
4. Port Status

Example 25 A switch port status screen

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

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

Actions->  Back   Intrusion log   Help

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

Configuring ports (Menu)

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

1. From the Main Menu, select:

2. Switch Configuration...

2. Port/Trunk Settings

Example 26 Port/trunk settings with a trunk group configured

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

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

Actions->  Cancel   Edit     Save     Help

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

2. Press [E] (for Edit).

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

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

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

Viewing port status and configuration (CLI)

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

Syntax:

| | |
|---|---|
| <code>show interfaces [brief config <port-list>]</code> | |
| <code>brief</code> | Lists the current operating status for all ports on the switch. |
| <code>config</code> | Lists a subset of configuration data for all ports on the switch; that is, for each port, the display shows whether the port is enabled, the operating mode, and whether it is configured for flow control. |
| <code><port-list></code> | Shows a summary of network traffic handled by the specified ports. |

Example 27 The show interfaces brief command listing

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

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

Example 28 The show interfaces config command listing

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

Port Settings

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

Dynamically updating the show interfaces command (CLI/Menu)

Syntax:

```
show interfaces display
```

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

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

Example:

HP Switch# show interfaces display

When using the **display** option in the CLI, the information stays on the screen and is updated every 3 seconds, as occurs with the display using the menu feature. The update is terminated with **Ctrl-C**. You can use the arrow keys to scroll through the screen when the output does not fit in one screen.

Figure 6 show interfaces display command with dynamically updating output

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

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

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

Customizing the show interfaces command (CLI)

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

Syntax:

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

Select the information that you want to display. Supported columns are shown in [Table 4 \(page 48\)](#).

Table 4 Supported columns, what they display, and examples:

| Parameter column | Displays | Examples |
|------------------|---|------------------------|
| port | Port identifier | A2 |
| type | Port type | 100/1000T |
| status | Port status | up or down |
| speed | Connection speed and duplex | 1000FDX |
| mode | Configured mode | auto, auto-100, 100FDX |
| mdi | MDI mode | auto, MDIX |
| flow | Flow control | on or off |
| name | Friendly port name | |
| vlanid | The vlan id this port belongs to, or "tagged" if it belongs to more than one vlan | 4 tagged |
| enabled | port is or is not enabled | yes or no intrusion |

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

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

Example 29 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 29](#), 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 \(page 49\)](#).

Error messages associated with the show interfaces command

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

Note on using pattern matching with the show interfaces custom command

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

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

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

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

Viewing port utilization statistics (CLI)

Use the `show interface port-utilization` command to view a real-time rate display for all ports on the switch. [Example 30](#) shows a sample output from this command.

Example 30 A show interface port-utilization command listing

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

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

Operating notes for viewing port utilization statistics

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

Viewing transceiver status (CLI)

The `show interfaces transceivers` command allows you to:

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

[Example 31](#) shows sample output from the `show tech transceivers` command.

NOTE: Part # column in [Example 31](#) enables you to determine the manufacturer for a specified transceiver and revision number.

Example 31 The show tech transceivers command

```
HP Switch# show tech transceivers
```

```
Transceiver Technical Information:
```

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

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

| Port # | Message |
|---------|--------------------|
| Port 23 | Self test failure. |

Operating notes

- The following information is displayed for each installed transceiver:
 - Port number on which transceiver is installed.
 - Type of transceiver.
 - Product number — Includes revision letter, such as A, B, or C. If no revision letter follows a product number, this means that no revision is available for the transceiver.
 -
 -
 -
 - Part number — Allows you to determine the manufacturer for a specified transceiver and revision number.
-
-
- TM
- For a non-HP switches installed transceiver (see line 23 [Example 31 \(page 51\)](#)), no transceiver type, product number, or part information is displayed. In the Serial Number field, non-operational is displayed instead of a serial number.
- The following error messages may be displayed for a non-operational transceiver:
 - Unsupported Transceiver. (SelfTest Err#060)
Check: www.hp.com/rnd/device_help/2_inform for more info.
 - This switch only supports revision B and above transceivers.
Check: www.hp.com/rnd/device_help/2_inform for more info.
 - Self test failure.
 - Transceiver type not supported in this port.
 - Transceiver type not supported in this software version.
 - Not an HP Switch Transceiver.
Go to: www.hp.com/rnd/device_help/2_inform for more info.

Enabling or disabling ports and configuring port mode (CLI)

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

See [Table 3 \(page 44\)](#).

Syntax:

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

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

```
speed-duplex
```

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

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

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

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

Examples:

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

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

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

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


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

```
HP Switch(config)# int e c6
```

```
HP Switch(eth-C6)# speed-duplex 100-full
```

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

Figure 7 Two methods for changing a port configuration



For more on flow control, see [“Enabling or disabling flow control \(CLI\)” \(page 52\)](#).

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 27 \(page 47\)](#).) Also, the port (`speed-duplex`) mode must be set to `Auto` (the default).

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

(You can find more information on flow control in [Table 3 \(page 44\)](#).)

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 32 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 | | | MDI Mode | Flow Ctrl | Bcast Limit | |
|------|---------|-----------|---------|--------|----------|-----------|-------------|---|
| | | Alert | Enabled | Status | | | | |
| 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 33 Continued from Example 32

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

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

Status and Counters - Port Status

| Port | Type | Intrusion | | | MDI Mode | Flow Ctrl | Bcast Limit | |
|------|---------|-----------|---------|--------|----------|-----------|-------------|---|
| | | Alert | Enabled | Status | | | | |
| 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 34 Continued from Example 33

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

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

Status and Counters - Port Status

| Port | Type | Intrusion | | | MDI Mode | Flow Ctrl | Bcast Limit | |
|------|---------|-----------|---------|--------|----------|-----------|-------------|---|
| | | Alert | Enabled | Status | | | | |
| 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 |

Port shutdown with broadcast storm

A LAN broadcast storm arises when an excessively high rate of broadcast packets flood the LAN. Occurrence of LAN broadcast storm disrupts traffic and degrades network performance. To prevent

LAN traffic from being disrupted, an enhancement of fault-finder commands adds new options, and the corresponding MIBs, that trigger a port disablement when a broadcast storm is detected on that port.

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

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

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

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

Syntax:

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

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

Syntax:

```
no fault-finder broadcast-storm [ethernet] <port-list>  
broadcast-storm      Configure broadcast storm control.  
pps                  Rising threshold level in number of broadcast packets  
                     per second.  
percent              Rising threshold level as a percentage of bandwidth of  
                     the port. The percentage is calculated on 64 byte packet  
                     size.  
warn                 Log the event only.  
warn-and-disable     Log the event and disable the port.  
seconds              Re-enable the port after waiting for the specified number  
                     of seconds. Default is not to re-enable.
```

Configuration examples:

```
HP Switch(config)# fault-finder broadcast-storm [ethernet] <A1> action [warn-and-disable  
<65535>]< percent 10>  
HP Switch(config)# fault-finder broadcast-storm [ethernet] <A2> action [warn-and-disable <pps  
100>  
HP Switch(config)# fault-finder broadcast-storm [ethernet] <A22> action [warn]  
<pps 100>
```

Viewing broadcast storm

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

Syntax:

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

Examples:

```
HP Switch# show fault-finder broadcast-storm [A1]
```

| Port | Bcast Storm | Port Status | Rising Threshold | Action | Disable Timer | Disable Timer Left |
|------|-------------|-------------|------------------|----------------|---------------|--------------------|
| A1 | Yes | Down | 10% | warnanddisable | 65535 | — |

```
HP Switch (config)# show fault-finder broadcast-storm
```

| Port | Bcast Storm | Port Status | Rising Threshold | Action | Disable Timer | Disable Timer Left |
|------|-------------|-------------|------------------|---------------|---------------|--------------------|
| A1 | Yes | Down | 200 pps | warnandisable | 10 | 9 |

```
HP Switch (config)# show fault-finder broadcast-storm A1
```

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

```
HP Switch (config)# show fault-finder broadcast-storm
```

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

SNMP MIB

SNMP support will be provided through the following MIB objects:

hpicFfbcastStormControlPortConfig OBJECT IDENTIFIER

:: = { hpicFaultFinder 5 }

hpicFfbcastStormControlPortConfigTable OBJECT-TYPE

- syntax sequence: **HpicFfbcastStormControlPortConfigEntry**
- max-access: not-accessible
- status: current
- description: This table provides information about broadcast storm control configuration of all ports.

::= {hpicFfbcastStormControlPortConfig 1}

hpicFfbcastStormControlPortConfigEntry OBJECT-TYPE

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

::= {hpicFfbcastStormControlPortConfigTable 1}

hpicFfbcastStormControlPortConfigEntry ::=

- Syntax sequence:
hpicFfbcastStormControlPortIndex InterfaceIndex,
hpicFfbcastStormControlMode Integer,
hpicFfbcastStormControlRisingpercent Integer32,
hpicFfbcastStormControlRisingpps Integer32,
hpicFfbcastStormControlAction Integer,
hpicFfbcastStormControlPortDisableTimer Unsigned32

hpicFfbcastStormControlPortIndex OBJECT-TYPE

- Syntax: Interfaceindex
- max-access: not-accessible
- status: current
- description: The Index Value Which Uniquely Identifies A Row In The Interfaces Table.
::= {hpicFfbcastStormControlPortConfigEntry 1}

hpicFfbcastStormControlMode OBJECT-TYPE

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

hpicFfbcastStormControlRisingpercent OBJECT-TYPE

- Syntax Integer32 (1..100)
- max-access: read-write
- status: current
- description: This Is The Rising Threshold Level in percent of bandwidth of the port. **hpicFfbcastStormControlAction** occurs when broadcast traffic reaches this level.
::= {hpicFfbcastStormControlPortConfigEntry 3}

hpicFfbcastStormControlRisingpps OBJECT-TYPE

- Syntax Integer32 (1..10000000)
- max-access: read-write
- status: current
- description: This object indicates the rising threshold for broadcast storm control. This value is in packets-per-second of received broadcast traffic. **hpicfffbcaststormcontrolaction** object takes action when broadcast traffic reaches this level.
::= {hpicFfbcastStormControlPortConfigEntry 4}

hpicffBcastStormControlAction OBJECT-TYPE

- Syntax integer: none(1), warn(2), warnanddisable(3)
- max-access: read-write
- status: current
- Description: This object defines the action taken by the switch when a broadcast storm occurs on a port. A value of none (1) indicates that no action is performed. A value of warn (2) indicates that an event is logged when broadcast traffic crosses the threshold value set on that port. A value of warn-and-disable (3) indicates that the port is disabled and an event is logged as soon as the broadcast traffic reaches the threshold value set on that port.
- DEFVAL: none
 ::= {hpicffBcastStormControlPortConfigEntry 5}

hpicffBcastStormControlPortDisableTimer OBJECT-TYPE

- Syntax Unsigned32 (0..604800)
- Units: seconds
- max-access: read-write
- status: current
- Description: This object specifies the time period for which the port remains in disabled state. A port is disabled when broadcast traffic reaches the threshold value set on that port. This time period is specified in seconds. The default value is zero which means that the port remains disabled and is not enabled again.
- DEFVAL {0}
 ::= {hpicffBcastStormControlPortConfigEntry 6}

Configuring auto-MDIX

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

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

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

Using the above ports:

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

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

For more information on this subject, see the *IEEE 802.3ab Standard Reference*. For more information on MDI-X, see the *Installation and Getting Started Guide* for your switch.

Manual override

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

- Manual MDI
- Manual MDI-X

Table 5 (page 59) shows the cabling requirements for the MDI/MDI-X settings.

Table 5 Cable types for auto and manual MDI/MDI-X settings

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

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

Configuring auto-MDIX (CLI)

The auto-MDIX features apply only to copper port switches using twisted-pair copper Ethernet cables. For information about auto-MDIX, see “Configuring auto-MDIX” (page 58).

Syntax:

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

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

Syntax:

```
show interfaces config
```

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

Syntax:

```
show interfaces brief
```

- Where a port is linked to another device, this command lists the MDI mode the port is currently using.
- In the case of ports configured for Auto (auto-mdix), the MDI mode appears as either MDI or MDIX, depending upon which option the port has negotiated with the device on the other end of the link.
- In the case of ports configured for MDI or MDIX, the mode listed in this display matches the configured setting.

- If the link to another device was up, but has gone down, this command shows the last operating MDI mode the port was using.
- If a port on a given switch has not detected a link to another device since the last reboot, this command lists the MDI mode to which the port is currently configured.

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

Example 35 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 36 Displaying the current MDI operating mode

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

Status and Counters - Port Status

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

Using friendly (optional) port names

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

Configuring and operating rules for friendly port names

- At either the global or context configuration level, you can assign a unique name to a port. You can also assign the same name to multiple ports.
- The friendly port names you configure appear in the output of the `show name [port-list]`, `show config`, and `show interface <port-number>` commands. They do not appear in the output of other `show` commands or in Menu interface screens. (See “[Displaying friendly port names with other port data \(CLI\)](#)” (page 62).)

- Friendly port names are not a substitute for port numbers in CLI commands or Menu displays.
- Trunking ports together does not affect friendly naming for the individual ports. (If you want the same name for all ports in a trunk, you must individually assign the name to each port.)
- A friendly port name can have up to 64 contiguous alphanumeric characters.
- Blank spaces within friendly port names are not allowed, and if used, cause an **invalid input** error. (The switch interprets a blank space as a name terminator.)
- In a port listing, **not assigned** indicates that the port does not have a name assignment other than its fixed port number.
- To retain friendly port names across reboots, you must save the current running-configuration to the startup-config file after entering the friendly port names. (In the CLI, use the `write memory` command.)

Configuring friendly port names (CLI)

For detailed information about friendly port names, see [“Using friendly \(optional\) port names” \(page 60\)](#).

Syntax:

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

Assigns a port name to port-list.

Syntax:

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

Deletes the port name from <port-list>.

Configuring a single port name (CLI)

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

Example 37 Configuring a friendly port name

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

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

Configuring the same name for multiple ports (CLI)

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

Example 38 Configuring one friendly port name on multiple ports

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

Port Names

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

```
Port : A6
Type : 10GbE-T
Name : Draft-Server:Trunk
```

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

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

Displaying friendly port names with other port data (CLI)

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

Syntax:

```
show name
```

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

Syntax:

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

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

Syntax:

```
show config
```

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

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

Syntax:

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

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

Example 39 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 40 Friendly port name data for specific ports on the switch

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

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

Syntax:

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

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

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

Example 41 A friendly port name in a per-port statistics listing

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

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

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

```
Name : not assigned
```

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

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

Syntax:

```
show config
```

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

See [Example 42 “Listing of the startup-config file with a friendly port name configured \(and saved\)”](#) to configure port A1 with a friendly port name. Notice that the command sequence saves the friendly port name for port A1 in the startup-config file. The name entered for port A2 is not saved because it was executed after `write memory`.

Example 42 Listing of the startup-config file with a friendly port name configured (and saved)

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

HP Switch(config)# show config

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

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

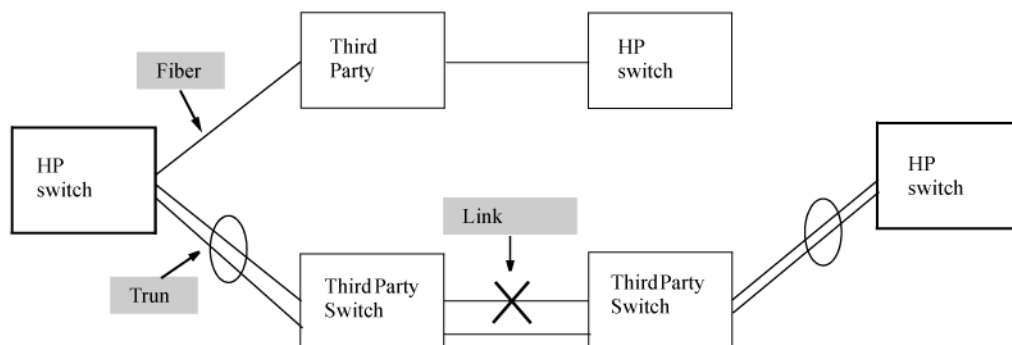
Uni-directional link detection (UDLD)

Uni-directional link detection (UDLD) monitors a link between two HP switches and blocks the ports on both ends of the link if the link fails at any point between the two devices. This feature is particularly useful for detecting failures in fiber links and trunks. [Figure 8 \(page 65\)](#) shows an Example:.

Figure 8 UDLD Example:

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

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



In this Example:, each HP switch load balances traffic across two ports in a trunk group. Without the UDLD feature, a link failure on a link that is not directly attached to one of the HP switches remains undetected. As a result, each switch continue 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-direction fibers to transmit and receive packets. Without UDLD, if a fiber breaks in one direction, a fiber port may assume the link is still good (because the other direction is operating normally) and continue to send traffic on the connected ports. UDLD-enabled ports; however, will prevent traffic from being sent across a bad link by blocking the ports in the event that either the individual transmitter or receiver for that connection fails.

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

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

Configuring UDLD

When configuring UDLD, keep the following considerations in mind:

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

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

For detailed information about UDLD, see [“Uni-directional link detection \(UDLD\)” \(page 65\)](#).

Syntax:

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

Enables UDLD on a port or range of ports.

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

Default: UDLD disabled

Syntax:

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

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

Default: 50 (5 seconds)

Syntax:

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

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

Default: 5

Syntax:

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

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

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

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

Enabling UDLD (CLI)

UDLD is enabled on a per-port basis.

Example:

To enable UDLD on port `a1`, enter:

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

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

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

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

Changing the keepalive interval (CLI)

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

Example:

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

```
HP Switch(config)# link-keepalive interval 70
```

Changing the keepalive retries (CLI)

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

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

Example:

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

```
HP Switch(config)# link-keepalive retries 4
```

Configuring UDLD for tagged ports

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

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

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

NOTE:

- You must configure the same VLANs that will be used for UDLD on all devices across the network; otherwise, the UDLD link cannot be maintained.
- If a VLAN ID is not specified, UDLD control packets are sent out of the port as untagged packets.
- To re-assign a VLAN ID, re-enter the command with the new VLAN ID number. The new command overwrites the previous command setting.
- When configuring UDLD for tagged ports, you may receive a warning message if there are any inconsistencies with the VLAN configuration of the port.
See [Table 3 \(page 44\)](#) for potential problems.

Viewing UDLD information (CLI)

Syntax:

```
show link-keepalive
```

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

Syntax:

```
show link-keepalive statistics
```

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

Syntax:

```
clear link-keepalive statistics
```

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

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

Enter the `show link-keepalive` command.

Example:

Figure 9 Example: of `show link-keepalive` command

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

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

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

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

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

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

Port 5 has been disabled by the System Administrator.

Viewing detailed UDLD information for specific ports (CLI)

Enter the `show link-keepalive statistics` command.

Example:

Figure 10 Example: of show link-keepalive statistics command

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

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

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

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

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

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

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

Clearing UDLD statistics (CLI)

Enter the following command:

```
HP Switch# clear link-keepalive statistics
```

This command clears the packets sent, packets received, and transitions counters in the show link keepalive statistics display (see [Figure 10 \(page 69\)](#) for an Example:).

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/PoE+ Planning and Implementation Guide*, which is available on the HP Networking website at www.hp.com/networking. Enter your Switch number.

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

PoE terminology

PoE and PoE+ operate similarly in most cases. Any differences between PoE and PoE+ operation are noted; otherwise, the term "PoE" is used to designate both PoE and PoE+ functionality.

PoE operation

Using the commands described in this chapter, you can:

- Enable or disable PoE operation on individual ports.
- Monitor PoE status and performance.
- 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.
- Specify the port priority you want to use for provisioning PoE power in the event that the PoE resources become oversubscribed.

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

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

Configuration options

In the default configuration, all ports in a HP switch covered in this guide are configured to support PoE operation. You can:

- Disable or re-enable per-port PoE operation on individual ports to help control power usage and avoid oversubscribing PoE resources.
- Configure per-port priority for allocating power in case a PoE device becomes oversubscribed and must drop power for some lower-priority ports to support the demand on other, higher-priority ports.
- 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. Depending on the amount of power the power supply device delivers to a PoE switch, there may or may not always be enough power available to connect and support PoE operation on all the ports. When a new PD connects to a PoE switch and the switch does not have enough power left for that port:

- If the new PD connects to a port “X” having a *higher* PoE priority than another port “Y” that is already supporting another PD, then 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.
- If the new PD connects to a port “X” having a *lower* priority than all other PoE ports currently providing power to PDs, then 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” \(page 71\)](#).)

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

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

Power priority operation

When is power allocation prioritized?

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, then 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 switch’s PoE-enabled ports.

How is power allocation prioritized?

There are two ways that PoE power is prioritized:

- Using a *priority class* method, a power priority of **Low** (the default), **High**, or **Critical** is assigned to each enabled PoE port.
- Using a *port-number priority* method, a lower-numbered port has priority over a higher-numbered port within the same configured priority class. For example, port A1 has priority over port A5 if both are configured with **High** priority.

Configuring PoE operation

In the default configuration, PoE support is enabled on the ports in a PoE 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

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.

Enabling support for pre-standard devices

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

Syntax:

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

Detects and powers pre-802.3af standard devices.

NOTE: The default setting for the `pre-std-detect` PoE parameter changed. In earlier software the default setting is "on". 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 are prioritized by port number in ascending order. For example, if ports A1-A24 have a priority level of critical, port A1 has priority over ports A2-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. PoE priorities

are invoked only when all active PoE ports cannot be provisioned (supplied with PoE power)

| | |
|----------|--|
| Critical | Specifies the highest-priority PoE support for <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. |

Controlling PoE allocation

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.

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.

| | |
|--------------------|---|
| <code>usage</code> | The automatic allocation by a PD |
| <code>class</code> | Uses the power ramp-up signature of the PD to identify which power class the device will be in. Classes and their ranges are shown in Table 6 . |
| <code>value</code> | <i>A user-defined level of PoE power allocated for that port.</i> |

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.

To configure a port by value:

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

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

or in interface context:

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

2. Select a value:

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

or in interface context:

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

To view the settings, enter the `show power-over-ethernet` command, shown in [Example 43](#).

Example 43 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   :
Value            : 15 W 1
Power Class       : 2

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

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

1 Maximum power delivered.

If you set the PoE maximum value to less than what the PD requires, a fault occurs, as shown in [Example 44](#).

Example 44 PoE power value set too low for the PD

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

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

Status and Counters - Port Power Status for port A7

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

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

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

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

Configuring PoE redundancy

When PoE redundancy is enabled, PoE redundancy occurs automatically. The switch keeps track of power use and will 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 if one of the power supplies should fail.

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/PoE+ Planning and Implementation Guide*, available on the HP website at www.hp.com/networking. Auto search the model number for your switch, For example, "HP Switch 2920", then select the device from the list, and click on **Product manuals**. Click on the "Setup and install — general" link under **Manuals**.

Changing the threshold for generating a power notice

You can configure one of the following thresholds:

- A global power threshold that applies to all ports on the switch. This setting acts as a trigger for sending a notice when the PoE power consumption on any PoE port installed in the switch crosses the configured global threshold level. (Crossing the threshold level in either

direction—PoE power usage either increasing or decreasing— triggers the notice.) The default setting is 80%.

- 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.

Syntax:

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

This command 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.

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 destination(s).

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.

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%.
```

For a message listing, please see the *Event Log Message Reference Guide* for your switch. Go to www.hp.com/networking; auto search the model number for your switch, for Example: "HP Switch 2920", then select the device from the list and click on **Product manuals**. Click on the "User guide" link under **Manuals**.

(Default Global PoE Power Threshold: **80**). 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
```

NOTE:

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%.

PoE/PoE+ allocation using LLDP information

LLDP with PoE

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

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

Enabling or disabling ports for allocating power using LLDP

Syntax:

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

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

Default: Disabled

Example:

You can enter this command to enable LLDP detection:

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

or in interface context:

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

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
```

LLDP with PoE+

Overview

The data link layer classification DLC for PoE provides more exact control over the power requirement between a PSE and PD. The DLC works in conjunction with the physical layer classification 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.

The power negotiation between a PSE and a PD can be implemented 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 repeatedly query the PD 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

LLDP can negotiate power with a PD by using LLDP MED TLVs (disabled by default). This 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.

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

Syntax:

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

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

Default: Disabled

Example:

You can enter this command to enable LLDP detection:

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

or in interface context:

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

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

You can view the settings by entering the `show power-over-ethernet brief` command, as shown in [Example 45](#).

Example 45 Port with LLDP configuration information obtained from the device

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

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

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

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

```
Module A Power
```

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

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

Viewing PoE when using LLDP information

Viewing LLDP port configuration

To view information about LLDP port configuration, use the `show lldp config` command.

Syntax:

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

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

Example 46 LLDP port configuration information with PoE

```
HP Switch(config)# show lldp config 4
```

```
LLCP Port Configuration Detail
```

```
Port : 4
```

```
AdminStatus [Tx_Rx] : Tx_Rx
```

```
NotificationsEnabled [False] : False
```

```
Med Topology Trap Enabled [False] : False
```

```
TLVS Advertised:
```

```
* port_descr
```

```
* system_name
```

```
* system_descr
```

```
* system_cap
```

```
* capabilities
```

```
* network_policy
```

```
* location_id
```

```
* poe
```

```
* macphy_config
```

```
* poeplus_config
```

```
IpAddress Advertised:
```

Example 47 shows an Example: of the local device power information using the `show lldp info local-device <port-list>` command.

Example 47 Local power information

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

LLCP Local Port Information Detail

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

Poe Plus Information Detail

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

Example 48 shows the remote device power information using the `show lldp info remote-device <port-list>` command.

Example 48 Remote power information

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

LLCP Remote Device Information Detail

```
Local Port      : A3
ChassisType     : mac-address
ChassisId       : 00 16 35 ff 2d 40
PortType        : local
PortId          : 23
SysName         : HPSwitch
System Descr    : HP Switch 3500-24, revision W.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
```

Operating note

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

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


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

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

Viewing the global PoE power status of the switch

Syntax:

```
show power-over-ethernet [ 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 82) . |
| <code><port-list></code> | Displays PoE information for the ports in port-list. See “Viewing the PoE status on specific ports” (page 84) . |
| <code><slot-id-range></code> | Displays PoE information for the selected slots. See Example 51 (page 84)). Enter the <code>all</code> option to display the PoE information for all slots. |

The `show power-over-ethernet` displays data similar to that shown in [Example 49 \(page 82\)](#).

Example 49 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:

| | |
|------------------------|--|
| PoE Port | Lists all PoE-capable ports on the switch. |
| Power Enable | Shows Yes for ports enabled to support PoE (the default) and No for ports on which PoE is disabled. |
| Power Priority | Lists the power priority (Low , High , and Critical) configured on ports enabled for PoE. (For more information on this topic, see “Configuring PoE operation” (page 72).) |
| Alloc by | Displays how PoE is allocated (usage , class , value). |
| Alloc Power | The maximum amount of PoE power allocated for that port (expressed in watts). Default: 17 watts for PoE; 33 watts for PoE+. |
| Actual Power | The power actually being used on that port. |
| Configured Type | If configured, shows the user-specified identifier for the port. If not configured, this field is empty. |

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

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

Example 50 show power-over-ethernet brief command output

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

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

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

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

```
Module A Power
```

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

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

You can also show the PoE information by slot:

Example 51 Showing the PoE information by slot

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

Status and Counters - System Power Status for slot A

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

Viewing the PoE status on specific ports

Syntax:

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

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

| | |
|-------------------------|---|
| Power Enable | Shows Yes for ports enabled to support PoE (the default) and No for ports on which PoE is disabled. For ports on which power is disabled, this is the only field displayed by <code>show power-over-ethernet <port-list></code> . |
| Priority | Lists the power priority (Low , High , and Critical) configured on ports enabled for PoE. (For more on this topic, see "Configuring PoE operation" (page 72) .) |
| Allocate by | How PoE is allocated (usage , class , value). |
| Detection Status | <ul style="list-style-type: none">• Searching: The port is trying to detect a PD connection.• Delivering: The port is delivering power to a PD.• Disabled: On the indicated port, either PoE support is disabled or PoE power is enabled but the PoE module does not have enough power available to supply the port's power needs.• Fault: The switch detects a problem with the connected PD.• Other Fault: The switch has detected an internal fault that prevents it from supplying power on that port. |
| Over Current Cnt | Shows the number of times a connected PD has attempted to draw more than 15.4 watts for PoE or 24.5 watts for PoE+. Each occurrence generates an Event Log message. |
| Power Denied Cnt | Shows the number of times PDs requesting power on the port have been denied because of insufficient power available. Each occurrence generates an Event Log message. |
| Voltage | The total voltage, in volts, being delivered to PDs. |
| Power | The total power, in watts, being delivered to PDs. |
| LLDP Detect | Port is enabled or disabled for allocating PoE power, based on the link-partner's capabilities via LLDP. |
| Configured Type | If configured, shows the user-specified identifier for the port. If not configured, the field is empty. |
| Value | The maximum amount of PoE power allocated for that port (expressed in watts). Default: 17 watts for PoE; 33 watts for PoE+ |
| Power Class | Shows the power class of the PD detected on the indicated port. Classes include: <ul style="list-style-type: none">0: 0.44 to 12.95 watts1: 0.44 to 3.84 watts2: 3.84 to 6.49 watts3: 6.49 to 12.95 watts4: For PoE+; up to 25.5 watts can be drawn by the PD |

| | |
|-----------------------|---|
| MPS Absent Cnt | Shows the number of times a detected PD has no longer requested power from the port. Each occurrence generates an Event Log message. ("MPS" refers to the "maintenance power signature.") |
| Short Cnt | Shows the number of times the switch provided insufficient current to a connected PD. |
| Current | The total current, in mA, being delivered to PDs. |

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

Example 52 `show power-over-ethernet <port-list> output`

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

Status and Counters - Port Power Status for port A6

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

Status and Counters - Port Power Status for port A7

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

Using the HP 2920 Switch with an external power supply

Overview

The HP 640 Redundant/External Power Supply Shelf (J9805A) is an external shelf that can house up to three power supplies (PSUs). The PSUs installed in the HP 640 RPS/EPS Shelf can supply redundant power to HP 2920 PoE and non-PoE Switches in the event of an HP 2920 Switch internal power supply failure, and can provide additional PoE power to the 2920 PoE Switches. This section discusses the switch CLI `external-power-supply` command options used to configure the HP 2920 Switches for operation with the HP 640 RPS/EPS. The HP 640 RPS/EPS Shelf is also identified as the **XPS**, since that is how it is identified in the HP 2920 Switch software.

For complete information on the HP 640 RPS/EPS Shelf installation, physical setup options, and troubleshooting, see the *HP 640 RPS/EPS Shelf Installation and Power Setup Guide* online at www.hp.com/networking. Auto search on "640", select the device in the list, and click on **Display selected**. Then click on the links that have "manuals" in them to get to the web page that lists the available manuals.


Supported PSUs

The same PSUs can be used in both the HP 2920 Switches and the XPS. Each XPS zone can hold any of the supported PSUs. The supported PSUs are these:

- HP X332 1050W PSU (J9737A) is a 54V power supply unit that can provide 740W of PoE power and a maximum power rating of 1050W (combined system and PoE power).
- HP X332 575W PSU (J9738A) is a 54V power supply unit that can provide 370W of PoE power and a maximum power rating of 575W (combined system and PoE power).
- HP X3312 165W PSU (J9739) is a 12V power supply unit providing non-PoE power. It is not accepted in PoE switches.

Figure 11 shows an Example: of the three PSUs installed in the XPS zones and the power that they provide.

Figure 11 HP 640 RPS/EPS with supported power supplies



In addition to the voltage and power differences between the three PSUs, the non-PoE J9739A PSU has a mechanical key that is different from the PoE PSUs. The mechanical key prevents the insertion of a PoE PSU into a non-PoE switch, or a non-PoE PSU into a PoE switch. This keying function is not needed for the HP 640 RPS/EPS as it can accept all three types of PSUs—PoE and non-PoE.

Using the XPS for additional PoE power

The XPS can be used to provide PoE power to an HP 2920-PoE switch in addition to the power from the switch's internal power supply (IPS). The amount of available external power depends on which external power supplies are installed in the XPS and how the power zones have been configured.

Determining the maximum available PoE power

The information in tables below shows the maximum amount of PoE power that is available for various power supply configurations. It is important to use the information displayed in these tables when determining the power supplied for a configuration, as they accurately represent the maximum power that is available.

Table 7 Maximum PoE power available with 575W PSU in 640 RPS/EPS

| Number of ports enabled in the zone | For 2920 Switch with 575W PSU and 640 RPS/EPS with 575W PSU | | | For 2920 Switch with 1050W PSU and 640 RPS/EPS with 575W PSU | | |
|-------------------------------------|---|---------------------|-----------|--|---------------------|-----------|
| | PoE from 640 RPS/EPS PSU | PoE from switch PSU | Total PoE | PoE from 640 RPS/EPS PSU | PoE from switch PSU | Total PoE |
| 1 port (zones 1,2, or 3) | 370W | 370W | 740W | 0W (not supported) | 740W | 740W |
| 2 ports (zones 1, 2, or 3) | 140W | 370W | 510W | The PSU in the 640 RPS/EPS must have equal to or greater power (Watts) than the PSU in the switch. | 740W | 740W |
| 3 ports (zone 1 only) | 60W | 370W | 430W | | 740W | 740W |
| 4 ports (zone 1 only) | 0W | 370W | 370W | | 740W | 740W |

Table 8 Maximum PoE power available with 1050W PSU in 640 RPS/EPS

| Number of ports enabled in the zone | For 2920 Switch with 575W PSU and 640 RPS/EPS with 1050W PSU | | | For 2920 Switch with 1050W PSU and 640 RPS/EPS with 1050W PSU | | |
|-------------------------------------|--|---------------------|-----------|---|---------------------|-----------|
| | PoE from 640 RPS/EPS PSU | PoE from switch PSU | Total PoE | PoE from 640 RPS/EPS PSU | PoE from switch PSU | Total PoE |
| 1 port (zones 1,2, or 3) | 700W | 370W | 1070W | 700W | 740W | 1440W |
| 2 ports (zones 1, 2, or 3) | 370W | 370W | 740W | 370W | 740W | 1110W |
| 3 ports (zone 1 only) | 130W | 370W | 500W | 130W | 740W | 870W |
| 4 ports (zone 1 only) | 60W | 370W | 430W | 60W | 740W | 800W |

For example, the internal 1050W PSU can supply 740 watts of internal PoE power to the PoE ports. With the addition of an XPS containing a 1050W PSU, an additional 700 watts of external PoE power can be delivered to the PoE ports, for a total of 1440W of PoE power. This is the maximum amount of PoE power that can be supplied to the switch ports (30W per port x 48 ports = 1440W).

As shown in [Table 7](#), though, when a 575W PSU is installed in Zone 1 and all four ports are enabled, there is redundancy protection, but zero watts of external PoE power from the XPS. The following table illustrates three basic setups for HP 2920 Switches and using an HP 640 RPS/EPS for extra PoE power.

Table 9 Example: basic setups for switches using the XPS

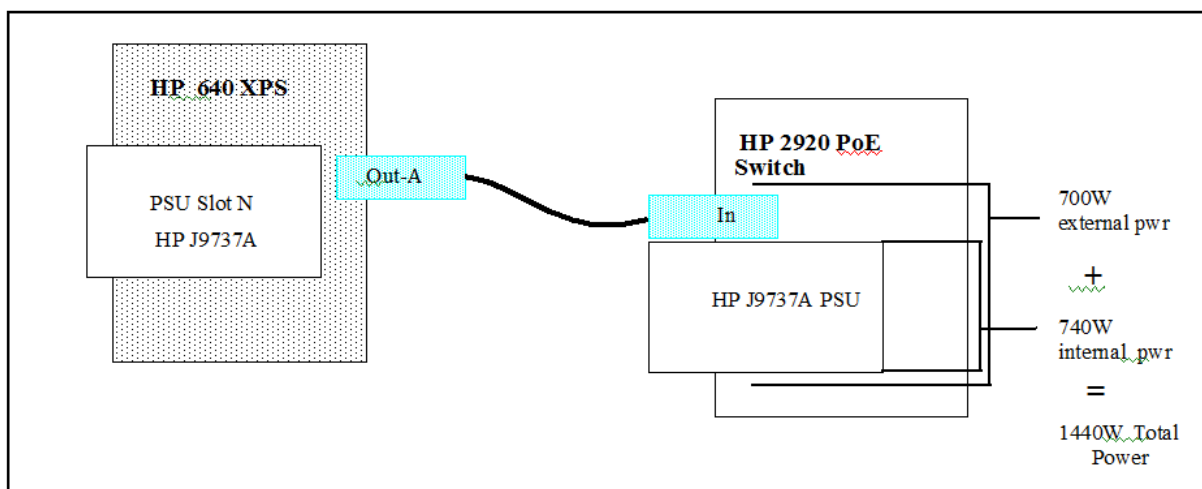
| Power | Total power Available per switch | # of Switches/Zone | Switch PSU Model | RPS/EPS PSU Model | Description |
|------------------|--|--------------------|------------------|-------------------|---|
| 1440W | 1440W/30W for 48 ports | 1 | 1050W | 1050W | 740W internal power and 700W external PoE power |
| 740W, one switch | 30W for 24 ports or 15.4W for 48 ports | 1 | 575W | 575W | 370W internal power and 370W |

Table 9 Example: basic setups for switches using the XPS (continued)

| Power | Total power Available per switch | # of Switches/Zone | Switch PSU Model | RPS/EPS PSU Model | Description |
|----------------------------|---|--------------------|------------------|-------------------|--|
| | | | | | external PoE power |
| 740W, zone with 2 switches | 30W for 24 ports, or 15.4W for 48 ports | 2 | 575W | 1050W | 370W internal power and 370W of external power for each switch |

Figure 12 shows an Example: physical connection between an HP 640 RPS/EPS Shelf with a 1050W PSU installed and an HP 2920 Switch also with a 1050W PSU installed. The resulting PoE Power is indicated.

Figure 12 An external PSU and an internal PSU combined to provide 1440W of total PoE power



For complete information about configuration options, see the *HP 640 RPS/EPS Shelf Installation and Power Setup Guide* online at www.hp.com/networking. Auto search on "640", select the device in the list, and click on **Display selected**. Then click on the links that have "manuals" in them to get to the web page that lists the available manuals.

Operating rules

There will be power flow between the switch and the XPS if these conditions are met:

- PSUs in the HP 2920 switch and the HP 640 XPS are valid, recognized models.
- All PSUs in a zone are mutually compatible.
- The zone configuration for all the switches in that zone is supported.
- The power of the PSU in the XPS must always be equal to or larger than the power of the internal power supply in any switch in the same zone.

A 575W PSU installed in an HP 640 XPS zone cannot be used to provide power to any switch that contains a 1050W PSU. An error message displays and there is no flow of power. You must use a 1050W PSU in the HP 640 XPS to supply the PoE power for that zone.

It is OK to have a 1050W PSU in the XPS and a 575W PSU in the connected switch.

- Any necessary reduction of external PoE power in any switch is performed in an orderly and configured way before another switch is added.

NOTE: By default, the HP 2920 Switch ports have PoE power priority starting with the lowest numbered port. That is, port one has the highest PoE priority.

You should always connect PoE devices that have the highest requirement for uninterrupted PoE power to the lowest numbered HP 2920 ports.

Using redundant (N+1) power

Using the XPS as a redundant power supply provides N+1 redundancy to the first switch that fails in the zone. For example, if two HP 2920 switches are connected to the same zone and the PSU in the first switch fails, the XPS then provides 12V power to that switch to keep it operating. The 12V power from the XPS to the second switch is disabled, and that switch continues to operate under power from its own IPS, but it no longer has N+1 redundancy.

The XPS continues to provide PoE power to both switches but the total PoE power for both switches is reduced. As a result, some of the PoE devices connected to both switches might lose power, depending on how many devices are connected and how much PoE power they are using.

NOTE: Only one PSU failure is supported. Multiple failures are not supported.

Providing non-PoE redundant power

The HP 165W non-PoE PSU (J9739A) is the only power supply used with non-PoE HP 2920 switches. If a J9739A PSU is installed in an XPS zone, then only HP 2920 non-PoE Switches can be connected to that zone. The XPS provides redundant (N+1) power if the power supply in one of the non-PoE switches fails. The power flows to the switch with the failed power supply. For any other switches connected to that zone, the power flow is disabled until the failed PSU is replaced. The switches continue to operate without interruption and continue to communicate with the XPS.

If the HP J9739A power supply that fails is installed in the XPS, all power flow to all switches connected to that zone is disabled until the failed PSU is replaced. The connected switches continue to operate without interruption, but communication with the XPS may stop, depending on the severity of the failure.

For more information about supported configurations and redundant behavior, see the *HP 640 RPS/EPS Shelf Installation and Power Setup Guide* online at www.hp.com/networking. Auto search on "640", select the device in the list, and click on **Display selected**. Then click on the links that have "manuals" in them to get to the web page that lists the available manuals.

Configuring the HP 2920 PoE switches to use the XPS

To configure the HP 2920 PoE Switches to use the PoE power from the XPS, you will issue `external-power-supply` commands to the switches. By default, all the available PoE power is shared equally by all the switches connected to a given XPS zone. To cause a redistribution of this power, you must issue the `external-power-supply` commands to *all* of the switches that are connected to that zone.

NOTE: Configuring HP 2920 Switches that are Members of a 2920 Stack. If the `external-power-supply` commands that are used to configure the HP 2920 Switches for how they interact with the XPS, there is a `member-id` option. When the switch is a member of a 2920 Stack of switches, the `member-id` must be included in order to configure the XPS Shelf zone to which this member switch is connected.

Enabling and disabling power from the XPS

In the XPS default configuration, the switch automatically receives backup power and external PoE power (for PoE switches that require it) when an HP 2920 Switch is connected to an XPS port. Auto-recovery is also enabled in the default configuration. The following command lets you control whether the switch receives the XPS power, if you need to do so.

Syntax:

```
external-power-supply [member <member-id>] <enable | disable>
```

Permits power to be supplied or discontinued from the XPS to the switch or to a member of a stack of switches.

| | |
|---------|---|
| enable | Turns on the XPS port to provide power to the switch. When the switch is connected to an XPS port, it automatically receives backup power and PoE power. The XPS is enabled by default with the auto-recovery feature. NOTE: If the <code>external-power-supply disable</code> command is executed, auto-recovery is disabled and you must execute the <code>external-power-supply auto-recovery enable</code> command to re-enable auto-recovery. Executing the <code>external-power-supply enable</code> command does not re-enable auto-recovery. |
| disable | Turns off the XPS port. Auto-recovery on the switch is turned off as well. The disable option can be used to turn off the XPS even if the cable is part of the current distribution map. This can be useful for troubleshooting. |

Example 53 Enabling and disabling the power for a specified 2920 stack member

```
HP Switch(config)# external-power-supply member 1 enable
```

```
HP Switch(config)# external-power-supply member 1 disable
```

This will stop the power supply to this member and disable auto recovery.

```
Continue (y/n)? y
```

Configuring auto-recovery

Syntax:

```
external-power-supply [member <member-id>] auto-recovery  
<enable | disable>
```

When enabled, the auto-recovery feature allows the switch to configure itself if an internal PSU or an external PSU has a power failure and is replaced. The switch begins to receive backup power.

| | |
|---------|--|
| enable | When the switch is connected to an XPS port and the port is part of the distribution map, the XPS can provide redundant or external power. Default: Enabled |
| disable | When auto-recovery is disabled, the switch must be reconfigured to obtain backup power in case of a power supply failure or the hotswap of the XPS cable. |

Example 54 Disabling auto-recovery for a specified 2920 stack member

```
HP Switch(config)# external-power-supply member 1 auto-recovery disable
This will disable the auto recovery feature enabled on this member. External
power supply needs to be re-enabled in case of power supply failure or hot swap
of power supply cable or change in distribution map with 'force'.
Continue (y/n)? y
```

Example 55 Disabling auto-recovery for the switch

```
HP Switch(config)# external-power-supply auto-recovery disable
This will disable the auto recovery feature enabled on this switch. External
power supply needs to be re-enabled in case of power supply failure or hot swap
of power supply cable or change in distribution map with 'force'.
Continue (y/n)? y
```

Restoring the default external power supply settings

Syntax:

```
external-power-supply [member <member-id>]
reset
```

Restores the XPS configuration on the current zone to factory default settings. This may power down some PoE ports.

Default: All XPS ports are operational.

For a stack of switches, the zone connected to the specified member is reset to its factory default configurations. Specify the member-id to configure the zone to which the member is connected.

NOTE: This command is not available in stacking member context.

Example 56 Restoring the default external power supply settings

```
HP Switch(config)# external-power-supply reset
```

This will reset the external power supply to factory default configurations. This might shutdown powered PoE ports on the connected switches.

```
Continue (y/n)? y
Configuring external power supply, this might take up to a minute...
```

Distributing power to specified ports

Syntax:

```
external-power-supply [member <member-id>]
power-share <xps ports> [force]
```

Configures the XPS to distribute power to the ports specified. The amount of XPS power received by each XPS port depends on the number of ports that have been specified.

NOTE: This command is not available in stacking member context.

| | |
|---------|--|
| [force] | When the <code>force</code> option is selected, the zone can be re-configured. This allows an additional switch to be added to the existing setup. External PoE power is distributed to the newly added switch, however, this will result in the temporary shutdown of all PoE devices connected to PoE ports on the affected switches that are receiving their PoE power from the XPS. PoE devices that are receiving their PoE power from the switch's IPS will continue operation. For more information, see "Example: of using the force option" (page 92) . |
|---------|--|

Example: of the power-share option

This Example: is for a configuration with a distribution map of 1A, 1B, 1C, and 1D, but you want power from the XPS to be available to only ports 1A and 1C.

Example 57 Configuration for power allocation

```
HP Switch(config)# external-power-supply power-share 1A, 1C
```

This would change the allocated power for XPS port 1A, 1C to 370W, disable XPS ports 1B, 1D and change their allocated power to 0W. This might cause PoE power ports connected in system 1B, 1D to be shut down.

```
Continue (y/n)? y
```

```
Configuring external power supply, this might take up to a minute...
```

Example: of adding a switch

This Example: illustrates adding a new switch to 1D with a current distribution map of 1A, 1C.

Example 58 Configuring power allocation when adding a switch

```
HP Switch-1A(config)# external-power-supply power-share allow 3
```

```
HP Switch-1C(config)# external-power-supply power-share allow 3
```

```
HP Switch(config)# external-power-supply power-share 1A,1C,1D
```

This would change allocated power for XPS port 1A, 1C, 1D to 130W.

```
Continue (y/n)? y
```

```
Configuring external power supply, this might take up to a minute...
```

Example: of using the force option

The `force` option allows you to force an immediate change to the PoE power distribution for a specified XPS zone. The `force` option can be convenient in that it needs to be issued to only the switch that is being added to the zone. For the “graceful” method of power redistribution, using the `external-power-supply allow` and `external-power-supply <xps ports>` command sequence, you must issue these commands to all of the affected switches in the zone. But, using the `force` option has consequences in the PoE power delivery to the affected switches. See the important note below.

- ❗ **IMPORTANT:** Using the `force` option causes all PoE power coming from the XPS to be temporarily discontinued while the XPS and connected switches negotiate the new power configuration. PoE PDs connected to lower-priority PoE ports on the affected switches, and which are getting their power from the XPS, will lose power. Only the PDs that are connected to higher-priority PoE ports, up to the PoE capacity of the switch’s IPS, will retain their power. The lowered number switch ports have a higher PoE priority.

HP recommends that you should not use the `force` option at times when PoE power to the PDs must be maintained. Use the `external-power-supply allow` and `external-power-supply <xps ports>` command sequence instead, which causes a more controlled redistribution of the power.

This Example: uses the `force` option to change the power allocation.

Example 59 Non-graceful method for adding a switch and distributing external power

```
HP Switch(config)# external-power-supply power-share 1D force
This would change allocated power for XPS port 1D to 370W,
disable XPS ports 1A, 1B, 1C and change their allocated power to 0W.
This might result in PoE powered ports connected in system 1A, 1B, 1C, 1D to be
shutdown.
Continue (y/n)? y
Configuring external power supply, this might take up to a minute...
```

Reducing allocated external power

Syntax:

```
external-power-supply [member <member-id>]
power-share allow <num-of-switches-in-zone>
```

Provides a graceful way to reduce the allocated external power when a switch is added to an existing XPS setup.

NOTE: This command is not available in stacking member context.

This command is executed when a new switch is connected to an existing XPS setup. To distribute power to the newly added switch, execute this command on each connected switch to reduce the allocated power so that the new switch can draw power from the XPS. This may cause some PoE devices connected to the switches to be powered down because the total PoE power going to each switch will be reduced.

For example, if a switch is connected to an XPS zone with one other switch, and a third switch is added to that zone (must be Zone 1 which is the only zone with more than two ports), then the following command should also be executed on all the switches connected to this zone.

Example 60 Reallocating external power

```
HP Switch(config)# external-power-supply power-share allow 3
```

This would change allocated power for current system from 370W to 130W.
The following PoE powered ports 1,2 4-7 would be shutdown.

```
Continue (y/n)? y
```

Example: configurations

Non-PoE configuration

If the non-PoE switch and the XPS are in their default configurations, run the `show external-power-supply brief` command to verify that there is adequate XPS power to provide redundancy power to the switch.

If the non-PoE switch has auto-recovery disabled and the XPS is not providing redundancy support to the switch, execute the commands as shown in [Example 61](#).

Example 61 Enabling an XPS for a non-PoE switch configuration

```
HP Switch(config)# external-power-supply enable

HP Switch(config)# show external-power-supply brief
External Power Supply Type       : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number : CN36FX201L
External Power Supply Module      : J9805A
External Power Supply PSU Revision : 0
External Power Supply PSU Module  : J9739A
Voltage / Wattage                 : 12V / 165W
Current Zone                       : 2
Zone State                         : Powered
Zone Record Version               : 3

Cable Port  Connection      XPS      Mbr System Name
Id    Allow Status          Enabled Id
-----
2A*   Yes   Available      Yes     -   HP-2920-24G
2B    Yes   Not Connected
```

If you want to enable auto-recovery as well, execute the `external-power-supply auto-recovery enable` command.

Example 62 Enabling an XPS and auto-recovery for a non-PoE switch configuration

```
HP Switch(config)# external-power-supply auto-recovery enable

HP Switch(config)# show external-power-supply detail

External Power Supply Type       : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number : CN36FX201L
External Power Supply Module      : J9805A
External Power Supply PSU Revision : 0
External Power Supply PSU Module  : J9739A
Voltage / Wattage                 : 12V / 165W
Current Zone                       : 1
Zone State                         : Powered
Zone Record Version               : 3

Cable ID : 1A

System Name           : HP 2920-24G-PoE+ Switch
Stack Id              : 00010021-f73bdd81
Member Id             : 1
Module                : J9727A
MAC Address           : 0021f7-78d6d0
Software Version      : WB.15.13.0000x
Serial Number         : SG2ZFLX098
Internal Power Supply Rating : 12V / 165W
External Power        : 0 W
Connection Status     : Available
Auto Recovery         : Yes
Cable Record Version  : 3
Supported Zone Record Version: 3
```

PoE configuration for full PoE power to one XPS port

[Example 63](#) shows the configuration for an HP 2920 switch with a 1050W IPS, and an XPS with a 1050W PSU. Execute the `show external-power-supply brief` command to view the current status of the power distribution. The output shows that the XPS is providing 60W of external PoE power to each XPS port and the port's connected switch.

Example 63 Distribution of PoE power

```
HP Switch(config)# show external-power-supply brief
```

```
External Power Supply Type       : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number : CN36FX201L
External Power Supply Module      : J9805A
External Power Supply PSU Revision : 0
External Power Supply PSU Module  : J9737A
Voltage / Wattage                 : 54V / 1050W
Current Zone                       : 1
Zone State                         : Powered
Zone Record Version               : 3
```

| Cable Id | Port Allow | Connection Status | XPS Enabled | Ext. Power | Mbr Id | System Name |
|----------|------------|-------------------|-------------|------------|--------|------------------|
| 1A* | Yes | Available | Yes | 60 W | - | HP-2920-48G-POE+ |
| 1B | Yes | Available | Yes | 60 W | - | HP-2920-48G-POE+ |
| 1C | Yes | Available | Yes | 60 W | - | HP-2920-24G-PoEP |
| 1D | Yes | Available | Yes | 60 W | - | HP-2920-24G-PoEP |

As shown in [Example 64](#), executing the `power-share` command to cause all power to be distributed to port 1A changes the allocated power 700W for that port. XPS ports 1B, 1C, and 1D are disabled and the allocated power for each is now zero watts.

Example 64 Distribution of PoE power after redistribution

```
HP Switch(config)# external-power-supply power-share 1A
This would change allocated power for XPS port 1A to 700W,
disable XPS ports 1B, 1C, 1D and change their allocated power to 0W.
This might result PoE powered ports connected in system 1B, 1C, 1D to be
shutdown.
Continue (y/n)? y
Configuring external power supply, this might take up to a minute...
```

```
HP Switch(config)# show external-power-supply brief
```

```
External Power Supply Type      : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number : CN36FX201L
External Power Supply Module     : J9805A
External Power Supply PSU Revision : 0
External Power Supply PSU Module  : J9737A
Voltage / Wattage               : 54V / 1050W
Current Zone                    : 1
Zone State                      : Powered
Zone Record Version             : 3
```

| Cable Id | Port Allow | Connection Status | XPS Enabled | Ext. Power | Mbr Id | System Name |
|----------|------------|-------------------|-------------|------------|--------|------------------|
| 1A* | Yes | Available | Yes | 700 W | - | HP-2920-48G-POE+ |
| 1B | No | Unavailable | No | 0 W | - | HP-2920-48G-POE+ |
| 1C | No | Unavailable | No | 0 W | - | HP-2920-24G-PoEP |
| 1D | No | Unavailable | No | 0 W | - | HP-2920-24G-PoEP |

Example 65 Output displaying PoE power available

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

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

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

```
Chassis power-over-ethernet:
```

```
Total Available Power   : 1440 W
Total Failover Power    : 740 W
Total Redundancy Power   : 0 W
Total Used Power        : 0 W +/- 6W
Total Remaining Power   : 1440 W
```

```
Internal Power
  1 740W/POE+ /Connected.
External Power
  EPS1 700W /Connected.
```

PoE configuration for multiple switches

Before configuring external PoE power for multiple HP 2920 switches, execute the `show external-power-supply brief` command to determine the current XPS configuration. [Example 66](#) shows an XPS with a 1050W PSU (J9737A) in zone 1 connected to four HP 2920 switches. Assume, for this Example: that each of the switches also contains a 1050W (J9737A) IPS. The figure shows a current configuration with 700W of PoE power being delivered only to XPS port 1A.

Example 66 XPS PoE power delivered to a single switch

```
HP Switch(config)# show external-power-supply brief
```

```
External Power Supply Type       : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number : CN36FX201L
External Power Supply Module      : J9805A
External Power Supply PSU Revision : 0
External Power Supply PSU Module  : J9737A
Voltage / Wattage                 : 54V / 1050W
Current Zone                       : 1
Zone State                         : Powered
Zone Record Version               : 3
```

| Cable Id | Port Allow | Connection Status | XPS Enabled | Ext. Power | Mbr Id | System Name |
|----------|------------|-------------------|-------------|------------|--------|------------------|
| 1A* | Yes | Available | Yes | 700 W | - | HP-2920-48G-POE+ |
| 1B | No | Unavailable | No | 0 W | - | HP-2920-48G-POE+ |
| 1C | No | Unavailable | No | 0 W | - | HP-2920-24G-PoEP |
| 1D | No | Unavailable | No | 0 W | - | HP-2920-24G-PoEP |

To change the power distribution to deliver power to all four XPS ports and their connected switches, execute the `external-power-supply power-share` commands as shown in [Example 67 “Distributing XPS PoE power to multiple switches”](#). Note that the `allow 4` command must be executed on all three of the switches that are currently sharing the power – the switches connected to ports 1A, 1B, and 1C. Then, the command to specify the new distribution map is executed on the switch that is being added – the switch connected to port 1D.

After executing those commands, the `show external-power-supply brief` command now displays 60W of PoE power being delivered to all four XPS ports and their connected switches.

Example 67 Distributing XPS PoE power to multiple switches

```
HP Switch-1A(config)# external-power-supply power-share allow 4
HP Switch-1B(config)# external-power-supply power-share allow 4
HP Switch-1C(config)# external-power-supply power-share allow 4
HP Switch-1D(config)# external-power-supply power-share 1A,1B,1C,1D
This would change allocated power for XPS port 1A,1B,1C,1D to 60W.
Continue (y/n) y
Configuring external power supply, this might take up to a minute...

HP Switch(config)# show external-power-supply brief

External Power Supply Type           : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number  : CN36FX201L
External Power Supply Module         : J9805A
External Power Supply PSU Revision   : 0
External Power Supply PSU Module     : J9737A
Voltage / Wattage                    : 54V / 1050W
Current Zone                          : 1
Zone State                            : Powered
Zone Record Version                  : 3

Cable Port  Connection  XPS      Ext.    Mbr System Name
Id    Allow Status      Enabled Power  Id
-----
1A*   Yes   Available  Yes     60 W   -   HP-2920-48G-POE+
1B    Yes   Available  Yes     60 W   -   HP-2920-48G-POE+
1C    Yes   Available  Yes     60 W   -   HP-2920-24G-PoEP
1D    Yes   Available  Yes     60 W   -   HP-2920-24G-PoEP
```

NOTE: As shown in [Example 68](#), the same results could be accomplished by using a single command issued to the switch connected to port 1D, and by using the force option. As noted in the message provided by the switch software though, PoE power that is being provided to any of the XPS ports might be temporarily shut down while the new power distribution is activated. Port 1A was the only port receiving power, so it is the only one listed: This might result in PoE powered ports connected in system 1A to be shutdown.

Example 68 Distributing XPS PoE power to multiple switches using the force option

```
HP Switch(config)# external-power-supply power-share 1A,1B,1C,1D force
This would change allocated power for XPS port 1A,1B,1C,1D to 60W.
This might result in PoE powered ports connected in system 1A to be shutdown.
Continue (y/n) y
```

For more information, see [“Example: of using the force option”](#) (page 92).

Viewing power information

Syntax:

```
show external-power-supply [member <member-id>] <brief |
detail | info>
```

Displays information about the XPS operational and configuration parameters.

If the switch is a member of a stack of switches, the member-id must be specified to obtain information about the zone to which the member is connected. In the output, an asterisk (“*”) next to the cable ID denotes the current member from which the command is executed.

NOTE: This command is not available in stacking member context.

| | |
|--------|---|
| brief | Displays brief information about the XPS operational and configuration parameters. |
| detail | Displays detailed information about the XPS operational and configuration parameters. |
| info | Displays the power received per switch based on the number of switches connected to the zone. |

XPS parameter information includes:

- **External Power Supply PSU Revision:** The current revision of the PSU.
- **Voltage/Wattage:** The total voltage and wattage available with that PSU.
- **Current Zone:** The zone where this switch is connected.
- **Zone State:** Powered or not powered.
- **Zone Record Version:** The current version of the zone record.
- **Cable Id:** The XPS port designation.
- **Connection Status:** The connection is available, unavailable, not connected, or mismatched. Mismatched connections occur when the PSU is not supported in that configuration.
- **XPS Enabled:** The XPS port is enabled or disabled for power delivery.
- **Ext. Power:** The amount of external power that is allocated, in watts.
- **Mbr Id:** The number of the switch member in the stack.
- **System Name:** The system name of the switch or switch stack.
- **Auto Recovery:** Yes for enabled, no for disabled.

The amount of power received by a port is determined by the distribution map and the type of power supplies used.

Examples for show external-power-supply

Example 69 Output when 3 PoE switches are connected to an EPS/RPS 640 power supply

```
HP Switch(config)# show external-power-supply member 1 brief
```

```
External Power Supply Type       : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number : CN36FX202L
External Power Supply Module      : J9805A
External Power Supply PSU Revision : 1
External Power Supply PSU Module  : J9738A
Voltage / Wattage                 : 54V / 575W
Current Zone                       : 1
Zone State                         : Powered
Zone Record Version               : 3
```

| Cable Id | Port Allow | Connection Status | XPS Enabled | Ext. Power | Mbr Id | System Name |
|----------|------------|-------------------|-------------|------------|--------|------------------|
| 1A* | Yes | Available | Yes | 0 W | 1 | 2-mbr-stack |
| 1B | Yes | Available | Yes | 0 W | 2 | 2-mbr-stack |
| 1C | Yes | Not Connected | | | | |
| 1D | Yes | Available | Yes | 0 W | - | HP-2920-48G-POE+ |

The asterisk beside the cable ID, For example, 1A*, indicates the switch that is communicating with the XPS for information. [Example 69](#) indicates that the switch connected to XPS port 1A is communicating with the XPS. For a stack of switches, all XPS ports in the same stack will display the asterisk beside the cable ID.

Example 70 Output for a 4-member stack of switches when no member is specified

```
HP Switch(config)# show external-power-supply brief
```

```
External power supply information for members 1,2,3,4
```

```
External Power Supply Type       : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number : CN2ZFX2027
External Power Supply Module      : J9805A
External Power Supply PSU Revision : 0
External Power Supply PSU Module  : J9737A
Voltage / Wattage                 : 54V / 1050W
Current Zone                       : 1
Zone State                         : Powered
Zone Record Version               : 3
```

| Cable Id | Port Allow | Connection Status | XPS Enabled | Ext. Power | Mbr Id | System Name |
|----------|------------|-------------------|-------------|------------|--------|---------------|
| 1A* | Yes | Available | Yes | 60 W | 4 | HP-Stack-2920 |
| 1B* | Yes | Available | Yes | 60 W | 3 | HP-Stack-2920 |
| 1C* | Yes | Available | Yes | 60 W | 2 | HP-Stack-2920 |
| 1D* | Yes | Available | Yes | 60 W | 1 | HP-Stack-2920 |

The output varies depending on the switch from which the command is executed. An asterisk next to the port ID indicates where the command was executed. In [Example 71](#) the command is executed from a non-Stack PoE switch connected to XPS port 1C in a PoE zone. In [Example 72](#) the command is executed from a non-PoE switch connected to XPS port 1B in a PoE zone.

Example 71 Output when command is executed from PoE switch 1C connected to a PoE zone

```
HP Switch(config)# show external-power-supply brief

External Power Supply Type      : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number : CN36FX201L
External Power Supply Module     : J9805A
External Power Supply PSU Revision : 0
External Power Supply PSU Module  : J9737A
Voltage / Wattage               : 54V / 1050W
Current Zone                    : 1
Zone State                      : Powered
Zone Record Version             : 3

Cable Port  Connection      XPS      Ext.    Mbr System Name
Id          Allow Status    Enabled Power  Id
-----
1A         Yes   Available   Yes     60 W   -   HP-2920-48G-POE+
1B         Yes   Unavailable No      0 W   -   HP-2920-24G
1C*        Yes   Available   Yes     60 W   -   HP-2920-24G-PoEP
1D         Yes   Available   Yes     60 W   -   HP-2920-24G-PoEP
```

Example 72 Output when command is executed from non-PoE switch 1B connected to a PoE zone

```
HP Switch(config)# show external-power-supply brief

External Power Supply Type      : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number : CN36FX201L
External Power Supply Module     : J9805A
External Power Supply PSU Revision : 0
External Power Supply PSU Module  : J9737A
Voltage / Wattage               : 54V / 1050W
Current Zone                    : 1
Zone State                      : Powered
Zone Record Version             : 3

Cable Port  Connection      XPS      Ext.    Mbr System Name
Id          Allow Status    Enabled Power  Id
-----
1A         Yes   Available   Yes     60 W   -   HP-2920-48G-POE+
1B*        Yes   Mismatch   No      0 W   -   HP-2920-24G
1C         Yes   Available   Yes     60 W   -   HP-2920-24G-PoEP
1D         Yes   Available   Yes     60 W   -   HP-2920-24G-PoEP
```

Example 73 Output for info option with the 575W PSU (J9738A) installed in zone 1

```
HP Switch(config)# show external-power-supply info

External Power Supply Type      : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number : CN36FX202L
External Power Supply Module     : J9805A
External Power Supply PSU Revision : 1
External Power Supply PSU Module  : J9738A
Voltage / Wattage               : 54V / 575W
Current Zone                    : 1
Zone State                      : Powered
Zone Record Version             : 3

Number of Switches      Power Received
Connected               Per Switch
-----
1                       370 W
2                       140 W
3                       60 W
4                       0 W
```

Example 74 Truncated output for detail option with a 575W PSU (J9738A) installed in zone 1 with 3 HP 2920 Switches configured

HP Switch(config)# show external-power-supply detail

```
External Power Supply Type      : HP 640 Redundant/External PS Shelf
External Power Supply Serial Number : CN36FX202L
External Power Supply Module     : J9805A
External Power Supply PSU Revision : 1
External Power Supply PSU Module  : J9738A
Voltage / Wattage               : 54V / 575W
Current Zone                     : 1
Zone State                       : Powered
Zone Record Version              : 3
```

Cable ID : 1A

```
System Name      : HP 2920-24G-PoE+ Switch
Stack Id         : 00010021-f73bdd81
Member Id        : 1
Module           : J9727A
MAC Address      : 0021f7-78d6d0
Software Version : WB.15.13.0000x
Serial Number    : SG2ZFLX098
Internal Power Supply Rating : 54V / 575W
External Power   : 0 W
Connection Status : Available
Auto Recovery    : Yes
Cable Record Version : 3
Supported Zone Record Version: 3
```

Cable ID : 1B

```
System Name      : HP 2920-24G-PoE+ Switch
Stack Id         : 00010021-f73bdd81
Member Id        : 2
Module           : J9727A
MAC Address      : 0021f7-78c6c1
Software Version : WB.15.13.0000x
Serial Number    : SG2ZFLX099
Internal Power Supply Rating : 54V / 575W
External Power   : 0 W
Connection Status : Available
Auto Recovery    : Yes
Cable Record Version : 3
Supported Zone Record Version: 3
```

.
.
.

Examples for show power-over-ethernet commands

Example 75 Output showing both internal and external power supplies connected

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

Status and Counters - System Power Status for member 1

  System Power Status      : Full redundancy
  PoE Power Status        : No redundancy

Chassis power-over-ethernet:

  Total Available Power   : 740 W
  Total Failover Power    : 370 W
  Total Redundancy Power  : 0 W
  Total used Power        : 0 W +/- 6W
  Total Remaining Power   : 740 W

Internal Power
  1 370W/POE+ /Connected.
External Power
  EPS1 370W/POE+ /Connected.
```

Example 76 Output showing failed internal power supply

```
HP Switch# show power-over-ethernet

Status and Counters - System Power Status for member 1

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

Chassis power-over-ethernet:

  Total Available Power   : 370 W
  Total Failover Power    : 0 W
  Total Redundancy Power  : 0 W
  Total used Power        : 0 W +/- 6W
  Total Remaining Power   : 370 W

Internal Power
  1 0W/POE+ /Connected - Faulted.
External Power
  EPS1 370W/POE+ /Connected.
```

Example 77 Output for show power-over-ethernet brief command

HP Switch# show power-over-ethernet brief

Status and Counters - Port Power Status

System Power Status : Full redundancy
PoE Power Status : No redundancy

Available: 1440 W Used: 1439 W Remaining: 1 W

Module 1-48 Power

Available: 1440 W Used: 1439 W Remaining: 1 W

| PoE Port | Power Enable | Power Priority | Alloc By | Alloc Power | Actual Power | Configured Type | Detection Status | Power Class | Pre-std Detect |
|----------|--------------|----------------|----------|-------------|--------------|-----------------|------------------|-------------|----------------|
| 1 | Yes | low | usage | 17 W | 31.9 W | | Delivering | 4 | off |
| 2 | Yes | low | usage | 17 W | 32.3 W | | Delivering | 4 | off |
| 3 | Yes | low | usage | 17 W | 32.3 W | | Delivering | 4 | off |
| 4 | Yes | low | usage | 17 W | 32.3 W | | Delivering | 4 | off |
| | | | | | | | | | |
| 44 | Yes | low | usage | 17 W | 31.7 W | | Delivering* | 4 | off |
| 45 | Yes | low | usage | 17 W | 32.5 W | | Delivering* | 4 | off |
| 46 | Yes | low | usage | 17 W | 0.0 W | | Disabled | 4 | off |
| 47 | Yes | low | usage | 17 W | 0.0 W | | Disabled | 4 | off |
| 48 | Yes | low | usage | 17 W | 0.0 W | | Disabled | 4 | off |

Delivering* - Ports not backed up in the event of Power Supply Failure

Example: for show running-config command

Example 78 Output of running-config file for stack member 1 with auto-recovery disabled

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

Running configuration:

; J9727A Configuration Editor; Created on release #WB.15.13.0000x
; Ver #04:e3.ff.35.0d:20

stacking
 member 1 Type "J9587A" mac-address bb99cc-554433
 exit
hostname "HP-2920-24G-PoEP"
module 1 type j9727a
ip access-list extended "aaa"
 exit
ipv6 ra-guard ports 6
interface 10
 lACP active
 exit
interface 11
 lACP active
 exit
snmp-server community "public" unrestricted
oobm
 ip address dhcp-bootp
 exit
vlan 1
 name "DEFAULT_VLAN"
 no untagged 5-6
 untagged 1-4,7-24,A1-A2,B1-B2
 ip address dhcp-bootp
 exit
vlan 2
 name "VLAN2"
 untagged 5-6
 no ip address
 ipv6 enable
 ipv6 mld enable
 exit
external-power-supply member 1 auto disable
```

Planning and implementing a PoE configuration

This section provides an overview of some considerations for planning a PoE application. For additional information on this topic, refer to the *HP PoE/PoE+ Planning and Implementation Guide* which is available on the HP Networking web site at www.hp.com/networking.

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

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

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

Power requirements

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

By connecting an external power supply you can optionally provision more PoE wattage per port and or supply the switch with redundant 12V power to operate should an internal power supply fail.

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

Assigning PoE ports to VLANs

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

Applying security features to PoE configurations

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

Assigning priority policies to PoE traffic

You can use the configurable QoS (Quality of Service) features in the switch to create prioritization policies for traffic moving through PoE ports. [Table 10](#) lists the available classifiers and their order of precedence.

Table 10 Classifiers for prioritizing outbound packets

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

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

PoE Event Log messages


Please see the *Event Log Message Reference Guide* for information about Event Log messages. To see these manuals, go to www.hp.com/networking. Auto search the model number for your switch, for Example: “HP Switch 2920”, then select the device from the list and click on **Product manuals**. Click on the “User guide” link under **Manuals**.

4 Port Trunking

Overview of port trunking

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

Figure 13 Conceptual Example: of port trunking



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

Port connections and configuration

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

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

NOTE:

Link connections

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

Port security restriction

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

Port trunk features and operation

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

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

Up to 60 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” \(page 118\)](#).)

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.

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.

Syntax:

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

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

Example 79 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
```

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

Example 80 An 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
```

| 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 | A5 | Up | No | Success | 250 | 250 |

Static trunk

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

Table 11 Trunk types used in static and dynamic trunk groups

| Trunking method | LACP | Trunk |
|-----------------|------|-------|
| Dynamic | Yes | No |
| Static | Yes | Yes |

Table 12 describes the trunking options for LACP and Trunk protocols.

Table 12 Trunk configuration protocols

| Protocol | Trunking Options |
|----------------------|---|
| LACP (802.3ad) | <p>Provides dynamic and static LACP trunking options.</p> <ul style="list-style-type: none"> • Dynamic LACP — Use the switch-negotiated dynamic LACP trunk when: <ul style="list-style-type: none"> • The port on the other end of the trunk link is configured for Active or Passive LACP. • You want fault-tolerance for high-availability applications. If you use an eight-link trunk, you can also configure one or more additional links to operate as standby links that will activate only if another active link goes down. • Static LACP — Use the manually configured static LACP trunk when: <ul style="list-style-type: none"> • The port on the other end of the trunk link is configured for a static LACP trunk. • You want to configure non-default spanning tree or IGMP parameters on an LACP trunk group. • You want an LACP trunk group to operate in a VLAN other than the default VLAN and GVRP is disabled. (See “VLANs and dynamic LACP” (page 121).) • You want to use a monitor port on the switch to monitor an LACP trunk. <p>For more information, see “Trunk group operation using LACP” (page 118).</p> |
| Trunk (non-protocol) | <p>Provides manually configured, static-only trunking to:</p> <ul style="list-style-type: none"> • Most HP Switch and routing switches not running the 802.3ad LACP protocol. • Windows NT and HP-UX workstations and servers <p>Use the Trunk option when:</p> <ul style="list-style-type: none"> • The device to which you want to create a trunk link is using a non-802.3ad trunking protocol. • You are unsure which type of trunk to use, or the device to which you want to create a trunk link is using an unknown trunking protocol. • You want to use a monitor port on the switch to monitor traffic on a trunk. <p>See “Trunk group operation using the “trunk” option” (page 123).</p> |

Table 13 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. |

Table 13 General operating rules for port trunks (continued)

| Example 81 Recommended port mode setting for LACP | | | | | | | | | | | | | | | | | | | |
|--|--|---------|------|-----------|------|-----------|-----|---|----------|-----|------|--------|------|---|----------|-----|------|--------|-----|
| | <pre>HP Switch(config)# show interfaces config</pre> <p>Port Settings</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Port</th> <th style="text-align: left;">Type</th> <th style="text-align: left;">Enabled</th> <th style="text-align: left;">Mode</th> <th style="text-align: left;">Flow Ctrl</th> <th style="text-align: left;">MDI</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>10/100TX</td> <td>Yes</td> <td>Auto</td> <td>Enable</td> <td>Auto</td> </tr> <tr> <td>2</td> <td>10/100TX</td> <td>Yes</td> <td>Auto</td> <td>Enable</td> <td>MDI</td> </tr> </tbody> </table> <p>All of the following operate on a per-port basis, regardless of trunk membership:</p> <ul style="list-style-type: none"> • Enable/Disable • Flow control (Flow Ctrl) <p>LACP is a full-duplex protocol. See “Trunk group operation using LACP” (page 118).</p> | Port | Type | Enabled | Mode | Flow Ctrl | MDI | 1 | 10/100TX | Yes | Auto | Enable | Auto | 2 | 10/100TX | Yes | Auto | Enable | MDI |
| Port | Type | Enabled | Mode | Flow Ctrl | MDI | | | | | | | | | | | | | | |
| 1 | 10/100TX | Yes | Auto | Enable | Auto | | | | | | | | | | | | | | |
| 2 | 10/100TX | Yes | Auto | Enable | MDI | | | | | | | | | | | | | | |
| Trunk configuration: | <p>All ports in the same trunk group must be the same trunk type (LACP or trunk). All LACP ports in the same trunk group must be either all static LACP or all dynamic LACP.</p> <p>A trunk appears as a single port labeled <code>Dyn1</code> (for an LACP dynamic trunk) or <code>Trk1</code> (for a static trunk of type LACP, Trunk) on various menu and CLI screens. For a listing of which screens show which trunk types, see “How the switch lists trunk data” (page 123).</p> <p>For spanning-tree or VLAN operation, configuration for all ports in a trunk is done at the trunk level. (You cannot separately configure individual ports within a trunk for spanning-tree or VLAN operation.)</p> | | | | | | | | | | | | | | | | | | |
| Traffic distribution: | <p>All of the switch trunk protocols use the SA/DA (source address/destination address) method of distributing traffic across the trunked links. See “Outbound traffic distribution across trunked links” (page 123).</p> | | | | | | | | | | | | | | | | | | |
| Spanning Tree: | <p>802.1D (STP) and 802.1w (RSTP) Spanning Tree operate as a global setting on the switch (with one instance of Spanning Tree per switch). 802.1s (MSTP) Spanning Tree operates on a per-instance basis (with multiple instances allowed per switch). For each Spanning Tree instance, you can adjust Spanning Tree parameters on a per-port basis.</p> <p>A static trunk of any type appears in the Spanning Tree configuration display, and you can configure Spanning Tree parameters for a static trunk in the same way that you would configure Spanning Tree parameters on a non-trunked port. (Note that the switch lists the trunk by name—such as <code>Trk1</code>—and does not list the individual ports in the trunk.) For example, if ports C1 and C2 are configured as a static trunk named <code>Trk1</code>, they are listed in the Spanning Tree display as <code>Trk1</code> and do not appear as individual ports in the Spanning Tree displays. See Example 82 (page 112).</p> <p>When Spanning Tree forwards on a trunk, all ports in the trunk will be forwarding. Conversely, when Spanning Tree blocks a trunk, all ports in the trunk are blocked.</p> <p>NOTE: A dynamic LACP trunk operates only with the default Spanning Tree settings. Also, this type of trunk appears in the CLI show spanning-tree display, but not in the Spanning Tree Operation display of the Menu interface.</p> <p>If you remove a port from a static trunk, the port retains the same Spanning Tree settings that were configured for the trunk.</p> <p>In the below Example:, ports C1 and C2 are members of <code>TRK1</code> and do not appear as individual ports in the port configuration part of the listing.</p> | | | | | | | | | | | | | | | | | | |

Table 13 General operating rules for port trunks *(continued)*

| | | Example 82 A port trunk in a Spanning Tree listing | | | | | |
|--------------------------------------|--|---|-----------|-------|----------|------------|-------------------|
| | | Port | Type | Cost | Priority | State | Designated Bridge |
| | | ----- | ----- | ----- | ----- | ----- | ----- |
| | | C3 | 100/1000T | 5 | 12B | Forwarding | 0020c1-b27ac0 |
| | | C4 | 100/1000T | 5 | 12B | Forwarding | 0060b0-889e00 |
| | | C5 | 100/1000T | 5 | 12B | Disabled | |
| | | C6 | 100/1000T | 5 | 12B | Disabled | |
| | | Trk1 | | 1 | 64 | Forwarding | 0001e7-a0ec00 |
| | | ----- | ----- | ----- | ----- | ----- | ----- |
| IP multicast protocol (IGMP): | <p>A static trunk of any type appears in the IGMP configuration display, and you can configure IGMP for a static trunk in the same way that you would configure IGMP on a non-trunked port. (Note that the switch lists the trunk by name—such as Trk1—and does not list the individual ports in the trunk.) Also, creating a new trunk automatically places the trunk in IGMP Auto status if IGMP is enabled for the default VLAN.</p> <p>A dynamic LACP trunk operates only with the default IGMP settings and does not appear in the IGMP configuration display or <code>show ip igmp</code> listing.</p> | | | | | | |
| VLANs: | <p>Creating a new trunk automatically places the trunk in the DEFAULT_VLAN, regardless of whether the ports in the trunk were in another VLAN. Similarly, removing a port from a trunk group automatically places the port in the default VLAN. You can configure a static trunk in the same way that you configure a port for membership in any VLAN.</p> <p>NOTE: For a dynamic LACP trunk to operate in a VLAN other than the default VLAN (DEFAULT_VLAN), GVRP must be enabled. See “Trunk group operation using LACP” (page 118).</p> | | | | | | |
| Port security: | <p>Trunk groups (and their individual ports) cannot be configured for port security, and the switch excludes trunked ports from the <code>show port-security</code> listing. If you configure non-default port security settings for a port, then subsequently try to place the port in a trunk, you see the following message and the command is not executed:</p> <pre>< port-list> Command cannot operate over a logical port.</pre> | | | | | | |
| Monitor port: | <p>NOTE: A trunk cannot be a monitor port. A monitor port can monitor a static trunk but cannot monitor a dynamic LACP trunk.</p> | | | | | | |

Viewing and configuring a static trunk group (Menu)

- ① **IMPORTANT:** Configure port trunking *before* you connect the trunked links to another switch, routing switch, or server. Otherwise, a broadcast storm could occur. (If you need to connect the ports before configuring them for trunking, you can temporarily disable the ports until the trunk is configured. See "Enabling or Disabling Ports and Configuring Port Mode".)

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

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

Figure 14 Example: of the menu screen for configuring a port trunk group

```

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

Port   Type   Enabled   Mode   Flow Ctrl   Group   Type
-----+-----
C1    10/100TX | Yes      Auto   Disable     |      |
C2    10/100TX | Yes      Auto   Disable     |      |
C3    10/100TX | Yes      Auto   Disable     |      |
C4    10/100TX | Yes      Auto   Disable     |      |
C5    10/100TX | Yes      Auto   Disable     |      |
C6    10/100TX | Yes      Auto   Disable     |      |

Actions->  Cancel   Edit   Save   Help

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

```

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

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

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

Figure 15 Example: of the Configuration for a Two-Port Trunk Group

```

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

Port   Type   Enabled   Mode   Flow Ctrl   Group   Type
-----+-----
C1    10/100TX | Yes      Auto   Disable     |      |
C2    10/100TX | Yes      Auto   Disable     |      |
C3    10/100TX | Yes      Auto   Disable     |      |
C4    10/100TX | Yes      Auto   Disable     |      |
C5    10/100TX | Yes      Auto   Disable     | Trk1 | Trunk
C6    10/100TX | Yes      Auto   Disable     | Trk1 | Trunk

Actions->  Cancel   Edit   Save   Help

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

```

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

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

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

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

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

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

Viewing and configuring port trunk groups (CLI)

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

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

Syntax:

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

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

Example:

In a switch where ports A4 and A5 belong to Trunk 1 and ports A7 and A8 belong to Trunk 2, you have the options shown in figures [Example 83 \(page 114\)](#) and [Example 84](#) for displaying port data for ports belonging to static trunks.

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

Example 83 Listing specific ports belonging to static trunks

```
HP Switch> show trunks e 5-7
```

```
Load Balancing
```

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

The `show trunks <port-list>` command in the above Example: includes a port list, and thus shows trunk group information only for specific ports that have membership in a static trunk. In [Example 84](#), the command does not include a port list, so the switch lists all ports having static trunk membership.

Example 84 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 15 \(page 120\)](#).)

Example 85 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 15 \(page 120\)](#).

Dynamic LACP Standby Links

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

Example 86 A Dynamic LACP trunk with one standby link

```
HP Switch> show lacp
```

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

Configuring a static trunk or static LACP trunk group

- ① **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 11](#) describes the maximum number of trunk groups you can configure on the switch. An individual trunk can have up to eight links, with additional standby links if you're using LACP. You can configure trunk group types as follows:

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

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

Syntax:

```
trunk <port-list> <trk1 ... trk60> <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:

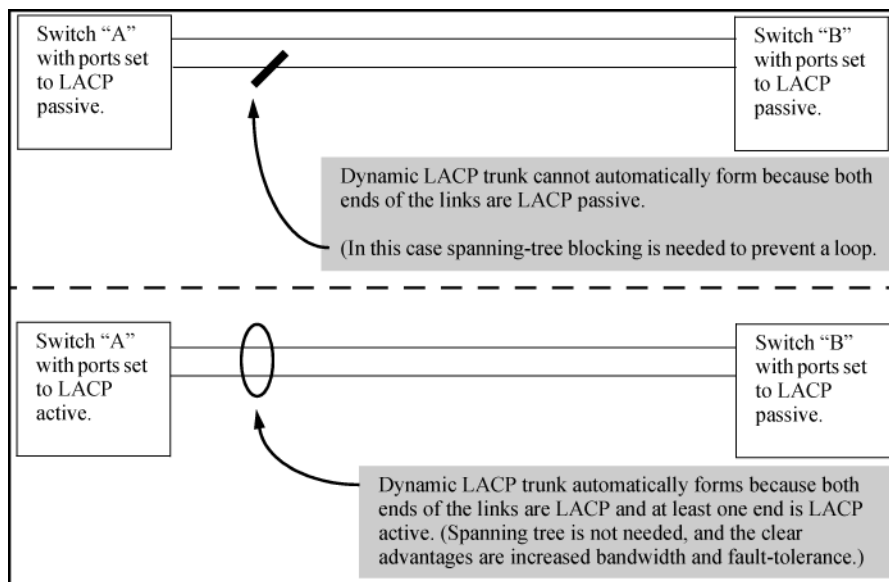
```
HP Switch(config)# no trunk c4-c5
```

Enabling 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 16 Criteria for automatically forming a dynamic LACP trunk



Syntax:

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

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

Example:

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

```
HP Switch(config)# interface c4-c5 lacp active
```

Removing ports from a dynamic LACP trunk group

To remove a port from dynamic LACP trunk operation, you must turn off LACP on the port. (On a port in an operating, dynamic LACP trunk, you cannot change between LACP Active and LACP passive without first removing LACP operation from the port.)

- △ CAUTION:** Unless spanning tree is running on your network, removing a port from a trunk can result in a loop. To help prevent a broadcast storm when you remove a port from a trunk where spanning tree is not in use, HP recommends that you first disable the port or disconnect the link on that port.

Syntax:

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

Removes <port-list> from any dynamic LACP trunk and returns the ports in <port-list> to passive LACP.

Example:

Port C6 belongs to an operating, dynamic LACP trunk. To remove port C6 from the dynamic trunk and return it to passive LACP, do the following:

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

In the above Example:, if the port on the other end of the link is configured for active LACP or static LACP, the trunked link will be re-established almost immediately.

Viewing existing port trunk groups (WebAgent)

While the WebAgent does not enable you to configure a port trunk group, it does provide a view of an existing trunk group.

To view any port trunk groups:

1. In the navigation pane, click **Interface**.
2. Click **Port Info/Config**. The trunk information for the port displays in the **Port Properties** box.

Trunk group operation using LACP

The switch can automatically configure a dynamic LACP trunk group, or you can manually configure a static LACP trunk group.

NOTE: LACP requires full-duplex (FDx) links of the same media type (10/100Base-T, 100FX, and so on) and the same speed and enforces speed and duplex conformance across a trunk group. For most installations, HP Switch recommends that you leave the port mode settings at `Auto` (the default). LACP also operates with `Auto-10`, `Auto-100`, and `Auto-1000` (if negotiation selects FDx), and `10FDx`, `100FDx`, and `1000FDx` settings.

LACP trunk status commands include:

| Trunk display method | Static LACP trunk | Dynamic LACP trunk |
|--|----------------------|----------------------|
| CLI <code>show lacp</code> command | Included in listing. | Included in listing. |
| CLI <code>show trunk</code> command | Included in listing. | Not included. |
| Port/Trunk Settings screen in menu interface | Included in listing. | Not included |

Thus, to display a listing of dynamic LACP trunk ports, you must use the `show lacp` command. In most cases, trunks configured for LACP on the switches operate as described in [Table 14 \(page 118\)](#).

Table 14 LACP trunk types

| LACP port trunk configuration | Operation |
|-------------------------------|--|
| Dynamic LACP | This option automatically establishes an 802.3ad-compliant trunk group, with LACP for the port Type parameter and DynX for the port Group name, where X is an automatically assigned value from 1 to 60, depending on how many dynamic and static trunks are currently on the switch. (The switch allows a maximum of 60 trunk groups in any combination of static and dynamic trunks.) |

Table 14 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 121).</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="630 653 1010 768" data-label="Diagram"> <pre> graph LR subgraph Switch1 [Switch 1] direction TB P1[Port X: LACP Enable: Active] P2[Port Y: LACP Enable: Active] end subgraph Switch2 [Switch 2] direction TB P3[Port A: LACP Enable: Active] P4[Port B: LACP Enable: Passive] end P1 --- Active-to-Active P3 P2 --- Active-to-Passive P4 </pre> </div> <p>Either of the above link configurations allows a dynamic LACP trunk link.</p> <p>Backup Links: A maximum of eight operating links are allowed in the trunk, but, with dynamic LACP, you can configure one or more additional (backup) links that the switch automatically activates if a primary link fails. To configure a link as a standby for an existing eight-port dynamic LACP trunk, ensure that the ports in the standby link are configured as either active-to-active or active-to-passive between switches.</p> <p>Displaying dynamic LACP trunk data: To list the configuration and status for a dynamic LACP trunk, use the CLI <code>show lacp</code> command.</p> <p>NOTE: The dynamic trunk is automatically created by the switch and is not listed in the static trunk listings available in the menu interface or in the CLI <code>show trunk</code> listing.</p> |
| <p>Static LACP</p> | <p>Provides a manually configured, static LACP trunk to accommodate these conditions:</p> <ul style="list-style-type: none"> • The port on the other end of the trunk link is configured for a static LACP trunk. • You want to configure non-default Spanning Tree or IGMP parameters on an LACP trunk group. • You want an LACP trunk group to operate in a VLAN other than the default VLAN and GVRP is disabled. (See “VLANs and dynamic LACP” (page 121).) • You want to use a monitor port on the switch to monitor an LACP trunk. <p>The trunk operates if the trunk group on the opposite device is running one of the following trunking protocols:</p> <ul style="list-style-type: none"> • Active LACP • Passive LACP • Trunk <p>This option uses LACP for the port Type parameter and TrkX for the port Group parameter, where X is an automatically assigned value in a range corresponding to the maximum number of trunks the switch allows. (The table on Table 11 (page 110) 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 15 \(page 120\)](#) 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 15 LACP port status data

| Status name | Meaning |
|--------------|--|
| Port Numb | Shows the physical port number for each port configured for LACP operation (C1, C2, C3 ...). Unlisted port numbers indicate that the missing ports that are assigned to a static trunk group are not configured for any trunking. |
| LACP Enabled | <p>Active: The port automatically sends LACP protocol packets.</p> <p>Passive: The port does not automatically send LACP protocol packets and responds only if it receives LACP protocol packets from the opposite device.</p> <p>A link having either two active LACP ports or one active port and one passive port can perform dynamic LACP trunking. A link having two passive LACP ports does not perform LACP trunking because both ports are waiting for an LACP protocol packet from the opposite device.</p> <p>NOTE: In the default switch configuration, LACP is disabled for all ports.</p> |
| Trunk Group | <p>TrkX: This port has been manually configured into a static LACP trunk.</p> <p>Trunk group same as port number: The port is configured for LACP, but is not a member of a port trunk.</p> |
| Port Status | <p>Up: The port has an active LACP link and is not blocked or in standby mode.</p> <p>Down: The port is enabled, but an LACP link is not established. This can indicate, For example, a port that is not connected to the network or a speed mismatch between a pair of linked ports.</p> <p>Disabled: The port cannot carry traffic.</p> <p>Blocked: LACP, Spanning Tree has blocked the port. (The port is not in LACP standby mode.) This may be caused by a (brief) trunk negotiation or a configuration error, such as differing port speeds on the same link or trying to connect the switch to more trunks than it can support. (See the table on Table 12.)</p> <p>NOTE: Some older devices are limited to four ports in a trunk. When eight LACP-enabled ports are connected to one of these older devices, four ports connect, but the other four ports are blocked.</p> <p>Standby: The port is configured for dynamic LACP trunking to another device, but the maximum number of ports for the dynamic trunk to that device has already been reached on either the switch or the other device. This port will remain in reserve, or "standby" unless LACP detects that another, active link in the trunk has become disabled, blocked, or down. In this case, LACP automatically assigns a standby port, if available, to replace the failed port.</p> |
| LACP Partner | <p>Yes: LACP is enabled on both ends of the link.</p> <p>No: LACP is enabled on the switch, but either LACP is not enabled or the link has not been detected on the opposite device.</p> |
| LACP Status | <p>Success: LACP is enabled on the port, detects and synchronizes with a device on the other end of the link, and can move traffic across the link.</p> <p>Failure: LACP is enabled on a port and detects a device on the other end of the link, but is not able to synchronize with this device, and therefore is not able to send LACP packets across the link. This can be caused, For example, by an intervening device on the link (such as a hub), a bad hardware connection, or if the LACP operation on the opposite device does not comply with the IEEE 802.3ad standard.</p> |

LACP notes and restrictions

802.1X (Port-based access control) configured on a port

To maintain security, LACP is not allowed on ports configured for 802.1X authenticator operation. If you configure port security on a port on which LACP (active or passive) is configured, the switch

removes the LACP configuration, displays a notice that LACP is disabled on the port, and enables 802.1X on that port.

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

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

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

To restore LACP to the port, you must first remove the 802.1X configuration of the port and then re-enable LACP active or passive on the port.

Port security configured on a port

To maintain security, LACP is not allowed on ports configured for port security. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port, and enables port security on that port. For Example:

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

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

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

To restore LACP to the port, you must remove port security and re-enable LACP active or passive.

Changing trunking methods

To convert a trunk from static to dynamic, you must first eliminate the static trunk.

Static LACP trunks

When a port is configured for LACP (active or passive), but does not belong to an existing trunk group, you can add that port to a static trunk. Doing so disables dynamic LACP on that port, which means you must manually configure both ends of the trunk.

Dynamic LACP trunks

You can configure a port for LACP-active or LACP-passive, but on a dynamic LACP trunk you cannot configure the other options that you can on static trunks. If you want to manually configure a trunk, use the `trunk` command.

VLANs and dynamic LACP

A dynamic LACP trunk operates only in the default VLAN (unless you have enabled GVRP on the switch and use `Forbid` to prevent the ports from joining the default VLAN).

If you want to use LACP for a trunk on a non-default VLAN and GVRP is disabled, configure the trunk as a static trunk.

Blocked ports with older devices

Some older devices are limited to four ports in a trunk. When eight LACP-enabled ports are connected to one of these older devices, four ports connect, but the other four ports are blocked. The LACP status of the blocked ports is shown as "Failure."

If one of the other ports becomes disabled, a blocked port replaces it (Port Status becomes "Up"). When the other port becomes active again, the replacement port goes back to blocked (Port Status is "Blocked"). It can take a few seconds for the switch to discover the current status of the ports.

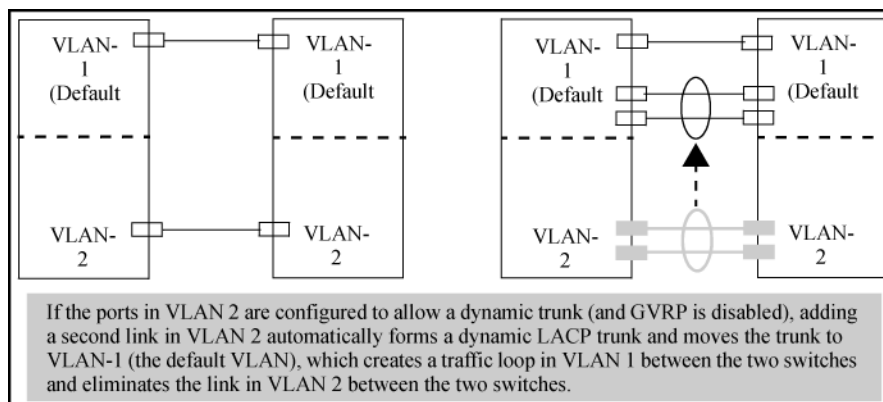
Example 87 Blocked ports with LACP

```
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 17 A dynamic LACP trunk forming in a VLAN can cause a traffic loop



Easy control methods include either disabling LACP on the selected ports or configuring them to operate in static LACP trunks.

Spanning Tree and IGMP

If Spanning Tree, IGMP, or both are enabled in the switch, a dynamic LACP trunk operates only with the default settings for these features and does not appear in the port listings for these features.

Half-duplex, different port speeds, or both not allowed in LACP trunks

The ports on both sides of an LACP trunk must be configured for the same speed and for full-duplex (FDx). The 802.3ad LACP standard specifies a full-duplex (FDx) requirement for LACP trunking. (10-gigabit ports operate only at FDx.)

A port configured as LACP passive and not assigned to a port trunk can be configured to half-duplex (HDx). However, in any of the following cases, a port cannot be reconfigured to an HDx setting:

- If the port is a 10-gigabit port.
- If a port is set to LACP Active, you cannot configure it to HDx.
- If a port is already a member of a static or dynamic LACP trunk, you cannot configure it to HDx.
- If a port is already set to HDx, the switch does not allow you to configure it for a static or dynamic LACP trunk.

Dynamic/static LACP interoperation

A port configured for dynamic LACP can properly interoperate with a port configured for static (TrkX) LACP, but any ports configured as standby LACP links are ignored.

Trunk group operation using the "trunk" option

This method creates a trunk group that operates independently of specific trunking protocols and does not use a protocol exchange with the device on the other end of the trunk. With this choice, the switch simply uses the SA/DA method of distributing outbound traffic across the trunked ports without regard for how that traffic is handled by the device at the other end of the trunked links. Similarly, the switch handles incoming traffic from the trunked links as if it were from a trunked source.

When a trunk group is configured with the `trunk` option, the switch automatically sets the trunk to a priority of "4" for Spanning Tree operation (even if Spanning Tree is currently disabled). This appears in the running-config file as `spanning-tree Trkn priority 4`. Executing `write memory` after configuring the trunk places the same entry in the startup-config file.

Use the `trunk` option to establish a trunk group between a switch and another device, where the other device's trunking operation fails to operate properly with LACP trunking configured on the switches.

How the switch lists trunk data

| | |
|--------------------------|--|
| Static trunk group | Appears in the menu interface and the output from the CLI <code>show trunk</code> and <code>show interfaces</code> commands. |
| Dynamic LACP trunk group | Appears in the output from the CLI <code>show lacp</code> command. |

| Interface option | Dynamic LACP trunk group | Static LACP trunk group | Static non-protocol |
|-------------------------------------|--------------------------|-------------------------|---------------------|
| Menu interface | No | Yes | Yes |
| CLI <code>show trunk</code> | No | Yes | Yes |
| CLI <code>show interfaces</code> | No | Yes | Yes |
| CLI <code>show lacp</code> | Yes | Yes | No |
| CLI <code>show spanning-tree</code> | No | Yes | Yes |
| CLI <code>show igmp</code> | No | Yes | Yes |
| CLI <code>show config</code> | No | Yes | Yes |


Outbound traffic distribution across trunked links

The two trunk group options (LACP and trunk) use SA/DA pairs for distributing outbound traffic over trunked links. That is, the switch sends traffic from the same source address to the same destination address through the same trunked link, and may also send traffic from the same source

address to a different destination address through the same link or a different link, depending on the mapping of path assignments among the links in the trunk. Likewise, the switch distributes traffic for the same destination address but from different source addresses through links depending on the path assignment.

The load-balancing is done on a per-communication basis. Otherwise, traffic is transmitted across the same path as shown in [Figure 18 \(page 124\)](#). 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 18 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 19 \(page 124\)](#) showing a three-port trunk, traffic could be assigned as shown in [Table 16 \(page 124\)](#).

Figure 19 Example: of port-trunked network




Table 16 Example: of link assignments in a trunk group (SA/DA distribution)

| Source | Destination | Link |
|--------|-------------|------|
| Node A | Node W | 1 |
| Node B | Node X | 2 |
| Node C | Node Y | 3 |
| Node D | Node Z | 1 |
| Node A | Node Y | 2 |
| Node B | Node W | 3 |

Because the amount of traffic coming from or going to various nodes in a network can vary widely, it is possible for one link in a trunk group to be fully utilized while other links in the same trunk have unused bandwidth capacity, even if the assignments were evenly distributed across the links in a trunk.

Trunk load balancing using port layers

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

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

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

Enabling trunk load balancing

Enter the following command to enable load balancing.

Syntax:

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

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

Default: L3-based load balancing

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

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

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

Example 88 Enabling L4-based trunk load balancing

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

Example 89 Output when L4-based trunk load balancing is enabled

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

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

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

Note in [Example 90 “Running config file when L4-based trunk load balancing is enabled”](#) that in if L4 trunk load balancing is enabled, a line appears in the running-config file. If it is not enabled, nothing appears as this is the default and the default values are not displayed.

Example 90 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 #XX.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
```

5 Port Traffic Controls

Rate-limiting

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

All traffic rate-limiting

Rate-limiting for all traffic operates on a per-port basis to allow only the specified bandwidth to be used for inbound 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.

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

The switches also support ICMP rate-limiting to mitigate the effects of certain ICMP-based attacks.

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 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" \(page 127\)](#).

Notes:

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

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

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

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

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

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

Displaying the current rate-limit configuration

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

Syntax:

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

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

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

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

Figure 20 Listing the rate-limit configuration

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

| 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 |

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-)

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

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

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

The `show running` command displays the currently applied setting for any interfaces in the switch configured for all traffic rate-limiting and ICMP rate limiting.

The `show config` command displays this information for the configuration currently stored in the `startup-config` file. (Note that configuration changes performed with the CLI, but not followed by a `write mem` command, do not appear in the `startup-config` file.)

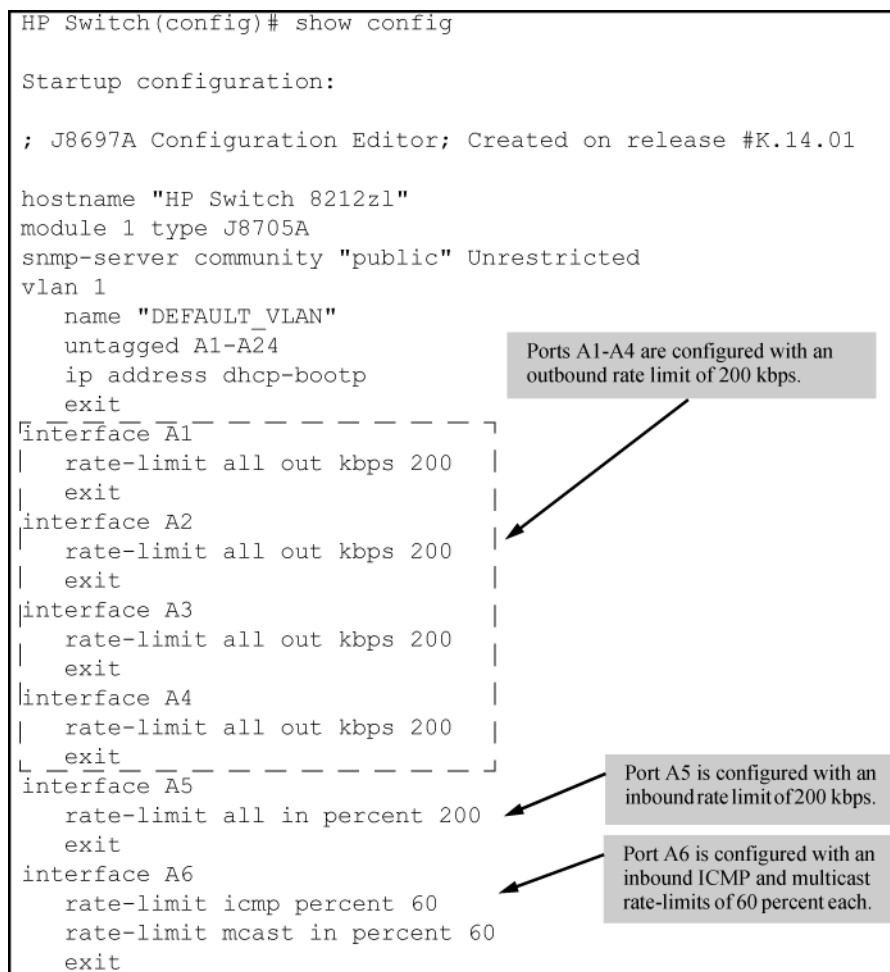
Figure 21 Example: of rate-limit settings listed in the `show config` output

```
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
```



Operating notes for rate-limiting

- **Rate-limiting operates on a per-port basis, regardless of traffic priority.** Rate-limiting is available on all types of ports (other than trunked ports) and at all port speeds configurable for these switches.
- **Rate-limiting is not allowed on trunked ports.** Rate-limiting is not supported on ports configured in a trunk group (including mesh ports). Configuring a port for rate-limiting and then adding

it to a trunk suspends rate-limiting on the port while it is in the trunk. Attempting to configure rate-limiting on a port that already belongs to a trunk generates the following message:

```
<port-list>: Operation is not allowed for a trunked port.
```

- **Rate-limiting and hardware.** The hardware will round the actual Kbps rate down to the nearest multiple of 64 Kbps.
- **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: For the HP Switch 2920-series, it is required for user to use a single PDU when trying to configure through SNMP.

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

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:

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

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 17 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 17 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 [132](#).)

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 GMB for outbound traffic](#)" (page [133](#)) 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 18 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 chapter "Quality of Service (QoS): Managing Bandwidth More Effectively" in the *Advanced Traffic Management Guide* for your switch.

Configuring 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 131).

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 131).

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 (queues 5 and 6) 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. |
| 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.)

Example 91 “Listing the GMB configuration” displays the GMB configuration resulting from either of the above commands.

Example 91 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
```

Example 92 “GMB settings listed in the show config output” shows how the preceding listing of the GMB configuration would appear in the startup-config file.

Example 92 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 #XX.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
```

GMB operating notes

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 133).

Jumbo frames

The maximum transmission unit (MTU) is the maximum size IP frame the switch can receive for Layer 2 frames inbound on a port. The switch drops any inbound frames larger than the MTU allowed on the port. Ports operating at a minimum of 1 Gbps can accept forward frames of up to 9220 bytes (including four bytes for a VLAN tag) when configured for jumbo traffic. You can enable inbound jumbo frames on a per-VLAN basis. That is, on a VLAN configured for jumbo traffic, all ports belonging to that VLAN and *operating* at a minimum of 1 Gbps allow inbound jumbo frames of up to 9220 bytes.

Operating rules

- **Required port speed:** This feature allows inbound and outbound jumbo frames on ports operating at a minimum of 1 Gbps.
- **GVRP operation:** A VLAN enabled for jumbo traffic cannot be used to create a dynamic VLAN. A port belonging to a statically configured, jumbo-enabled VLAN cannot join a dynamic VLAN.
- **Port adds and moves:** If you add a port to a VLAN that is already configured for jumbo traffic, the switch enables that port to receive jumbo traffic. If you remove a port from a jumbo-enabled VLAN, the switch disables jumbo traffic capability on the port only if the port is not currently a member of another jumbo-enabled VLAN. This same operation applies to port trunks.
- **Jumbo traffic sources:** A port belonging to a jumbo-enabled VLAN can receive inbound jumbo frames through any VLAN to which it belongs, including non-jumbo VLANs. For example, if VLAN 10 (without jumbos enabled) and VLAN 20 (with jumbos enabled) are both configured on a switch, and port 1 belongs to both VLANs, port 1 can receive jumbo traffic from devices on either VLAN. For a method to allow only some ports in a VLAN to receive jumbo traffic, see [“Configuring a maximum frame size” \(page 139\)](#).

Configuring jumbo frame operation

For detailed information about jumbo frames, see [“Jumbo frames” \(page 137\)](#).

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” \(page 139\)](#).) See [Figure 22](#).

Figure 22 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 23.)

Figure 23 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 24.)

Figure 24 Example: of listing the port membership and jumbo status for a VLAN

```
HP Switch(config)# show vlan 100
Status and Counters - VLAN Information - VLAN 100
VLAN ID : 100
Name : VLAN100
Status : Port-based Voice : No
Jumbo : No
Port Information Mode Unknown VLAN Status
-----
A1 Tagged Learn Up
A2 Tagged Learn Up
A3 Tagged Learn Up
A4 Tagged Learn Down
A5 Tagged Learn Up
```

Lists the ports belonging to VLAN 100 and whether the VLAN is enabled for jumbo frame traffic.

Enabling or disabling jumbo traffic on a VLAN

Syntax:

```
vlan <vid> jumbo
[no] vlan <vid> jumbo
```

Configures the specified VLAN to allow jumbo frames on all ports on the switch that belong to that VLAN. If the VLAN is not already configured on the switch, `vlan <vid> jumbo` also creates the VLAN.

A port belonging to one jumbo VLAN can receive jumbo frames through any other VLAN statically configured on the switch, regardless of whether the other VLAN is enabled for jumbo frames.

The `[no]` form of the command disables inbound jumbo traffic on all ports in the specified VLAN that do not also belong to another VLAN that is enabled for jumbo traffic. In a VLAN context, the command forms are `jumbo` and `no jumbo`.

(Default: Jumbos disabled on the specified VLAN.)

Configuring a maximum frame size

You can globally set a maximum frame size for jumbo frames that will support values from 1518 bytes to 9216 bytes for untagged frames.

Syntax:

```
jumbo max-frame-size <size>
```

Sets the maximum frame size for jumbo frames. The range is from 1518 bytes to 9216 bytes. (Default: 9216 bytes)

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

Default: 9216 bytes

Configuring IP MTU

NOTE: The following feature is available on the switches covered in this guide. `jumbos` support is required for this feature. On switches that do not support this command, the IP MTU value is derived from the maximum frame size and is not configurable.

You can set the IP MTU globally by entering this command. The value of `max-frame-size` must be greater than or equal to 18 bytes more than the value selected for `ip-mtu`. For example, if `ip-mtu` is set to 8964, the `max-frame-size` is configured as 8982.

Syntax:

```
jumbo ip-mtu <size>
```

Globally sets the IP MTU size. Values range between 1500 and 9198 bytes. This value must be 18 bytes less than the value of `max-frame-size`.

(Default: 9198 bytes)

SNMP implementation

Jumbo maximum frame size

The maximum frame size for jumbos is supported with the following proprietary MIB object:

```
hpSwitchMaxFrameSize OBJECT-TYPE
```

This is the value of the global `max-frame-size` supported by the switch. The default value is set to 9216 bytes.

Jumbo IP MTU

The IP MTU for jumbos is supported with the following proprietary MIB object:

```
hpSwitchIpMTU OBJECT-TYPE
```

This is the value of the global jumbos IP MTU (or L3 MTU) supported by the switch. The default value is set to 9198 bytes (a value that is 18 bytes less than the largest possible maximum frame size of 9216 bytes). This object can be used only in switches that support `max-frame-size` and `ip-mtu` configuration.

Displaying the maximum frame size

Use the `show jumbos` command to display the globally configured untagged maximum frame size for the switch, as shown in the following Example:.

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

```
Jumbos Global Values
```

```
Configured : MaxFrameSize : 9216   Ip-MTU : 9198  
In Use      : MaxFrameSize : 9216   Ip-MTU : 9198
```

For more information about frame size, see [“Jumbo frames” \(page 137\)](#).

Operating notes for maximum frame size

- When you set a maximum frame size for jumbo frames, it must be on a global level. You cannot use the `jumbo max-frame-size` command on a per-port or per-VLAN basis.
- The original way to configure jumbo frames remains the same, which is per-VLAN, but you cannot set a maximum frame size per-VLAN.
- Jumbo support must be enabled for a VLAN from the CLI or through SNMP.
- Setting the maximum frame size does not require a reboot.
- When you upgrade to a version of software that supports setting the maximum frame size from a version that did not, the `max-frame-size` value is set automatically to 9216 bytes.
- Configuring a jumbo maximum frame size on a VLAN allows frames up to `max-frame-size` even though other VLANs of which the port is a member are not enabled for jumbo support.

Operating notes for jumbo traffic-handling

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

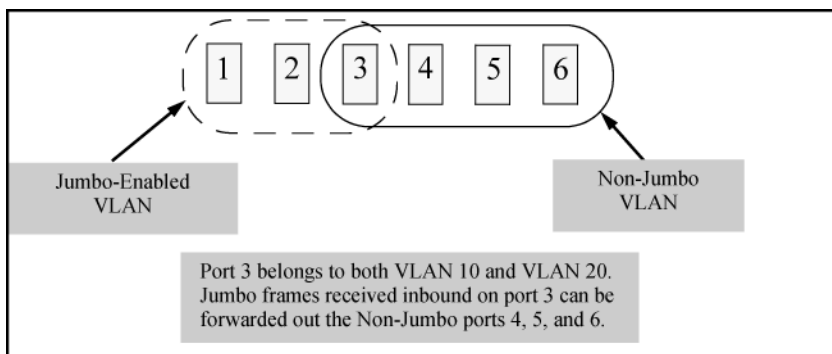
For example, suppose you want to allow inbound jumbo frames only on ports 6, 7, 12, and 13. However, these ports are spread across VLAN 100 and VLAN 200 and also share these VLANs with other ports you want excluded from jumbo traffic. A solution is to create a third VLAN with the sole purpose of enabling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

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

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

- **Outbound jumbo traffic.** Any port operating at 1 Gbps or higher can transmit outbound jumbo frames through any VLAN, regardless of the jumbo configuration. The VLAN is not required to be jumbo-enabled, and the port is not required to belong to any other, jumbo-enabled VLANs. This can occur in situations where a non-jumbo VLAN includes some ports that do not belong to another, jumbo-enabled VLAN and some ports that do belong to another, jumbo-enabled VLAN. In this case, ports capable of receiving jumbo frames can forward them to the ports in the VLAN that do not have jumbo capability, as shown in [Figure 25](#).

Figure 25 Forwarding jumbo frames through non-jumbo ports



Jumbo frames can also be forwarded out non-jumbo ports when the jumbo frames received inbound on a jumbo-enabled VLAN are routed to another, non-jumbo VLAN for outbound transmission on ports that have no memberships in other, jumbo-capable VLANs. Where either of the above scenarios is a possibility, the downstream device must be configured to accept the jumbo traffic. Otherwise, this traffic will be dropped by the downstream device.

Troubleshooting

A VLAN is configured to allow jumbo frames, but one or more ports drops all inbound jumbo frames

The port may not be operating at a minimum of 1 Gbps on the other switches covered in this guide. Regardless of a port's configuration, if it is actually operating at a speed lower than 1 Gbps for the other switches, it drops inbound jumbo frames. For example, if a port is configured for `Auto` mode (`speed-duplex auto`), but has negotiated a 7 Mbps speed with the device at the other end of the link, the port cannot receive inbound jumbo frames. To determine the actual operating speed of one or more ports, view the `Mode` field in the output for the following command:

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

A non-jumbo port is generating "Excessive undersize/giant frames" messages in the Event Log

The switches can transmit outbound jumbo traffic on any port, regardless of whether the port belongs to a jumbo VLAN. In this case, another port in the same VLAN on the switch may be jumbo-enabled through membership in a different, jumbo-enabled VLAN, and may be forwarding jumbo frames received on the jumbo VLAN to non-jumbo ports.

6 Configuring for Network Management Applications

Using SNMP tools to manage the switch

SNMP is a management protocol that allows an SNMP client application to retrieve device configuration and status information and to configure the device (*get* and *set*). You can manage the switch via SNMP from a network management station running an application such as PCM+. For more information on PCM+, see the HP website at: www.hp.com/networking.

From the **Products** menu, select **Network Management**. Then click on **PCM+ Network Management** under the **HP Network Management** bar.

To implement SNMP management, the switch must have an IP address configured either manually or dynamically (using DHCP or Bootp). If multiple VLANs are configured, each VLAN interface should have its own IP address. For DHCP use with multiple VLANs, see section "The Primary VLAN" in the "Static Virtual LANs (VLANs)" chapter of the *Advanced Traffic Management Guide* for your switch.

NOTE: If you use the switch's Authorized IP Managers and Management VLAN features, ensure that the SNMP management station, the choice of switch port used for SNMP access to the switch, or both, are compatible with the access controls enforced by these features. Otherwise, SNMP access to the switch will be blocked.

For more information on Authorized IP Managers, see the *Access Security Guide* for your switch. (The latest version of this guide is available on the HP Networking website.) For information on the Management VLAN feature, see the section "The Secure Management VLAN" in the "Static Virtual LANs (VLANs)" chapter of the *Advanced Traffic Management Guide* for your switch.

SNMP management features

SNMP management features on the switch include:

- SNMP version 1, version 2c, or version 3 over IP
- Security via configuration of SNMP communities ("[SNMPv3 communities](#)" (page 149))
- Security via authentication and privacy for SNMPv3 access
- Event reporting via SNMP
 - Version 1 traps
 - RMON: groups 1, 2, 3, and 9
- PCM/PCM+
- Flow sampling using sFlow
- Standard MIBs, such as the Bridge MIB (RFC 1493), Ethernet MAU MIB (RFC 1515), and others.

The switch SNMP agent also uses certain variables that are included in an HP proprietary MIB (management information base) file. If you are using HP OpenView, you can ensure that it is using the latest version of the MIB file by downloading the file to the OpenView database. To do so, go to the HP Networking website at: www.hp.com/networking.

1. Type a model number of your switch (For example, 8212) or product number in the **Auto Search** text box.
2. Select an appropriate product from the drop down list.
3. Click the Display selected button.
4. From the options that appear, select Software downloads.
5. MIBs are available with switch software in the Other category.

Click on `software updates`, then `MIBs`.

SNMPv1 and v2c access to the switch

SNMP access requires an IP address and subnet mask configured on the switch. If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address.

Once an IP address is configured, the main steps for configuring SNMPv1 and v2c access management features are:

1. Configure the appropriate SNMP communities. (See “SNMPv3 communities” (page 149).)
2. Configure the appropriate trap receivers.

In some networks, authorized IP manager addresses are not used. In this case, all management stations using the correct community name may access the switch with the View and Access levels that have been set for that community. If you want to restrict access to one or more specific nodes, you can use the switch's IP Authorized Manager feature. (See the *Access Security Guide* for your switch.)

-
- △ CAUTION:** For PCM/PCM+ version 1.5 or earlier (or any TopTools version), deleting the “public” community disables some network management functions (such as traffic monitoring, SNMP trap generation, and threshold setting). If network management security is a concern, and you are using the above software versions, HP recommends that you change the write access for the “public” community to “Restricted.”
-

SNMPv3 access to the switch

SNMPv3 access requires an IP address and subnet mask configured on the switch. (See “IP Configuration” on page 8-2.) If you are using DHCP/Bootp to configure the switch, ensure that the DHCP/Bootp process provides the IP address. (See “DHCP/Bootp Operation”.)

Once you have configured an IP address, the main steps for configuring SNMPv3 access management features are the following:

1. Enable SNMPv3 for operation on the switch (see “Enabling SNMPv3” (page 145)).
2. Configure the appropriate SNMP users (see “SNMPv3 users” (page 146)).
3. Configure the appropriate SNMP communities (see “SNMPv3 communities” (page 149)).
4. Configure the appropriate trap receivers (see “SNMP notifications” (page 153)).

In some networks, authorized IP manager addresses are not used. In this case, all management stations using the correct User and community name may access the switch with the View and Access levels that have been set for that community. If you want to restrict access to one or more specific nodes, you can use the IP Authorized Manager feature for the switch. (See the *Access Security Guide* for your switch.)

SNMP version 3 (SNMPv3) adds some new commands to the CLI for configuring SNMPv3 functions. To enable 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.
-

Enabling and disabling switch for access from SNMPv3 agents

This includes the creation of the initial user record.

Syntax:

```
[no] snmpv3 enable
```

Enabling or disabling restrictions to access from only SNMPv3 agents

When enabled, the switch rejects all non-SNMPv3 messages.

Syntax:

```
[no] snmpv3 only
```

Enabling or disabling restrictions from all non-SNMPv3 agents to read-only access

Syntax:

```
[no] snmpv3 restricted-access
```

Viewing the operating status of SNMPv3

Syntax:

```
show snmpv3 enable
```

Viewing status of message reception of non-SNMPv3 messages

Syntax:

```
show snmpv3 only
```

Viewing status of write messages of non-SNMPv3 messages

Syntax:

```
show snmpv3 restricted-access
```

Enabling SNMPv3

The `snmpv3 enable` command allows the switch to:

- Receive SNMPv3 messages.
- Configure initial users.
- Restrict non-version 3 messages to "read only" (optional).

△ CAUTION: Restricting access to only version 3 messages makes the community named "public" inaccessible to network management applications (such as autodiscovery, traffic monitoring, SNMP trap generation, and threshold setting) from operating in the switch.

Example:

Example 93 SNMP version 3 enable command

```
HP Switch(config)# snmpv3 enable
SHMPv3 Initialization process.
Creating user 'initial'
Authentication Protocol: MD5
Enter authentication password: *****
Privacy protocol is DES
Enter privacy password: *****

User 'initial' is created
Would you like to create a user that uses SHA? y
Enter user name: templateSHA
Authentication Protocol: SHA
Enter authentication password: *****
Privacy protocol is DES
Enter privacy password: *****

User creation is done. SHMPv3 is now functional.
Would you like to restrict SHMPv1 and SNMPv2c messages to have read only
access (you can set this later by the command 'snmp restrict-access'): n
```

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

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

SNMPv3 users

NOTE: To create new users, most SNMPv3 management software requires an initial user record to clone. The initial user record can be downgraded and provided with fewer features, but not upgraded by adding new features. For this reason, HP recommends that when you enable SNMPv3, you also create a second user with SHA authentication and DES privacy.

To use SNMPv3 on the switch, you must configure the users that will be assigned to different groups:

1. Configure users in the User Table with the `snmpv3 user` command.
To view the list of configured users, enter the `show snmpv3 user` command (see [“Adding users” \(page 146\)](#)).
2. Assign users to Security Groups based on their security model with the `snmpv3 group` command (see [“Assigning users to groups \(CLI\)” \(page 148\)](#)).

CAUTION: If you add an SNMPv3 user without authentication, privacy, or both, to a group that requires either feature, the user will not be able to access the switch. Ensure that you add a user with the appropriate security level to an existing security group.

Adding users

To configure an SNMPv3 user, you must first add the user name to the list of known users with the `snmpv3 user` command, as shown in [Figure 26 \(page 147\)](#).

Figure 26 Adding SNMPv3 users and displaying SNMPv3 configuration

```
HP Switch(config)# snmpv3 user NetworkAdmin ← Add user Network Admin with
no authentication or privacy.
HP Switch(config)# snmpv3 user NetworkMgr auth md5 authpass priv privpass
Add user Network Mgr with authentication and privacy.
MD5 authentication is enabled and the password is set to "authpass".
Privacy is enabled and the password is set to "privpass".
HP Switch(config)# show snmpv3 user

Status and Counters - SNMP v3 Global Configuration Information

User Name          Auth. Protocol      Privacy Protocol
-----
initial            MD5                  CFB AES-128
NetworkAdmin       MD5                  CBC-DES
```

SNMPv3 user commands

Syntax:

```
[no] snmpv3 user <user_name>
```

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

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

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

Default: None

Listing Users

To display the management stations configured to access the switch with SNMPv3 and view the authentication and privacy protocols that each station uses, enter the `show snmpv3 user` command.

Syntax:

```
show snmpv3 user
```

Example 94 "Display of the management stations configured on VLAN 1" displays information about the management stations configured on VLAN 1 to access the switch.

Example 94 Display of the management stations configured on VLAN 1

```
HP Switch# configure terminal
HP Switch(config)# vlan 1
HP Switch(vlan-1)# show snmpv3 user
```

Status and Counters - SNMPv3 Global Configuration Information

| User Name | Auth. Protocol | Privacy Protocol |
|--------------|----------------|------------------|
| initial | MD5 | CFB AES-128 |
| NetworkAdmin | MD5 | CBC-DES |

Assigning users to groups (CLI)

Next you must set the group access level for the user by assigning the user to a group. This is done with the `snmpv3 group` command, as shown in [Figure 27 \(page 148\)](#). For more details on the MIBs access for a given group, see [“Group access levels” \(page 148\)](#).

Figure 27 Example: of assigning users to groups

```
Switch(config)# snmpv3 group operatornoauth user NetworkAdmin sec-model ver3
Switch(config)# snmpv3 group managerpriv user NetworkMgr sec-model ver3
Switch(config)# show snmpv3 group
```

Status and Counters - SNHP v3 Global Configuration Information

| Security Name | Security Model | Group Name |
|----------------------------|----------------|----------------|
| CommunityManagerReadOnly | ver1 | ComManagerR |
| CommunityManagerReadWrite | ver1 | ComManagerRW |
| CommunityOperatorReadOnly | ver1 | ComOperatorRW |
| CommunityOperatorReadWrite | ver1 | ComOperatorRW |
| CommunityManagerReadOnly | ver2c | ComManagerR |
| CommunityManagerReadWrite | ver2c | ComManagerRW |
| CommunityOperatorReadOnly | ver2c | ComOperatorRW |
| CommunityOperatorReadWrite | ver2c | ComOperatorRW |
| NetworkMgr | ver3 | ManagerPriv |
| NetworkAdmin | ver3 | OperatorNoAuth |

Annotations in the image: 'Add NetworkAdmin to operator noauth group' points to the first command; 'Add NetworkMgr to managerpriv group' points to the second 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 148) . |
| <code>user <user_name></code> | Identifies the user to be added to the access group. This must match the user name added with the <code>snmpv3 user</code> command. |
| <code>sec-model <ver1 ver2c ver3></code> | Defines which security model to use for the added user. An SNMPv3 access group should use only the ver3 security model. |

Group access levels

The switch supports eight predefined group access levels, shown in [Table 6-3 \(page 149\)](#). 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 19 Predefined group access levels

| Group name | Group access type | Group read view | Group write view |
|----------------|---|------------------|------------------|
| managerpriv | Ver3 Must have Authentication and Privacy | ManagerReadView | ManagerWriteView |
| managerauth | Ver3 Must have Authentication | ManagerReadView | ManagerWriteView |
| operatorauth | Ver3 Must have Authentication | OperatorReadView | DiscoveryView |
| operatornoauth | Ver3 No Authentication | OperatorReadView | DiscoveryView |
| commanagerrw | Ver2c or Ver1 | ManagerReadView | ManagerWriteView |
| commanagerrr | Ver2c or Ver1 | ManagerReadView | DiscoveryView |
| comoperatorrw | Ver2c or Ver1 | OperatorReadView | OperatorReadView |
| comoperatorrr | Ver2c or Ver1 | OperatorReadView | DiscoveryView |

Each view allows you to view or modify a different set of MIBs:

- **Manager Read View** – access to all managed objects
- **Manager Write View** – access to all managed objects except the following:
 - vacmContextTable
 - vacmAccessTable
 - vacmViewTreeFamilyTable
- **OperatorReadView** – no access to the following:
 - icfSecurityMIB
 - hpSwitchIpTftpMode
 - vacmContextTable
 - vacmAccessTable
 - vacmViewTreeFamilyTable
 - usmUserTable
 - snmpCommunityTable
- **Discovery View** – Access limited to samplingProbe MIB.

NOTE: All access groups and views are predefined on the switch. There is no method to modify or add groups or views to those that are predefined on the switch.

SNMPv3 communities

SNMP communities are supported by the switch to allow management applications that use version 2c or version 1 to access the switch. The communities are mapped to Group Access Levels that are used for version 2c or version 1 support. This mapping happens automatically based on the communities access privileges, but special mappings can be added with the `snmpv3 community` command (see [“Mapping SNMPv3 communities \(CLI\)”](#) (page 149)).

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 149).

Syntax:

```
[no] snmpv3 community
```

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

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

Example:

Figure 28 (page 150) shows the assigning of the Operator community on MgrStation1 to the CommunityOperatorReadWrite group. Any other Operator has an access level of CommunityOperatorReadOnly.

Figure 28 Assigning a community to a group access level

```
Add mapping to allow write access for Operator community on MgrStation1
HP Switch(config)# snmpv3 Community index 30 name Operator sec-name
CommunityManagerReadWrite tag MgrStation1
HP Switch(config)# show snmpv3 community

snmpCommunityTable [rfc2576]
-----
Index Name          Community Name      Security Name
-----
1                  public             CommunityManagerReadWrite
2                  Operator          CommunityOperatorReadOnly
3                  Manager          CommunityManagerReadWrite
30                 Operator          CommunityManagerReadWrite
```

Two Operator Access Levels

SNMP community features

Use SNMP communities to restrict access to the switch by SNMP management stations by adding, editing, or deleting SNMP communities. You can configure up to five SNMP communities, each with either an operator-level or a manager-level view and either restricted or unrestricted write access.

Using SNMP requires that the switch have an IP address and subnet mask compatible with your network.

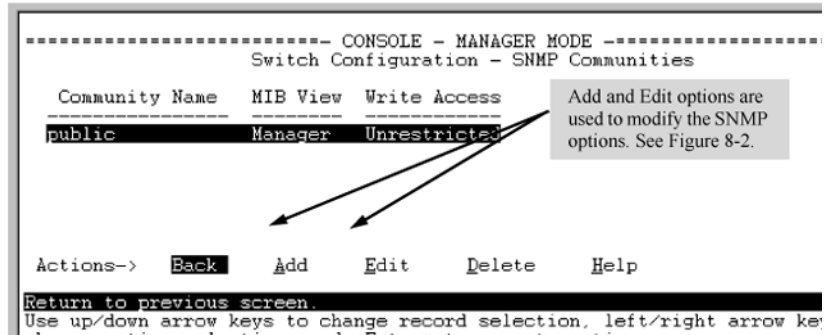
- △ **CAUTION:** For PCM/PCM+ version 1.5 or earlier (or any TopTools version), deleting the "public" community disables some network management functions (such as traffic monitoring, SNMP trap generation, and threshold setting). If network management security is a concern, and if you are using the above software versions, HP recommends that you change the write access for the "public" community to "Restricted."

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

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

Figure 29 The SNMP Communities screen (default values)

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



2. Press **[A]** (for **Add**).
If you need information on the options in each field, press **[Enter]** to move the cursor to the Actions line, then select the Help option. When you are finished with Help, press **[E]** (for Edit) to return the cursor to the parameter fields.
3. Enter the name you want in the Community Name field, and use the Space bar to select the appropriate value in each of the other fields. (Use the **[Tab]** key to move from one field to the next.)
4. Press **[Enter]**, then **[S]** (for **Save**).

Listing community names and values (CLI)

This command lists the data for currently configured SNMP community names (along with trap receivers and the setting for authentication traps—see “SNMP notifications” (page 153)).

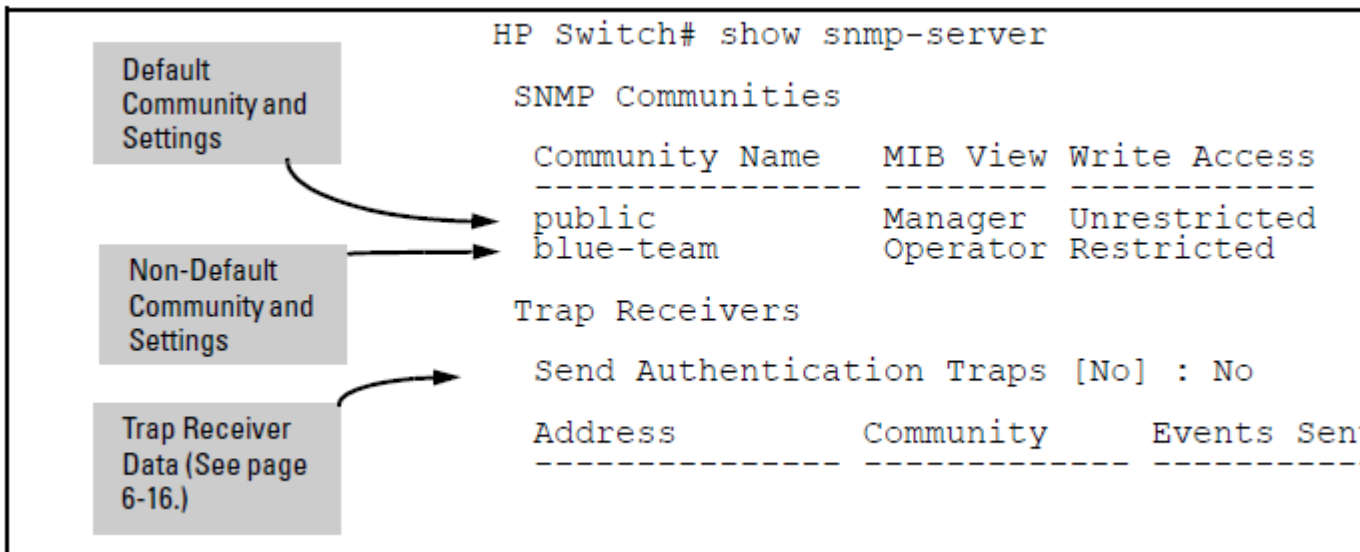
Syntax:

```
show snmp-server [ <community-string> ]
```

Example:

Lists the data for all communities in a switch; that is, both the default "public" community name and another community named "blue-team."

Figure 30 Example: of the SNMP community listing with two communities



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

```
HP Switch# show snmp-server public
```

Configuring community names and values (CLI)

The `snmp-server` command enables you to add SNMP communities with either default or specific access attributes, and to delete specific communities.

Syntax:

```
[no] snmp-server community <community-name>
```

Configures a new community name.

- If you do not also specify `operator` or `manager`, the switch automatically assigns the community to the `operator` MIB view.
- If you do not specify `restricted` or `unrestricted`, the switch automatically assigns the community to `restricted` (read-only) access.

The `no` form uses only the `<community-name>` variable and deletes the named community from the switch.

| | |
|-------------------------------|--|
| [operator manager] | <p>Optionally assigns an access level.</p> <ul style="list-style-type: none"> • At the <code>operator</code> level, the community can access all MIB objects except the CONFIG MIB. • At the <code>manager</code> level, the community can access all MIB objects. |
| [restricted unrestricted] | <p>Optionally assigns MIB access type.</p> <ul style="list-style-type: none"> • Assigning the <code>restricted</code> type allows the community to read MIB variables, but not to set them. • Assigning the <code>unrestricted</code> type allows the community to read and set MIB variables. |

Example:

To add the following communities:

| Community | Access Level | Type of Access |
|-----------|---|---------------------------|
| red-team | manager (Access to all MIB objects.) | unrestricted (read/write) |
| blue-team | operator (Access to all MIB objects except the CONFIG MIB.) | restricted (read-only) |

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

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

```
HP Switch(config) # no snmp-server community gold-team
```

SNMP notifications

The switches:

- Fixed or "Well-Known" Traps: A switch automatically sends fixed traps (such as "coldStart", "warmStart", "linkDown", and "linkUp") to trap receivers using the public community name, which is the default. These traps can also be sent to non-public communities.
- SNMPv2c informs
- SNMP v3 notification process, including traps

This section describes how to configure a switch to send network security and link-change notifications to configured trap receivers.

Supported Notifications

By default, the following notifications are enabled on a switch:

- Manager password changes
- SNMP authentication failure
- Link-change traps: when the link on a port changes from up to down (linkDown) or down to up (linkUp)
- Port-security (web, MAC, or 802.1X) authentication failure
- Invalid password entered in a login attempt through a direct serial, Telnet, or SSH connection
- Inability to establish a connection with the RADIUS or TACACS+ authentication server
- DHCP snooping events
- ARP protection events

General steps for configuring SNMP notifications

1. Determine the versions of SNMP notifications that you want to use in your network.
If you want to use SNMPv1 and SNMPv2c traps, you must also configure a trap receiver. See the following sections and follow the required configuration procedures:
 - “SNMPv1 and SNMPv2c Traps” (page 154)
 - “Configuring an SNMP trap receiver (CLI)” (page 155)
 - “Enabling SNMPv2c informs (CLI)” (page 157)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 157).
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 162).
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 162)
 - “Viewing SNMP notification configuration (CLI)” (page 164)

SNMPv1 and SNMPv2c Traps

The switches support the following functionality from earlier SNMP versions (SNMPv1 and SNMPv2c):

- **Trap receivers:** A *trap receiver* is a management station to which the switch sends SNMP traps and (optionally) event log messages sent from the switch. From the CLI you can configure up to ten SNMP trap receivers to receive SNMP traps from the switch.
- **Fixed or "Well-Known" Traps:** A switch automatically sends fixed traps (such as "coldStart", "warmStart", "linkDown", and "linkUp") to trap receivers using the `public` community name. These traps cannot be redirected to other communities. If you change or delete the default `public` community name, these traps are not sent.
- **Thresholds:** A switch automatically sends all messages created when a system threshold is reached to the network management station that configured the threshold, regardless of the trap receiver configuration.

SNMP trap receivers

Use the `snmp-server host` command to configure a trap receiver that can receive SNMPv1 and SNMPv2c traps, and (optionally) Event Log messages. When you configure a trap receiver, you specify its community membership, management station IP address, and (optionally) the type of Event Log messages to be sent.

If you specify a community name that does not exist—that is, has not yet been configured on the switch—the switch still accepts the trap receiver assignment. However, no traps are sent to that trap receiver until the community to which it belongs has been configured on the switch.

NOTE: To replace one community name with another for the same IP address, you must first enter the

`no snmp-server host <community-name> <ipv4-address | ipv6-address>`
 command to delete the unwanted community name. Otherwise, if you add a new community name with an IP address that is already used with a different community name, two valid community name entries are created for the same management station.

If you do not specify the event level (`[none | all | not-info | critical | debug]`), the switch does not send Event Log messages as traps. However, "well-known" traps and threshold traps (if configured) are still sent.

Configuring an SNMP trap receiver (CLI)

For information about configuring SNMP trap receivers, see "SNMP trap receivers" (page 154).

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 155)).</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 157).</p> |

Table 20 Security levels for Event Log messages sent as traps

| Security Level | Action |
|----------------|--|
| None (default) | Sends no Event Log messages. |
| All | Sends all Event Log messages. |
| Not-Info | Sends all Event Log messages that are not for information only. |
| Critical | Sends only Event Log messages for critical error conditions. |
| Debug | Sends only Event Log messages needed to troubleshoot network- and switch-level problems. |

Example:

To configure a trap receiver in a community named "red-team" with an IP address of 10.28.227.130 to receive only "critical" event log messages, you can enter the following command:

```
HP Switch(config)# snmp-server host 10.28.227.130 red-team critical
```

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 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
```

SNMPv2c informs

On a switch enabled for SNMPv2c, you can use the `snmp-server host inform` command (“Enabling SNMPv2c informs (CLI)” (page 157)) to send inform requests when certain events occur. When an SNMP Manager receives an inform request, it can send an SNMP response back to the sending agent on the switch to let the agent know that the inform request reached its destination.

If the sending agent on the switch does not receive an SNMP response back from the SNMP Manager within the timeout period, the inform request may be resent, based on the retry count value.

When you enable SNMPv2c inform requests to be sent, you must specify the IP address and community name of the management station that will receive the inform notification.

Enabling SNMPv2c informs (CLI)

For information about enabling SNMPv2c informs, see “SNMPv2c informs” (page 156).

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.

| | |
|---------|---|
| retries | Maximum number of times to resend an <code>inform</code> request if no SNMP response is received. (Default: 3) |
| timeout | 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 95 \(page 157\)](#) (note indication of `inform` Notify Type in bold below):

Example 95 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

Snm Response Pdu Source-IP Information

Selection Policy : Default rfc1517

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

Configuring SNMPv3 notifications (CLI)

The SNMPv3 notification process allows messages that are passed via SNMP between the switch and a network management station to be authenticated and encrypted.

1. Enable SNMPv3 operation on the switch by entering the `snmpv3 enable` command (See "SNMP Version 3 Commands" on page N-7).

When SNMPv3 is enabled, the switch supports:

- Reception of SNMPv3 notification messages (traps and informs)
 - Configuration of initial users
 - (Optional) Restriction of non-SNMPv3 messages to "read only"
2. Configure SNMPv3 users by entering the `snmpv3 user` command (see "SNMPv3 users" (page 146)). 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 148)).
 4. Define the name of an SNMPv3 notification configuration by entering the `snmpv3 notify` command.

Syntax:

```
[no] snmpv3 notify <notify_name> tagvalue <tag_name>
```

Associates the name of an SNMPv3 notification configuration with a tag name used (internally) in SNMPv3 commands. To delete a notification-to-tag mapping, enter `no snmpv3 notify notify_name`.

| | |
|---|---|
| <code>notify</code> <code><notify_name></code> | Specifies the name of an SNMPv3 notification configuration. |
| <code>tagvalue <tag_name></code> | Specifies the name of a tag value used in other SNMPv3 commands, such as <code>snmpv3 targetaddress params taglist tag_name</code> in Step 5. |

5. Configure the target address of the SNMPv3 management station to which SNMPv3 informs and traps are sent by entering the `snmpv3 targetaddress` command.

Syntax:

```
[no] snmpv3 targetaddress <ipv4-addr | ipv6-addr>  
<name>
```

Configures the IPv4 or IPv6 address, name, and configuration filename of the SNMPv3 management station to which notification messages are sent.

| | |
|---|---|
| <code>params <parms_name></code> | Name of the SNMPv3 station's parameters file. The parameters filename configured with <code>params parms_name</code> must match the <code>params parms_name</code> value entered with the <code>snmpv3 params</code> command in Step 6. |
| <code>taglist <tag_name> [tag_name]</code> ... | Specifies the SNMPv3 notifications (identified by one or more <code>tag_name</code> values) to be sent to the IP address of the SNMPv3 management station. You can enter more than one <code>tag_name</code> value. Each <code>tag_name</code> value must be already associated with the name of an SNMPv3 notification configuration entered with the <code>snmpv3 notify</code> command in Step 4. Use a blank space to separate <code>tag_name</code> values. You can enter up to 103 characters in <code>tag_name</code> entries following the <code>taglist</code> keyword. |

| | |
|---|--|
| [filter <none debug all not-info critical>] | (Optional) Configures the type of messages sent to a management station. (Default: none.) |
| [udp-port <port>] | (Optional) Specifies the UDP port to use. (Default: 162.) |
| [port-mask <mask>] | (Optional) Specifies a range of UDP ports. (Default: 0.) |
| [addr-mask <mask>] | (Optional) Specifies a range of IP addresses as destinations for notification messages. (Default: 0.) |
| [retries <value>] | (Optional) Number of times a notification is retransmitted if no response is received. Range: 1-255. (Default: 3.) |
| [timeout <value>] | (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).] |
| [max-msg-size <size>] | (Optional) Maximum number of bytes supported in a notification message to the specified target. (Default: 1472) |

6. Create a configuration record for the target address with the `snmpv3 params` command.

Syntax:

```
[no] snmpv3 params <params_name> user <user_name>
```

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

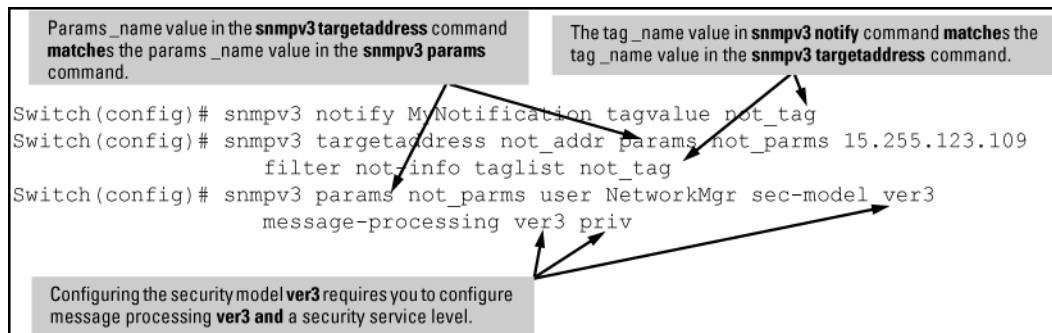
If you enter the `snmpv3 params user` command, you must also configure a security model (`sec-model`) and message processing algorithm (`msg-processing`).

| | |
|---|--|
| <sec-model [ver1 ver2c ver3>] | 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. If you configure the security model as <code>ver3</code> , you must also configure the message processing value as <code>ver3</code> . |
| msg-processing <ver1 ver2c ver3> [noauth auth priv] | Configures the algorithm used to process messages sent to the SNMPv3 target address. 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>). |

Example:

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

Figure 31 Example: of an SNMPv3 notification configuration



Network security notifications

By default, a switch is enabled to send the SNMP notifications listed in “Supported Notifications” (page 153) 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 155)
- “Configuring SNMPv3 notifications (CLI)” (page 157)

You can manage the default configuration of the switch to disable and re-enable notifications to be sent for the following types of security events:

- ARP protection events
- Inability to establish a connection with the RADIUS or TACACS+ authentication server
- DHCP snooping events
- Dynamic IP Lockdown hardware resources consumed
- Link change notification
- Invalid password entered in a login attempt through a direct serial, Telnet, or SSH connection
- Manager password changes
- Port-security (web, MAC, or 802.1X) authentication failure
- SNMP authentication failure
- Running configuration changes

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

For more information, see “Network security notifications” (page 160).

Syntax:

```
[no] snmp-server enable traps [ snmp-auth | password-change-mgr
| login-failure-mgr | port-security | auth-server-fail |
dhcp-snooping | arp-protect | running-config-change ]
```

Enables or disables sending one of the security notification types listed below to configured trap receivers. (Unless otherwise stated, all of the following notifications are enabled in the default configuration.)

The notification sends a trap:

| | |
|---|--|
| arp-protect | If ARP packets are received with an invalid source or destination MAC address, an invalid IP address, or an invalid IP-to-MAC binding. |
| auth-server-fail | If the connection with a RADIUS or TACACS+ authentication server fails. |
| dhcp-snooping | If DHCP packets are received from an untrusted source or if DHCP packets contain an invalid IP-to-MAC binding. |
| dyn-ip-lockdown | If the switch is out of hardware resources needed to program a dynamic IP lockdown rule |
| link-change <port-list> | When the link state on a port changes from up to down, or the reverse. |
| login-failure-mgr | For a failed login with a manager password. |
| password-change-mgr | When a manager password is reset. |
| mac-notify | Globally enables the generation of SNMP trap notifications upon MAC address table changes. |
| port-security | For a failed authentication attempt through a web, MAC, or 801.X authentication session. |
| running-config-change | When changes to the running configuration file are made. |
| snmp-authentication [extended standard] | For a failed authentication attempt via SNMP. (Default: extended.) |
| Startup-config-change | Sends a trap when changes to the startup configuration file are made. See "Enabling SNMP Traps on Startup Configuration Changes" on page 6-34. (Default: Disabled) |

To determine the specific cause of a security event, check the Event Log in the console interface to see why a trap was sent. For more information, see "Using the Event Log for Troubleshooting Switch Problems".

Viewing the current configuration for network security notifications (CLI)

Enter the `show snmp-server traps` command, as shown in [Example 96 \(page 162\)](#). Note that command output is a subset of the information displayed with the `show snmp-server` command in [Figure 32 \(page 165\)](#).

Example 96 Display of configured network security notifications

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

```
Trap Receivers
```

```
Link-Change Traps Enabled on Ports [All] : A1-A24
```

| Traps Category | Current Status |
|------------------------------|----------------|
| SNMP Authentication | : Extended |
| Password change | : Enabled |
| Login failures | : Enabled |
| Port-Security | : Enabled |
| Authorization Server Contact | : Enabled |
| DHCP Snooping | : Enabled |
| Dynamic ARP Protection | : Enabled |
| Dynamic IP Lockdown | : Enabled |

| Address | Community | Events Sent | Notify Type | Retry | Timeout |
|---|-----------|-------------|-------------|-------|---------|
| 15.255.5.225 | public | All | trap | 3 | 15 |
| 2001:0db8:0000:0001 :0000:0000:0000:0121 | user_1 | All | trap | 3 | 15 |

```
Excluded MIBs
```

Enabling Link-Change Traps (CLI)

By default, a switch is enabled to send a trap when the link state on a port changes from up to down (linkDown) or down to up (linkUp). To reconfigure the switch to send link-change traps to configured trap receivers, enter the `snmp-server enable traps link-change` command.

Syntax:

```
[no] snmp-server enable traps link-change <port-list> [ all ]
```

Enables or disables the switch to send a link-change trap to configured trap receivers when the link state on a port goes from up to down or down to up.

Enter `all` to enable or disable link-change traps on all ports on the switch.

Readable interface names in traps

The SNMP trap notification messages for linkup and linkdown events on an interface includes `IfDesc` and `IfAlias` var-bind information.

Source IP address for SNMP notifications

The switch uses an interface IP address as the source IP address in IP headers when sending SNMP notifications (traps and informs) or responses to SNMP requests.

For multi-netted interfaces, the source IP address is the IP address of the outbound interface of the SNMP reply, which may differ from the destination IP address in the IP header of the received request. For security reasons, it may be desirable to send an SNMP reply with the IP address of the destination interface (or a specified IP address) on which the corresponding SNMP request was received.

To configure the switch to use the source IP address on which an SNMP request was received in SNMP notification/traps and replies, enter the `snmp-server response-source` ((page 163)) and `snmp-server trap-source` ((page 163)) commands.

Configuring the source IP address for SNMP notifications (CLI)

For more information, see “Source IP address for SNMP notifications” (page 162).

Syntax:

```
[no] snmp-server response-source [ dst-ip-of-request [ ipv4-addr | ipv6-addr ] | loopback <0-7> ]
```

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

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

(Default: Interface IP address)

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

To use the IP address of the destination interface on which an SNMP request was received as the source IP address in the IP header of SNMP traps and replies, enter the following command:

```
HP Switch(config)# snmp-server response-source dst-ip-of-request
```

Syntax:

```
[no] snmp-server trap-source [ ipv4-addr | loopback <0-7> ]
```

Specifies the source IP address to be used for a trap PDU. To configure the switch to use a specified source IP address in generated trap PDUs, enter the `snmp-server trap-source` command.

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

(Default: Use the interface IP address in generated trap PDUs)

| | |
|-----------------------------------|---|
| <code>ipv4-addr</code> | User-defined interface IPv4 address that is used as the source IP address in generated traps. IPv6 addresses are not supported. |
| <code>loopback <0-7></code> | P address configured for the specified loopback interface that is used as the source IP address in a generated trap PDU. If multiple loopback IP addresses are configured, the lowest alphanumeric address is used. |

NOTE: When you use the `snmp-server response-source` and `snmp-server trap-source` commands, note the following behavior:

- The `snmp-server response-source` and `snmp-server trap-source` commands configure the source IP address for IPv4 interfaces only.
- You must manually configure the `snmp-server response-source` value if you wish to change the default user-defined interface IP address that is used as the source IP address in SNMP traps (RFC 1517).
- The values configured with the `snmp-server response-source` and `snmp-server trap-source` commands are applied globally to all interfaces that are sending SNMP responses or SNMP trap PDUs.
- Only the source IP address field in the IP header of the SNMP response PDU can be changed.
- Only the source IP address field in the IP header and the SNMPv1 Agent Address field of the SNMP trap PDU can be changed.

Verifying the configuration of the interface IP address used as the source IP address in IP headers for SNMP replies and traps sent from the switch (CLI)

Enter the `show snmp-server` command to display the SNMP policy configuration, as shown in [Example 97 \(page 164\)](#).

Example 97 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
Snm Response Pdu Source-IP Information
Selection Policy : dstIpOfRequest 1

Trap Pdu Source-IP Information
Selection Policy : Configured IP
```

- 1** `dstIpOfRequest`: The destination IP address of the interface on which an SNMP request is received is used as the source IP address in SNMP replies.

Viewing SNMP notification configuration (CLI)

Syntax:

```
show snmp-server
```

Displays the currently configured notification settings for versions SNMPv1 and SNMPv2c traps, including SNMP communities, trap receivers, link-change traps, and network security notifications.

Example:

In the following Example:, the `show snmp-server` command output shows that the switch has been configured to send SNMP traps and notifications to management stations that belong to the "public," "red-team," and "blue-team" communities.

Figure 32 Display of SNMP notification configuration

```

HP Switch(config)# show snmp-server

SNMP Communities
Community Name  MIB View Write Access
-----
public          Operator Restricted
blue-team      Manager Unrestricted
red-team       Manager Unrestricted

Trap Receivers

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

Trap Category          Current Trap Configuration
-----
SNMP Authentication   extended
Password change       enabled
Login failures         enabled
Port-Security          enabled
Authorization Server Contact enabled
ARP Protection         enabled
DHCP Snooping         enabled

Address      Community  Events Sent  Notify Type  Retry  Timeout
-----
10.28.227.200 public     All          trap         3      15
10.28.227.105 red-team   Critical     trap         3      15
10.28.227.120 blue-team  Not-INFO     trap         3      15
...

```

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 98 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 99 Information displayed for the `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 100 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 101 Information about SNMP traps, including MAC address count being Enabled/Disabled

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

```
Trap Receivers
```

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

```
Traps Category                Current Status
```

```
-----  
SNMP Authentication           : Extended  
Password change               : Enabled  
Login failures                : Enabled  
Port-Security                 : Enabled  
Authorization Server Contact  : Enabled  
DHCP-Snooping                : Enabled  
Dynamic ARP Protection        : Enabled  
Dynamic IP Lockdown           : Enabled
```

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

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

```
Excluded MIBs
```

1 The notify option is enabled.

Advanced management: RMON

The switch supports RMON (remote monitoring) on all connected network segments. This allows for troubleshooting and optimizing your network.

The following RMON groups are supported:

- Ethernet Statistics (except the numbers of packets of different frame sizes)
- Alarm
- History (of the supported Ethernet statistics)
- Event

The RMON agent automatically runs in the switch. Use the RMON management station on your network to enable or disable specific RMON traps and events. Note that you can access the Ethernet statistics, Alarm, and Event groups from the HP Switch Manager network management software. For more information on PCM+, see the HP Networking web site at www.hp.com/networking.

From the Products menu, select Network Management. Then click on PCM+ Network Management under the HP Network Management bar.

CLI-configured sFlow with multiple instances

sFlow can also be configured via the CLI for up to three distinct sFlow instances: once enabled, an sFlow receiver/destination can be independently configured for full flow-sampling and counter-polling. CLI-configured sFlow instances may be saved to the startup configuration to persist across a switch reboot.

Configuring sFlow (CLI)

The following sFlow commands allow you to configure sFlow instances via the CLI. For more information, see [“Advanced management: RMON” \(page 167\)](#).

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.

Configuring sFlow when destination is OOBM port

Syntax:

```
[no] sflow <receiver-instance> destination <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 102 sFlow destination is OOBM port

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

Example 103 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             : 6343
Owner                     : Administrator, CLI-Owned, Instance 1
Timeout (seconds)        : 2147479533
Max Datagram Size        : 1400
Datagram Version Support  : 5
OOBM Support              : Enabled
```

Example 104 Output of the running-config file showing the sFlow destination is the OOBM port

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

```
Running configuration:
```

```
; J9726A Configuration Editor; Created on release #WB.15.xx
; Ver #01:0d:0c
```

```
hostname "HP Switch"
module 1 type J8702A
module 7 type J870BA
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. [Example 106 \(page 170\)](#) is an Example: of sflow agent information.

Syntax:

```
show sflow agent
```

Displays sFlow agent information. The agent address is normally the IP address of the first VLAN configured.

The `show sflow agent` command displays read-only switch agent information. The version information shows the sFlow version, MIB support, and software versions; the agent address is typically the IP address of the first VLAN configured on the switch.

Example 105 Viewing sflow agent information

```
HP Switch# show sflow agent
```

```
Version          1.3;HP;XX.11.40
Agent Address    10.0.10.228
```

Syntax:

```
show sflow <receiver instance> destination
```

Displays information about the management station to which the sFlow sampling-polling data is sent.

The `show sflow instance destination` command includes information about the management-station's destination address, receiver port, and owner, as shown in [Example 106 \(page 170\)](#).

Example 106 Viewing sFlow destination information

```
HP Switch# show sflow 2 destination
```

```
Destination Instance      2
sflow                     Enabled
Datagrams Sent           221
Destination Address       10.0.10.41
Receiver Port             6343
Owner                     Administrator, CLI-owned, Instance 2
Timeout (seconds)        99995530
Max Datagram Size         1400
Datagram Version Support  5
```

Note the following details:

- **Destination Address** remains blank unless it has been configured.
- **Datagrams Sent** shows the number of datagrams sent by the switch agent to the management station since the switch agent was last enabled.
- **Timeout** displays the number of seconds remaining before the switch agent will automatically disable sFlow (this is set by the management station and decrements with time).
- **Max Datagram Size** shows the currently set value (typically a default value, but this can also be set by the management station).

Syntax:

```
show sflow <receiver instance> sampling-polling
<port-list/range>
```

Displays status information about sFlow sampling and polling.

The `show sflow instance sampling-polling [port-list]` command displays information about sFlow sampling and polling on the switch, as shown in [Figure 33 \(page 171\)](#). You can specify a list or range of ports for which to view sampling information.

Figure 33 Example: of viewing sFlow sampling and polling information

```

HP Switch# show sflow 2 sampling-polling A1-A4

```

Number denotes the sampling/polling instance to which the receiver is coupled.

| Port | Sampling | | Header | Dropped Samples | Polling | |
|------|----------|------|--------|-----------------|---------|----------|
| | Enabled | Rate | | | Enabled | Interval |
| A1 | Yes (2) | 40 | 128 | 1234567890 | --- | --- |
| A2 | --- | --- | --- | 0 | Yes (1) | 60 |
| A3 | No (1) | 0 | 100 | 898703 | No | 30 |
| A4 | Yes (3) | 50 | 128 | 0 | No (3) | 0 |

NOTE: The sampling and polling instances (noted in parentheses) coupled to a specific receiver instance are assigned dynamically, and so the instance numbers may not always match. The key thing to note is whether sampling or polling is enabled on a port, and the sampling rates or polling intervals for the receiver instance configured on each port.

Configuring UDLD Verify before forwarding

When an UDLD enabled port transitions to link-up, the port will begin with a UDLD blocking state. UDLD will probe via protocol packet exchange to determine the bidirectional state of the link. Until UDLD has completed the probe, all data traffic will be blocked. If the link is found to be bidirectional, UDLD will unblock the port for data traffic to pass. Once UDLD unblocks the port, other protocols will see the port as up and data traffic can be safely forwarded.

The default mode of a switch is “forward first then verify”. Enabling UDLD link-up will default to “forward first then verify”. To change the mode to “verify then forward”, you need to configure using the commands found in section 6.72.

NOTE: Link-UP data traffic will resumed after probing the link partner completes. All other protocols running will see the port as down.

UDLD time delay

UDLD protocol informs the link partner simultaneously as it detects a state change from unidirectional to bidirectional traffic. Additional packet exchanges will be carried out by UDLD in addition to the existing UDLD exchanges whenever state changes from unidirectional to bidirectional.

Table 21 Peer state transition timings

| Interval Time | Interval 1 | Interval 1 + delta | Interval 2 | Interval 3 |
|---------------------------|-------------------------------------|--|--|-----------------|
| | 5 sec | 5+(<5) sec* | 10 sec | 15 sec |
| With triggered updates | State = blockedPeer State = blocked | Inform PeerState = unblockedPeer State = unblocked | Regular UDLD TX | Regular UDLD TX |
| Without triggered updates | State = blockedPeer State = blocked | State = unblockedPeer State = blocked | Inform PeerState = unblockedPeer State = unblocked | Regular UDLD TX |

*delta is the time when the unblock event occurs on local side

Restrictions

- There is no support available when configuring this mode from the web and menu interface.
- There are no new packet types are introduced with UDLD.
- There are no new UDLD timers being introduced.

UDLD configuration commands

Syntax:

```
HP Switch(config)# link-keepalive mode [verify-then-forward  
| forward-then-verify]
```

This command configures the link-keepalive mode.

Link-keepalive provides two modes of operation; `verify-then-forward` and `forward-then-verify`.

When using the `verify-then-forward` mode, the port is in a blocking state until the link configured for UDLD establishes bidirectional communication. When using the `forward-then-verify` mode, the port forwards the data then verifies the status of the link-in state.

When a unidirectional state is detected, the port is moved to a blocked state.

When a bidirectional state is detected, the data is forwarded without interruption.

Syntax:

```
HP Switch(config)# link-keepalive mode verify-then-forward
```

Keeps the port in a logically blocked state until the link configured for UDLD has been successfully established in bi-directional communication.

Syntax:

```
HP Switch(config)# link-keepalive mode forward-then-verify
```

Forwards the data then verifies the status of the link. If a unidirectional state is detected, the port is then moved to a blocked state.

Syntax:

```
HP Switch(config)# link-keepalive interval <deciseconds>
```

Configure the interval for link-keepalive. The link-keepalive interval is the time between sending two UDLD packets. The time interval is entered in deciseconds (1/10 sec). The default keepalive interval is 50 deciseconds.

Example:

A value of 10 is 1 sec., 11 is 1.1 sec.

Syntax:

```
HP Switch(config)# link-keepalive retries <number>
```

Maximum number of sending attempts for UDLD packets before declaring the link as faulty.

Default keepalive attempt is 4.

Show commands

Syntax:

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

Sample output:

```
Total link-keepalive enabled ports: 8  
Keepalive Retries : 4  
Keepalive Interval: 5 sec
```

```
Keepalive Mode : verify-then-forward
Physical Keepalive Adjacent UDLD
```

| Port | Enabled | Status | Status | Switch | VLAN |
|------|---------|--------|----------|---------------|----------|
| 1 | Yes | down | off-line | 000000-000000 | untagged |
| 2 | Yes | down | off-line | 000000-000000 | untagged |
| 3 | Yes | down | off-line | 000000-000000 | untagged |
| 4 | Yes | down | off-line | 000000-000000 | untagged |
| 5 | Yes | down | off-line | 000000-000000 | untagged |
| 6 | Yes | down | off-line | 000000-000000 | untagged |
| 7 | Yes | down | off-line | 000000-000000 | untagged |
| 8 | Yes | down | off-line | 000000-000000 | untagged |

RMON generated when user changes UDLD mode

RMON events are generated when UDLD is configured. The first time you configure the mode, the UDLD states will be re-initialized. An event log entry is initiated to include the reason for the initial UDLD blocking state during link up.

Example:

UDLD mode [verify-then-forward | forward-then-verify] is configured

Severity: - Info.

LLDP

To standardize device discovery on all 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 to 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 188).

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 \(media-endpoint-discovery\)” \(page 188\)](#).

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 \(media-endpoint-discovery\)” \(page 188\)](#).

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 \(page 179\)](#)).

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 \(media-endpoint-discovery\)” \(page 188\)](#).

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 \(page 179\)](#)).

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 \(page 180\)](#)).

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 \(page 182\)](#)). Per-port configuration options include four modes:

- Transmit and receive (`tx_rx`): This is the default setting on all ports. It enables a given port to both transmit and receive LLDP packets and to store the data from received (inbound) LLDP packets in the switch's MIB.
- Transmit only (`txonly`): This setting enables a port to transmit LLDP packets that can be read by LLDP neighbors. However, the port drops inbound LLDP packets from LLDP neighbors without reading them. This prevents the switch from learning about LLDP neighbors on that port.
- Receive only (`rxonly`): This setting enables a port to receive and read LLDP packets from LLDP neighbors and to store the packet data in the switch's MIB. However, the port does not transmit outbound LLDP packets. This prevents LLDP neighbors from learning about the switch through that port.
- Disable (`disable`): This setting disables LLDP packet transmissions and reception on a port. In this state, the switch does not use the port for either learning about LLDP neighbors or informing LLDP neighbors of its presence.

SNMP notification

You can enable the switch to send a notification to any configured SNMP trap receiver(s) when the switch detects a remote LLDP data change on an LLDP-enabled port ([Configuring SNMP notification support \(page 182\)](#)).

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 \(page 183\)](#)).

Table 22 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, 6} | N/A | Always Enabled | Indicates the type of identifier used for Chassis ID. |
| Chassis ID ⁶ | N/A | Always Enabled | Uses base MAC address of the switch. |
| Port Type ^{3, 6} | N/A | Always Enabled | Uses "Local," meaning assigned locally by LLDP. |
| Port Id ⁶ | N/A | Always Enabled | Uses port number of the physical port. This is an internal number reflecting the reserved slot/port position in the chassis. For more information on this numbering scheme, see the appendix "MAC Address Management". |
| Remote Management Address | | | |
| Type ^{4, 6} | N/A | Always Enabled | Shows the network address type. |

Table 22 Data available for basic LLDP advertisements *(continued)*

| Data type | Configuration options | Default | Description |
|---|-----------------------|---|--|
| Address ⁴ | Default or Configured | Uses a default address selection method unless an optional address is configured. See “Remote management address” (page 176). | |
| System Name ⁶ | Enable/Disable | Enabled | Uses the switch's assigned name. |
| System Description ⁶ | Enable/Disable | Enabled | Includes switch model name and running software version, and ROM version. |
| Port Description ⁶ | Enable/Disable | Enabled | Uses the physical port identifier. |
| System capabilities supported ^{5, 6} | Enable/Disable | Enabled | Identifies the switch's primary capabilities (bridge, router). |
| System capabilities enabled ^{5, 6} | Enable/Disable | Enabled | Identifies the primary switch functions that are enabled, such as routing. |

¹ The Packet Time-to-Live value is included in LLDP data packets.

² Subelement of the Chassis ID TLV.

⁶ Populated with data captured internally by the switch. For more on these data types, refer to the IEEE P802.1AB Standard.

³ Subelement of the Port ID TLV.

⁴ Subelement of the Remote-Management-Address TLV.

⁵ 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 177).

Debug logging

You can enable LLDP debug logging to a configured debug destination (Syslog server, a terminal device, or both) by executing the `debug lldp` command. (For more information on Debug and Syslog, see the “Troubleshooting” appendix in this guide.) Note that the switch's Event Log does not record usual LLDP update messages.

Options for reading LLDP information collected by the switch

You can extract LLDP information from the switch to identify adjacent LLDP devices. Options include:

- Using the switch's `show lldp info` command options to display data collected on adjacent LLDP devices—as well as the local data the switch is transmitting to adjacent LLDP devices (“Displaying the global LLDP, port admin, and SNMP notification status (CLI)” (page 177)).
- 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 188).)

LLDP operating rules

For additional information specific to LLDP-MED operation, see [“LLDP-MED \(media-endpoint-discovery\)”](#) (page 188).

Port trunking

LLDP manages trunked ports individually. That is, trunked ports are configured individually for LLDP operation, in the same manner as non-trunked ports. Also, LLDP sends separate advertisements on each port in a trunk, and not on a per-trunk basis. Similarly, LLDP data received through trunked ports is stored individually, per-port.

IP address advertisements

In the default operation, if a port belongs to only one static VLAN, the port advertises the lowest-order IP address configured on that VLAN. If a port belongs to multiple VLANs, the port advertises the lowest-order IP address configured on the VLAN with the lowest VID. If the qualifying VLAN does not have an IP address, the port advertises 127.0.0.1 as its IP address. For example, if the port is a member of the default VLAN (VID=1), and there is an IP address configured for the default VLAN, the port advertises this IP address. In the default operation, the IP address that LLDP uses can be an address acquired by DHCP or Bootp.

You can override the default operation by configuring the port to advertise any IP address that is manually configured on the switch, even if the port does not belong to the VLAN configured with the selected IP address ([“Configuring a remote management address for outbound LLDP advertisements \(CLI\)”](#) (page 183)). (Note that LLDP cannot be configured through the CLI to advertise an addresses acquired through DHCP or Bootp. However, as mentioned above, in the default LLDP configuration, if the lowest-order IP address on the VLAN with the lowest VID for a given port is a DHCP or Bootp address, the switch includes this address in its LLDP advertisements unless another address is configured for advertisements on that port.) Also, although LLDP allows configuring multiple remote management addresses on a port, only the lowest-order address configured on the port will be included in outbound advertisements. Attempting to use the CLI to configure LLDP with an IP address that is either not configured on a VLAN or has been acquired by DHCP or Bootp results in the following error message.

```
xxx.xxx.xxx.xxx: This IP address is not configured or is a DHCP address.
```

Spanning-tree blocking

Spanning tree does not prevent LLDP packet transmission or receipt on STP-blocked links.

802.1X blocking

Ports blocked by 802.1X operation do not allow transmission or receipt of LLDP packets.

Configuring LLDP operation

Displaying the global LLDP, port admin, and SNMP notification status (CLI)

In the default configuration, LLDP is enabled and in both transmit and receive mode on all active ports. The LLDP configuration includes global settings that apply to all active ports on the switch, and per-port settings that affect only the operation of the specified ports.

The commands in this section affect both LLDP and LLDP-MED operation. for information on operation and configuration unique to LLDP-MED, refer to [“LLDP-MED \(Media-Endpoint-Discovery\)”](#).

Syntax:

```
show lldp config
```

Displays the LLDP global configuration, LLDP port status, and SNMP notification status. For information on port admin status, see [“Configuring per-port transmit and receive modes \(CLI\)”](#) (page 182).

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

Example 107 Viewing the general LLDP configuration

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

LLDP Global Configuration

```
LLDP Enabled [Yes] : Yes
LLDP Transmit Interval [30] : 30
LLDP Hold time Multiplier [4] : 4
LLDP Delay Interval [2] : 2
LLDP Reinit Interval [2] : 2
LLDP Notification Interval [5] : 5
LLDP Fast Start Count [5] : 5
```

LLDP Port Configuration

| Port | AdminStatus | NotificationEnabled | Med Topology Trap Enabled |
|------|-------------|---------------------|---------------------------|
| A1 | Tx_Rx | False | False |
| A2 | Tx_Rx | False | False |
| A3 | Tx_Rx | False | False |
| A4 | Tx_Rx | False | False |
| A5 | Tx_Rx | False | False |
| A6 | Tx_Rx | False | False |
| A7 | Tx_Rx | False | False |
| A8 | Tx_Rx | False | False |

NOTE: The values displayed in the LLDP column correspond to the `lldp refresh-interval` command

Viewing port configuration details (CLI)

Syntax:

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

Displays the LLDP port-specific configuration for all ports in `<port-list>`, including which optional TLVs and any non-default IP address that are included in the port's outbound advertisements.

For information on the notification setting, see [“Configuring SNMP notification support”](#) (page 182). For information on the other configurable settings displayed by this command, see [“Configuring per-port transmit and receive modes \(CLI\)”](#) (page 182).

Figure 34 Per-port configuration display

```
HP Switch(config)# show lldp config 1

LLDP Port Configuration Detail

Port : 1
AdminStatus [Tx_Rx] : Tx_Rx
NotificationEnabled [False] : False
Med Topology Trap Enabled [False] : False

TLVS Advertised:
* port_descr
* system_name
* system_descr
* system_cap

[* capabilities]
[* network_policy]
[* location_id]
[* poe]
[* macphy_config]

IpAddress Advertised:
```

These fields appear when medtlvenable is enabled on the switch, which is the default setting.

This field appears when dot3tlvenable is enabled on the switch, which is the default setting.

The blank IpAddress field indicates that the default IP address will be advertised from this port.

Configuring Global LLDP Packet Controls

The commands in this section configure the aspects of LLDP operation that apply the same to all ports in the switch.

LLDP operation on the switch

Enabling LLDP operation (the default) causes the switch to:

- Use active, LLDP-enabled ports to transmit LLDP packets describing itself to neighbor devices.
- Add entries to its neighbors table based on data read from incoming LLDP advertisements.

Enabling or disabling LLDP operation on the switch (CLI)

For more information, see “LLDP operation on the switch” (page 179).

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 108 Disabling LLDP

```
HP Switch(config)# no lldp run
```

Changing the packet transmission interval (CLI)

This interval controls how often active ports retransmit advertisements to their neighbors.

Syntax:

```
lldp refresh-interval <5-32768>
```

Changes the interval between consecutive transmissions of LLDP advertisements on any given port.

(Default: 30 seconds)

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

Time-to-Live for transmitted advertisements

The Time-to-Live value (in seconds) for all LLDP advertisements transmitted from a switch is controlled by the switch that generates the advertisement and determines how long an LLDP neighbor retains the advertised data before discarding it. The Time-to-Live value is the result of multiplying the `refresh-interval` by the `holdtime-multiplier`.

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

For more information, see “Time-to-Live for transmitted advertisements” (page 180).

Syntax:

```
lldp holdtime-multiplier <2-10>
```

Changes the multiplier an LLDP switch uses to calculate the Time-to-Live for the LLDP advertisements it generates and transmits to LLDP neighbors. When the Time-to-Live for a given advertisement expires, the advertised data is deleted from the neighbor switch's MIB.

(Default: 4; Range 2–10)

Example:

If the `refresh-interval` on the switch is 15 seconds and the `holdtime-multiplier` is at the default, the Time-to-Live for advertisements transmitted from the switch is 60 seconds (4×15).

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

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

Delay interval between advertisements generated by value or status changes to the LLDP MIB

The switch uses a `delay-interval` setting to delay transmitting successive advertisements resulting from these LLDP MIB changes. If a switch is subject to frequent changes to its LLDP MIB, lengthening this interval can reduce the frequency of successive advertisements. You can change the `delay-interval` by using either an SNMP network management application or the CLI `setmib` command.

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

Syntax:

```
setmib lldpTxDelay.0 -i <1-8192>
```

Uses `setmib` to change the minimum time (`delay-interval`) any LLDP port will delay advertising successive LLDP advertisements because of a change in LLDP MIB content.

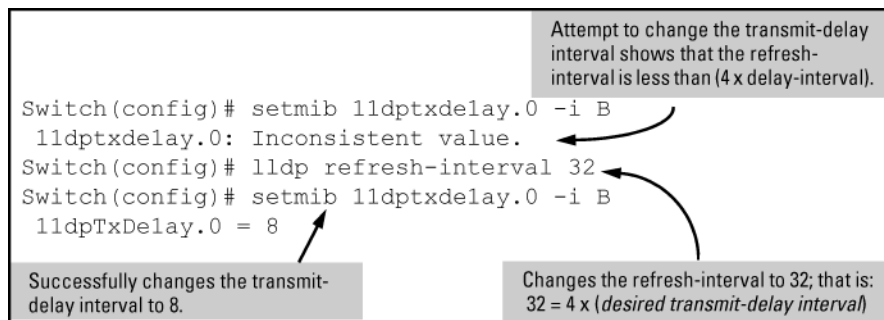
(Default: 2; Range 1–8192)

NOTE: The LLDP refresh-interval (transmit interval) must be greater than or equal to (4 x delay-interval). The switch does not allow increasing the delay interval to a value that conflicts with this relationship. That is, the switch displays `Inconsistent value` if (4 x delay-interval) exceeds the current transmit interval, and the command fails. Depending on the current refresh-interval setting, it may be necessary to increase the refresh-interval before using this command to increase the delay-interval.

Example:

To change the delay-interval from 2 seconds to 8 seconds when the refresh-interval is at the default 30 seconds, you must first set the refresh-interval to a minimum of 32 seconds ($32 = 4 \times 8$). (See Figure 35 (page 181).)

Figure 35 Changing the transmit-delay interval



Reinitialization delay interval

In the default configuration, a port receiving a `disable` command followed immediately by a `txonly`, `rxonly`, or `tx_rx` command delays reinitializing for two seconds, during which LLDP operation remains disabled. If an active port is subjected to frequent toggling between the LLDP disabled and enabled states, LLDP advertisements are more frequently transmitted to the neighbor device. Also, the neighbor table in the adjacent device changes more frequently as it deletes, then replaces LLDP data for the affected port which, in turn, generates SNMP traps (if trap receivers and SNMP notification are configured). All of this can unnecessarily increase network traffic. Extending the reinitialization-delay interval delays the ability of the port to reinitialize and generate LLDP traffic following an LLDP disable/enable cycle.

Changing the reinitialization delay interval (CLI)

For more information, see “Reinitialization delay interval” (page 181).

Syntax:

```
setmib lldpReinitDelay.0 -i <1-10>
```

Uses `setmib` to change the minimum time (reinitialization delay interval) an LLDP port will wait before reinitializing after receiving an LLDP disable command followed closely by a `txonly` or `tx_rx` command. The delay interval commences with execution of the `lldp admin-status port-list disable` command.

(Default: 2 seconds; Range 1–10 seconds)

Example:

The following command changes the reinitialization delay interval to five seconds:

```
HP Switch(config)# setmib lldpreinitdelay.0 -i 5
```

Configuring SNMP notification support

You can enable SNMP trap notification of LLDP data changes detected on advertisements received from neighbor devices, and control the interval between successive notifications of data changes on the same neighbor.

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

For more information, see Section 1.67.3.2.

Syntax:

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

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

(Default: Disabled)

For information on configuring trap receivers in the switch, see “SNMP notifications” (page 153).

Example:

This command enables SNMP notification on ports 1 - 5:

```
HP Switch(config)# lldp enable-notification 1-5
```

Changing the minimum interval for successive data change notifications for the same neighbor

If LLDP trap notification is enabled on a port, a rapid succession of changes in LLDP information received in advertisements from one or more neighbors can generate a high number of traps. To reduce this effect, you can globally change the interval between successive notifications of neighbor data change.

Syntax:

```
setmib lldpnotificationinterval.0 -i <1-3600>
```

Globally changes the interval between successive traps generated by the switch. If multiple traps are generated in the specified interval, only the first trap is sent. The remaining traps are suppressed. (A network management application can periodically check the switch MIB to detect any missed change notification traps. See IEEE P802.1AB or later for more information.)

(Default: 5 seconds)

Example:

The following command limits change notification traps from a particular switch to one per minute.

```
HP Switch(config)# setmib lldpnotificationinterval.0 -i 60 lldpNotificationInterval.0=60
```

Configuring per-port transmit and receive modes (CLI)

Syntax:

```
lldp admin-status <port-list> <txonly | rxonly | tx_rx |  
disable>
```

With LLDP enabled on the switch in the default configuration, each port is configured to transmit and receive LLDP packets. These options enable you to control which

ports participate in LLDP traffic and whether the participating ports allow LLDP traffic in only one direction or in both directions.

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

Basic LLDP per-port advertisement content

In the default LLDP configuration, outbound advertisements from each port on the switch include both mandatory and optional data.

Mandatory Data

An active LLDP port on the switch always includes the mandatory data in its outbound advertisements. LLDP collects the mandatory data, and, except for the Remote Management Address, you cannot use LLDP commands to configure the actual data.

- Chassis Type (TLV subelement)
- Chassis ID (TLV)
- Port Type (TLV subelement)
- Port ID (TLV)
- Remote Management Address (TLV; actual IP address is a subelement that can be a default address or a configured address)

Configuring a remote management address for outbound LLDP advertisements (CLI)

This is an optional command you can use to include a specific IP address in the outbound LLDP advertisements for specific ports. For more information, see [“Basic LLDP per-port advertisement content” \(page 183\)](#).

Syntax:

```
[no] lldp config <port-list> ipAddrEnable <ip-address>
```

Replaces the default IP address for the port with an IP address you specify. This can be any IP address configured in a static VLAN on the switch, even if the port does not belong to the VLAN configured with the selected IP address.

The `no` form of the command deletes the specified IP address.

If there are no IP addresses configured as management addresses, the IP address selection method returns to the default operation.

Default: The port advertises the IP address of the lowest-numbered VLAN (VID) to which it belongs. If there is no IP address configured on the VLANs to which the port belongs, and if the port is not configured to advertise an IP address from any other (static) VLAN on the switch, the port advertises an address of 127.0.0.1.)

NOTE: This command does not accept either IP addresses acquired through DHCP or Bootp, or IP addresses that are not configured in a static VLAN on the switch.

Example:

If port 3 belongs to a subnetted VLAN that includes an IP address of 10.10.10.100 and you want port 3 to use this secondary address in LLDP advertisements, you need to execute the following command:

```
HP Switch(config)# lldp config 3 ipAddrEnable 10.10.10.100
```

Syntax:

```
[no] lldp config <port-list> basicTlvEnable <TLV-Type>
```

| | |
|--------------|---|
| port_descr | For outbound LLDP advertisements, this TLV includes an alphanumeric string describing the port. (Default: Enabled) |
| system_name | For outbound LLDP advertisements, this TLV includes an alphanumeric string showing the assigned name of the system. (Default: Enabled) |
| system_descr | For outbound LLDP advertisements, this TLV includes an alphanumeric string describing the full name and version identification for the hardware type, software version, and networking application of the system. (Default: Enabled) |
| system_cap | For outbound advertisements, this TLV includes a bitmask of supported system capabilities (device functions). Also includes information on whether the capabilities are enabled. (Default: Enabled) |

Example:

If you want to exclude the system name TLV from the outbound LLDP advertisements for all ports on a switch, use this command:

```
HP Switch(config)# no lldp config 1-24 basicTlvEnable system_name
```

If you later decide to reinstate the system name TLV on ports 1-5, use this command:

```
HP Switch(config)# lldp config 1-5 basicTlvEnable system_name
```

Optional Data

You can configure an individual port or group of ports to exclude one or more of the following data types from outbound LLDP advertisements.

- Port description (TLV)
- System name (TLV)
- System description (TLV)
- System capabilities (TLV)
 - System capabilities Supported (TLV subelement)
 - System capabilities Enabled (TLV subelement)
- Port speed and duplex (TLV subelement)

Optional data types, when enabled, are populated with data internal to the switch; that is, you cannot use LLDP commands to configure their actual content.

Support for port speed and duplex advertisements

This feature is optional for LLDP operation, but is *required* for LLDP-MED operation.

Port speed and duplex advertisements are supported on the switches to inform an LLDP endpoint and the switch port of each other's port speed and duplex configuration and capabilities. Configuration mismatches between a switch port and an LLDP endpoint can result in excessive collisions and voice quality degradation. LLDP enables discovery of such mismatches by supporting SNMP access to the switch MIB for comparing the current switch port and endpoint settings. (Changing a current device configuration to eliminate a mismatch requires intervention by the system operator.)

An SNMP network management application can be used to compare the port speed and duplex data configured in the switch and advertised by the LLDP endpoint. You can also use the CLI to display this information. For more information on using the CLI to display port speed and duplex information, see [“Viewing the current port speed and duplex configuration on a switch port” \(page 199\)](#).

Configuring support for port speed and duplex advertisements (CLI)

For more information, see [“Support for port speed and duplex advertisements” \(page 184\)](#).

Syntax:

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

For outbound advertisements, this TLV includes the (local) switch port's current speed and duplex settings, the range of speed and duplex settings the port supports, and the method required for reconfiguring the speed and duplex settings on the device (autonegotiation during link initialization, or manual configuration).

Using SNMP to compare local and remote information can help in locating configuration mismatches.

(Default: Enabled)

NOTE: For LLDP operation, this TLV is optional. For LLDP-MED operation, this TLV is mandatory.

Port VLAN ID TLV support on LLDP

The `port-vlan-id` option enables advertisement of the port VLAN ID TLV as part of the regularly advertised TLVs. This allows discovery of a mismatch in the configured native VLAN ID between LLDP peers. The information is visible using `show` commands and is logged to the Syslog server.

Configuring the VLAN ID TLV

This TLV advertisement is enabled by default. To enable or disable the TLV, use this command. For more information, see [“Port VLAN ID TLV support on LLDP” \(page 185\)](#).

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 36 Enabling the VLAN ID TLV

```
HP Switch(config)# lldp config a1 dot1TlvEnable port-vlan-id
```

Viewing the TLVs advertised

The show commands display the configuration of the TLVs. The command `show lldp config` lists the TLVs advertised for each port, as shown in [Example 109 \(page 187\)](#) through [Example 111 \(page 188\)](#).

Example 109 Displaying the TLVs for a port

```
HP Switch(config)# show lldp config a1
```

```
LLDP Port Configuration Detail
```

```
Port      : A1  
AdminStatus [Tx_Rx] : Tx_Rx  
NotificationEnabled [False] : False  
Med Topology Trap Enabled [False] : False
```

```
TLVS Advertised:
```

```
* port_descr  
* system_name  
* system_descr  
* system_cap  
  
* capabilities  
* network_policy  
* location_id  
* poe  
  
* macphy_config  
  
* port_vlan_id 1
```

```
IpAddress Advertised:  
:  
:
```

1 The VLAN ID TLV is being advertised.

Example 110 Local device LLDP information

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

```
LLDP Port Configuration Information Detail
```

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

```
Port VLAN ID : 1 1
```

1 The information that LLDP used in its advertisement.

Example 111 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 XX.13.03, ROM...
PortDescr      : A1
```

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

```
Port VLAN ID : 200
```

```
Remote Management Address
```

```
Type      : ipv4
Address   : 192.168.1.1
```

SNMP support

The LLDP-EXT-DOT1-MIB has the corresponding MIB variables for the Port VLAN ID TLV. The TLV advertisement can be enabled or disabled using the MIB object

`lldpXdot1ConfigPortVlanTxEnable` in the `lldpXdot1ConfigPortVlanTable`.

The port VLAN ID TLV local information can be obtained from the MIB object

`lldpXdot1LocPortVlanId` in the local information table `lldpXdot1LocTable`.

The port VLAN ID TLV information about all the connected peer devices can be obtained from the MIB object `lldpXdot1RemPortVlanId` in the remote information table `lldpXdot1RemTable`.

LLDP-MED (media-endpoint-discovery)

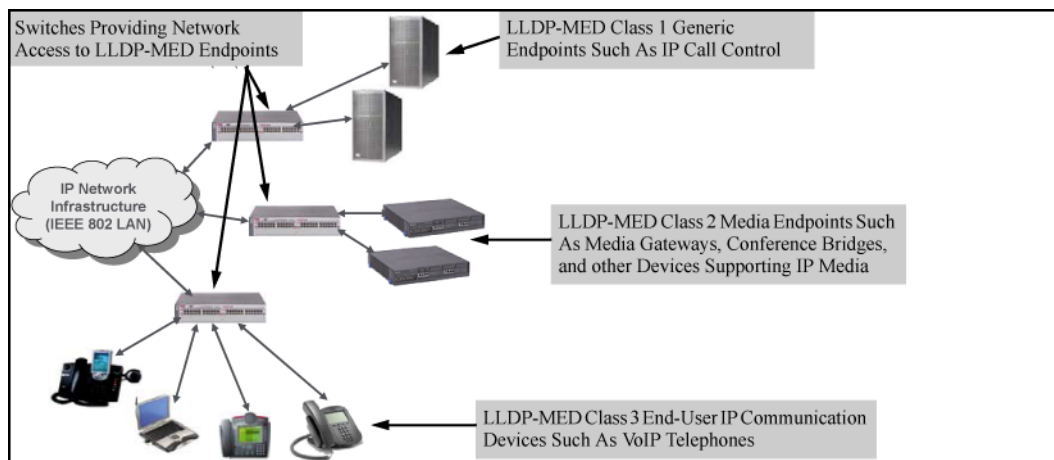
LLDP-MED (ANSI/TIA-1057/D6) extends the LLDP (IEEE 802.1AB) industry standard to support advanced features on the network edge for Voice Over IP (VoIP) endpoint devices with specialized capabilities and LLDP-MED standards-based functionality. LLDP-MED in the switches uses the standard LLDP commands described earlier in this section, with some extensions, and also introduces new commands unique to LLDP-MED operation. The `show` commands described elsewhere in this section are applicable to both LLDP and LLDP-MED operation. LLDP-MED benefits include:

- Plug-and-play provisioning for MED-capable, VoIP endpoint devices
- Simplified, vendor-independent management enabling different IP telephony systems to interoperate on one network
- Automatic deployment of convergence network policies (voice VLANs, Layer 2/CoS priority, and Layer 3/QoS priority)
- Configurable endpoint location data to support the Emergency Call Service (ECS) (such as Enhanced 911 service, 999, 112)
- Detailed VoIP endpoint data inventory readable via SNMP from the switch
- Power over Ethernet (PoE) status and troubleshooting support via SNMP
- support for IP telephony network troubleshooting of call quality issues via SNMP

This section describes how to configure and use LLDP-MED features in the switches to support VoIP network edge devices (media endpoint devices) such as:

- IP phones
- Voice/media gateways
- Media servers
- IP communications controllers
- Other VoIP devices or servers

Figure 37 Example: of LLDP-MED network elements



LLDP-MED endpoint support

LLDP-MED interoperates with directly connected IP telephony (endpoint) clients having these features and services:

- Autonegotiate speed and duplex configuration with the switch
- Use the following network policy elements configured on the client port
 - Voice VLAN ID
 - 802.1p (Layer 2) QoS
 - Diffserv codepoint (DSCP) (Layer 3) QoS
- Discover and advertise device location data learned from the switch
- Support ECS (such as E911, 999, and 112)
- Advertise device information for the device data inventory collected by the switch, including:

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

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

NOTE: LLDP-MED is intended for use with VoIP endpoints and is not designed to support links between network infrastructure devices, such as switch-to-switch or switch-to-router links.

LLDP-MED endpoint device classes

LLDP-MED endpoint devices are, by definition, located at the network edge and communicate using the LLDP-MED framework. Any LLDP-MED endpoint device belongs to one of the following three classes:

- Class 1 (generic endpoint devices): These devices offer the basic LLDP discovery services, network policy advertisement (VLAN ID, Layer 2/802.1p priority, and Layer 3/DSCP priority), and PoE management. This class includes such devices as IP call controllers and communication-related servers.
- Class 2 (media endpoint devices): These devices offer all Class 1 features plus media-streaming capability, and include such devices as voice/media gateways, conference bridges, and media servers.
- Class 3 (communication devices): These devices are typically IP phones or end-user devices that otherwise support IP media and offer all Class 1 and Class 2 features, plus location identification and emergency 911 capability, Layer 2 switch support, and device information management.

LLDP-MED operational support

The switches offer two configurable TLVs supporting MED-specific capabilities:

- `medTlvEnable` (for per-port enabling or disabling of LLDP-MED operation)
- `medPortLocation` (for configuring per-port location or emergency call data)

NOTE: LLDP-MED operation also requires the port speed and duplex TLV (`dot3TlvEnable`), which is enabled in the default configuration.

Topology change notifications provide one method for monitoring system activity. However, because SNMP normally employs UDP, which does not guarantee datagram delivery, topology change notification should not be relied upon as the sole method for monitoring critical endpoint device connectivity.

LLDP-MED fast start control

Syntax:

```
lldp fast-start-count <1-10>
```

An LLDP-MED device connecting to a switch port may use the data contained in the MED TLVs from the switch to configure itself. However, the `lldp refresh-interval` setting (default: 30 seconds) for transmitting advertisements can cause an unacceptable delay in MED device configuration.

To support rapid LLDP-MED device configuration, the `lldp fast-start-count` command temporarily overrides the `refresh-interval` setting for the `fast-start-count` advertisement interval. This results in the port initially advertising LLDP-MED at a faster rate for a limited time. Thus, when the switch detects a new LLDP-MED device on a port, it transmits one LLDP-MED advertisement per second out the port for the duration of the `fast-start-count` interval. In most cases, the default setting should provide an adequate `fast-start-count` interval.

(Default: 5 seconds)

NOTE: This global command applies only to ports on which a new LLDP-MED device is detected. It does not override the `refresh-interval` setting on ports where non-MED devices are detected.

Advertising device capability, network policy, PoE status and location data

The `medTlvEnable` option on the switch is enabled in the default configuration and supports the following LLDP-MED TLVs:

- LLDP-MED capabilities: This TLV enables the switch to determine:
 - Whether a connected endpoint device supports LLDP-MED
 - Which specific LLDP-MED TLVs the endpoint supports
 - The device class (1, 2, or 3) for the connected endpoint

This TLV also enables an LLDP-MED endpoint to discover what LLDP-MED TLVs the switch port currently supports.

- Network policy operating on the port to which the endpoint is connected (VLAN, Layer 2 QoS, Layer 3 QoS)
- PoE (MED Power-over-Ethernet)
- Physical location data (see [Configuring location data for LLDP-MED devices \(page 194\)](#))

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 \(CLI\)” \(page 185\)](#)).

Network policy advertisements

Network policy advertisements are intended for real-time voice and video applications, and include these TLV subelements:

- Layer 2 (802.1p) QoS
- Layer 3 DSCP (diffserv code point) QoS
- Voice VLAN ID (VID)

VLAN operating rules

These rules affect advertisements of VLANs in network policy TLVs:

- The VLAN ID TLV subelement applies only to a VLAN configured for voice operation (`vlan <vid> voice`).
- If there are multiple voice VLANs configured on a port, LLDP-MED advertises the voice VLAN having the lowest VID.
- The voice VLAN port membership configured on the switch can be tagged or untagged. However, if the LLDP-MED endpoint expects a tagged membership when the switch port is configured for untagged, or the reverse, a configuration mismatch results. (Typically, the endpoint expects the switch port to have a tagged voice VLAN membership.)
- If a given port does not belong to a voice VLAN, the switch does not advertise the VLAN ID TLV through this port.

Policy elements

These policy elements may be statically configured on the switch or dynamically imposed during an authenticated session on the switch using a RADIUS server and 802.1X or MAC authentication. (Web authentication does not apply to VoIP telephones and other telecommunications devices that are not capable of accessing the switch through a Web browser.) The QoS and voice VLAN policy elements can be statically configured with the following CLI commands:

```
vlan <vid> voice
```

```
vlan <vid> <tagged | untagged> <port-list>
int <port-list> qos priority <0-7>
vlan <vid> qos dscp <codepoint>
```

NOTE: A codepoint must have an 802.1p priority before you can configure it for use in prioritizing packets by VLAN-ID. If a codepoint you want to use shows `No Override` in the `Priority` column of the DSCP policy table (display with `show qos-dscp map`, then use `qos-dscp map <codepoint> priority <0-7>` to configure a priority before proceeding. For more information on this topic, see the chapter "Quality of Service (QoS): Managing Bandwidth More Effectively" in the *Advanced Traffic Management Guide* for your switch.

Enabling or Disabling medTlvEnable

In the default LLDP-MED configuration, the TLVs controlled by `medTlvEnable` are enabled. For more information, see ["Advertising device capability, network policy, PoE status and location data"](#) (page 191).

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 194).)
- Current PoE status TLV

(Default: All of the above TLVs are enabled.)

Helps to locate configuration mismatches by allowing use of an SNMP application to compare the LLDP-MED configuration on a port with the LLDP-MED TLVs advertised by a neighbor connected to that port.

| | |
|----------------|---|
| capabilities | <p>This TLV enables the switch to determine:</p> <ul style="list-style-type: none"> • Which LLDP-MED TLVs a connected endpoint can discover • The device class (1, 2, or 3) for the connected endpoint <p>This TLV also enables an LLDP-MED endpoint to discover what LLDP-MED TLVs the switch port currently supports.</p> <p>(Default: enabled)</p> <p>NOTE: This TLV cannot be disabled unless the <code>network_policy</code>, <code>poe</code>, and <code>location_id</code> TLVs are already disabled.</p> |
| network-policy | <p>This TLV enables the switch port to advertise its configured network policies (voice VLAN, Layer 2 QoS, Layer 3 QoS), and allows LLDP-MED endpoint devices to autoconfigure the voice network policy advertised by the switch. This also enables the use of SNMP applications to troubleshoot statically configured endpoint network policy mismatches.</p> <p>(Default: Enabled)</p> <p>NOTE: Network policy is advertised only for ports that are configured as members of the voice VLAN. If the port belongs to more than one voice VLAN, the voice VLAN with the lowest-numbered VID is selected as the VLAN for voice traffic. Also, this TLV cannot be enabled unless the <code>capability</code> TLV is already enabled.</p> <p>For more information, see "Network policy advertisements" (page 191).</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 194).</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 193).</p> |

PoE advertisements

These advertisements inform an LLDP-MED endpoint of the power (PoE) configuration on switch ports. Similar advertisements from an LLDP-MED endpoint inform the switch of the endpoint's power needs and provide information that can be used to identify power priority mismatches.

PoE TLVs include the following power data:

- **Power type:** indicates whether the device is a power-sourcing entity (PSE) or a PD. Ports on the J8702A PoE zl module are PSE devices. A MED-capable VoIP telephone is a PD.
- **Power source:** indicates the source of power in use by the device. Power sources for PDs include PSE, local (internal), and PSE/local. The switches advertise Unknown.
- **Power priority:** indicates the power priority configured on the switch (PSE) port or the power priority configured on the MED-capable endpoint.
- **Power value:** indicates the total power in watts that a switch port (PSE) can deliver at a particular time, or the total power in watts that the MED endpoint (PD) requires to operate.

Viewing PoE advertisements

To display the current power data for an LLDP-MED device connected to a port, use the following command:

```
show lldp info remote-device <port-list>
```

For more information on this command, see page A-60.

To display the current PoE configuration on the switch, use the following commands:

```
show power brief <port-list>
show power <port-list>
```

For more information on PoE configuration and operation, see Chapter 11, "Power Over Ethernet (PoE/PoE+) Operation".

Location data for LLDP-MED devices

You can configure a switch port to advertise location data for the switch itself, the physical wall-jack location of the endpoint (recommended), or the location of a DHCP server supporting the switch, endpoint, or both. You also have the option of configuring these different address types:

- **Civic address:** physical address data such as city, street number, and building information
- **ELIN (Emergency Location Identification Number):** an emergency number typically assigned to MLTS (Multiline Telephone System) Operators in North America
- **Coordinate-based location:** attitude, longitude, and altitude information (Requires configuration via an SNMP application.)

Configuring location data for LLDP-MED devices

For more information, see “Location data for LLDP-MED devices” (page 193).

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: <ul style="list-style-type: none">0: Location of DHCP server1: Location of switch2: Location of LLDP-MED endpoint (recommended application) This field is required in a <code>civic-addr</code> command. |
| Type/Value Pairs (CA-TYPE and CA-VALUE) | A series of data pairs, each composed of a location data "type" specifier and the corresponding location data for that type. That is, the first value in a pair is expected to be the civic address "type" number (<code>CA-TYPE</code>), and the second value in a pair is expected to be the corresponding civic address data (<code>CA-VALUE</code>). For example, if the <code>CA-TYPE</code> for "city name" is "3," the type/value pair to define the city of Paris is "3 Paris." Multiple type/value pairs can be entered in any order, although HP recommends that multiple pairs be entered in ascending order of the <code>CA-TYPE</code> . When an emergency call is placed from a properly configured class 3 endpoint device to an appropriate PSAP, the country code, device type, and type/value pairs configured on the switch port are included in the transmission. The "type" specifiers are used by the PSAP to identify and organize the location data components in an understandable format for response personnel to interpret. A <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. <code>CA-TYPE</code> : This is the first entry in a type/value pair and is a number defining the type of data contained in the |

| | |
|------------------------------|---|
| | <p>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 195).</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> |
| elin-addr <emergency-number> | <p>This feature is intended for use in ECS applications to support class 3 LLDP-MED VoIP telephones connected to a switch in an MLTS infrastructure.</p> <p>An ELIN is a valid NANP format telephone number assigned to MLTS operators in North America by the appropriate authority. The ELIN is used to route emergency (E911) calls to a PSAP.</p> <p>(Range: 1-15 numeric characters)</p> |

Configuring coordinate-based locations

Latitude, longitude, and altitude data can be configured per switch port using an SNMP management application. For more information, see the documentation provided with the application. A further source of information on this topic is *RFC 3825-Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information*.

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

Table 23 Some location codes used in CA-TYPE fields*

| Location element | Code | Location element | Code |
|----------------------|------|--------------------------|------|
| national subdivision | 1 | street number | 19 |
| regional subdivision | 2 | additional location data | 22 |
| city or township | 3 | unit or apartment | 26 |
| city subdivision | 4 | floor | 27 |
| street | 6 | room number | 28 |
| street suffix | 18 | | |

* The code assignments in this table are examples from a work-in-progress (the internet draft titled "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information

draft-ietf-geopriv-dhcp-civil-06" dated May 30, 2005.) For the actual codes to use, contact the PSAP or other authority responsible for specifying the civic addressing data standard for your network.

Example:

Suppose a system operator wants to configure the following information as the civic address for a telephone connected to her company's network through port A2 of a switch at the following location:

| CA-type | CA-type | CA-VALUE |
|----------------------|---------|-------------|
| national subdivision | 1 | CA |
| city | 3 | Widgitville |
| street | 6 | Main |
| street number | 19 | 1433 |
| unit | 26 | Suite 4-N |
| floor | 27 | 4 |
| room number | 28 | N4-3 |

Example 112 shows the commands for configuring and displaying the above data.

Example 112 A civic address configuration

```
HP Switch(config)# lldp config 2 medportlocation civic-addr US 2 1 CA 3
Widgitville 6 Main 19 1433 26 Suite_4-N 27 4 28 N4-3
```

```
HP Switch(config)# show lldp config 2
LLDP Port Configuration Detail
Port : A2
AdminStatus [Tx_Rx] : Tx_Rx
NotificationEnabled [False] : False
Med Topology Trap Enabled [False] : False
Country Name       : US
What               : 2
Ca-Type            : 1
Ca-Length          : 2
Ca-Value           : CA
Ca-Type            : 3
Ca-Length          : 11
Ca-Value           : Widgitville
Ca-Type            : 6
Ca-Length          : 4
Ca-Value           : Main
Ca-Type            : 19
Ca-Length          : 4
Ca-Value           : 1433
Ca-Type            : 26
Ca-Length          : 9
Ca-Value           : Suite_4-N
Ca-Type            : 27
Ca-Length          : 1
Ca-Value           : 4
Ca-Type            : 28
Ca-Length          : 4
Ca-Value           : N4-3
```

Viewing switch information available for outbound advertisements

Syntax:

```
show lldp info local-device [port-list]
```

Without the [*port-list*] option, displays the global switch information and the per-port information currently available for populating outbound LLDP advertisements.

With the [*port-list*] option, displays only the following port-specific information that is currently available for outbound LLDP advertisements on the specified ports:

- PortType
- PortId
- PortDesc

NOTE: This command displays the information available on the switch. Use the `lldp config <port-list>` command to change the selection of information that is included in actual outbound advertisements. In the default LLDP configuration, all information displayed by this command is transmitted in outbound advertisements.

In the default configuration, the switch information currently available for outbound LLDP advertisements appears similar to the display in [Example 113 \(page 198\)](#).

Example 113 Displaying the global and per-port information available for outbound advertisements

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

LLDP Local Device Information

Chassis Type : mac-address
Chassis Id : 00 23 47 4b 68 DD
System Name : HP Switch1
System Description : HP J9091A Switch 3500yl, revision XX.15.06...
System Capabilities Supported:bridge
System Capabilities Enabled:bridge

Management Address 1
  Type:ipv4
  Address:
```

LLDP Port Information

| Port | PortType | PortId | PortDesc |
|------|----------|--------|----------|
| 1 | local | 1 | 1 |
| 2 | local | 2 | 2 |
| 3 | local | 3 | 3 |
| 4 | local | 4 | 4 |
| 5 | local | 5 | 5 |

- 1** The Management Address field displays only the LLDP-configurable IP addresses on the switch. (Only manually-configured IP addresses are LLDP-configurable.) If the switch has only an IP address from a DHCP or Bootp server, then the Management Address field is empty (because there are no LLDP-configurable IP addresses available).

Example 114 The default per-port information content for ports 1 and 2

```
HP Switch(config)# show lldp info local 1-2

LLDP Local Port Information Detail

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

-----

Port      : 2
PortType  : local
PortId    : 2
PortDesc  : 2
```

Displaying the current port speed and duplex configuration on a switch port

You can compare port speed and duplex information for a switch port and a connected LLDP-MED

endpoint for configuration mismatches by using an SNMP application. You can also use the switch CLI to display this information, if necessary. The `show interfaces brief <port-list>` and `show lldp info remote-device [port-list]` (Example 48) 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 199).

Viewing the current port speed and duplex configuration on a switch port

Syntax:

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

Includes port speed and duplex configuration in the Mode column of the resulting display.

Viewing advertisements currently in the neighbors MIB

Syntax:

```
show lldp info remote-device [ port-list ]
```

Without the `[port-list]` option, provides a global list of the individual devices it has detected by reading LLDP advertisements. Discovered devices are listed by the inbound port on which they were discovered.

Multiple devices listed for a single port indicates that such devices are connected to the switch through a hub.

Discovering the same device on multiple ports indicates that the remote device may be connected to the switch in one of the following ways:

- Through different VLANs using separate links. (This applies to switches that use the same MAC address for all configured VLANs.)
- Through different links in the same trunk.
- Through different links using the same VLAN. (In this case, spanning-tree should be invoked to prevent a network topology loop. Note that LLDP packets travel on links that spanning-tree blocks for other traffic types.)

With the `[port-list]` option, provides a listing of the LLDP data that the switch has detected in advertisements received on the specified ports.

For descriptions of the various types of information displayed by these commands, see Table 22 (page 175).

Example 115 A global listing of discovered devices

```
HP Switch(config)# show lldp info remote
```

```
LLDP Remote Devices Information
```

| LocalPort | ChassisId | PortId | PortDescr | SysName |
|-----------|-------------------|--------|-----------|-----------|
| 1 | 00 11 85 35 3b 80 | 6 | 6 | HP Switch |
| 2 | 00 11 85 cf 66 60 | 8 | 8 | HP Switch |

Example 116 An LLLDP-MED listing of an advertisement received from an LLLDP-MED (VoIP telephone) source

```
HP Switch(config)# show lldp info remote-device 1
```

```
LLDP Remote Device Information Detail
```

```
Local Port      : A2
ChassisType    : network-address
ChassisId      : 0f ff 7a 5c
PortType       : mac-address
PortId         : 08 00 0f 14 de f2
SysName        : HP Switch
System Descr   : HP Switch, revision xx.15.06.0000x
PortDescr      : LAN Port
```

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

```
Remote Management Address
```

```
MED Information Detail 1
```

```
EndpointClass      :Class3
Media Policy Vlan id :10
Media Policy Priority :7
Media Policy Dscp   :44
Media Policy Tagged :False
Poe Device Type     :PD
Power Requested     :47
Power Source        :Unknown
Power Priority       :High
```

- 1** Indicates the policy configured on the telephone. A configuration mismatch occurs if the supporting port is configured differently.
-

Displaying LLDP statistics

LLDP statistics are available on both a global and a per-port levels. Rebooting the switch resets the LLDP statistics counters to zero. Disabling the transmit and/or receive capability on a port "freezes" the related port counters at their current values.

Viewing LLDP statistics

For more information, see ["Displaying LLDP statistics" \(page 200\)](#).

Syntax:

```
show lldp stats [port-list]
```


The *global LLDP* statistics command displays an overview of neighbor detection activity on the switch, plus data on the number of frames sent, received, and discarded per-port.

The *per-port LLDP* statistics command enhances the list of per-port statistics provided by the global statistics command with some additional per-port LLDP statistics.

Global LLDP Counters:

| | |
|------------------------------------|---|
| Neighbor Entries List Last Updated | The elapsed time since a neighbor was last added or deleted. |
| New Neighbor Entries Count | The total of new LLDP neighbors detected since the last switch reboot. Disconnecting, and then reconnecting a neighbor increments this counter. |
| Neighbor Entries Deleted Count | The number of neighbor deletions from the MIB for AgeOut Count and forced drops for all ports. For example, if the admin status for port on a neighbor device changes from tx_rx or txonly to disabled or rxonly, the neighbor device sends a "shutdown" packet out the port and ceases transmitting LLDP frames out that port. The device receiving the shutdown packet deletes all information about the neighbor received on the applicable inbound port and increments the counter. |
| Neighbor Entries Dropped Count | The number of valid LLDP neighbors the switch detected, but could not add. This can occur, For example, when a new neighbor is detected when the switch is already supporting the maximum number of neighbors. See "Neighbor maximum" (page 202) . |
| Neighbor Entries AgeOut Count | The number of LLDP neighbors dropped on all ports because of Time-to-Live expiring. |

Per-Port LLDP Counters:

| | |
|--------------------|---|
| NumFramesRecvd | The total number of valid, inbound LLDP advertisements received from any neighbors on <i>port-list</i> . Where multiple neighbors are connected to a port through a hub, this value is the total number of LLDP advertisements received from all sources. |
| NumFramesSent | The total number of LLDP advertisements sent from <i>port-list</i> . |
| NumFramesDiscarded | The total number of inbound LLDP advertisements discarded by <i>port-list</i> . 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 202) . This can also be an indication of advertisement formatting problems in the neighbor device. |
| Frames Invalid | The total number of invalid LLDP advertisements received on the port. An invalid advertisement can be caused by header formatting problems in the neighbor device. |
| TLVs Unrecognized | The total number of LLDP TLVs received on a port with a type value in the reserved range. This can be caused by a basic management TLV from a later LLDP version than the one currently running on the switch. |
| TLVs Discarded | The total number of LLDP TLVs discarded for any reason. In this case, the advertisement carrying the TLV may be accepted, but the individual TLV is not usable. |
| Neighbor Ageouts | The number of LLDP neighbors dropped on the port because of Time-to-Live expiring. |

Examples:

Example 117 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 118 A per-port LLDP statistics display

```
HP Switch(config)# show lldp stats 1
```

LLDP Port Statistics Detail

```
PortName : 1
Frames Discarded : 0
Frames Invalid : 0
Frames Received : 7309
Frames Sent : 7231
TLVs Unrecognized : 0
TLVs Discarded : 0
Neighbor Ageouts : 0
```

LLDP Operating Notes

Neighbor maximum

The neighbors table in the switch supports as many neighbors as there are ports on the switch. The switch can support multiple neighbors connected through a hub on a given port, but if the switch neighbor maximum is reached, advertisements from additional neighbors on the same or other ports will not be stored in the neighbors table unless some existing neighbors time-out or are removed.

LLDP packet forwarding

An 802.1D-compliant switch does not forward LLDP packets, regardless of whether LLDP is globally enabled or disabled on the switch.

One IP address advertisement per port

LLDP advertises only one IP address per port, even if multiple IP addresses are configured by `lldp config port-list ipAddrEnable` (see [syntax \(page 183\)](#)) 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 \(CLI\)" \(page 180\)](#).

Mandatory TLVs

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

LLDP and CDP data management

This section describes points to note regarding LLDP and CDP (Cisco Discovery Protocol) data received by the switch from other devices. LLDP operation includes both transmitting LLDP packets to neighbor devices and reading LLDP packets received from neighbor devices. CDP operation is limited to reading incoming CDP packets from neighbor devices. (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 llarp` 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. |
| LLDP Enabled ¹ | Generates and transmits LLDP packets out all ports on the switch. | Store inbound LLDP data. | No forwarding of inbound LLDP packets. |
| LLDP Disabled | No packet generation. | No storage of LLDP data from neighbor devices. | No forwarding of inbound LLDP packets. |

¹ Both CDP data collection and LLDP transmit/receive are enabled in the default configuration. If a switch receives CDP packets and LLDP packets from the same neighbor device on the same port, it stores and displays the two types of information separately if the chassis and port ID information in the two types of advertisements is different. In this case, if you want to use only one type of data from a neighbor sending both types, disable the unwanted protocol on either the neighbor device or on the switch. However, if the chassis and port ID information in the two types of advertisements is the same, the LLDP information overwrites the CDP data for the same neighbor device on the same port.

CDP operation and commands

By default the switches have CDP enabled on each port. This is a read-only capability, meaning that the switch can receive and store information about adjacent CDP devices but does not generate CDP packets.

When a CDP-enabled switch receives a CDP packet from another CDP device, it enters that device's data in the CDP Neighbors table, along with the port number where the data was received—and does not forward the packet. The switch also periodically purges the table of any entries that have expired. (The hold time for any data entry in the switch's CDP Neighbors table is configured in the device transmitting the CDP packet and cannot be controlled in the switch receiving the packet.) A switch reviews the list of CDP neighbor entries every three seconds and purges any expired entries.

NOTE: For details on how to use an SNMP utility to retrieve information from the switch's CDP Neighbors table maintained in the switch's MIB, see the documentation provided with the particular SNMP utility.

Viewing the current CDP configuration of the switch

CDP is shown as enabled/disabled both globally on the switch and on a per-port basis.

Syntax:

```
show cdp
```

Lists the global and per-port CDP configuration of the switch.

Example 119 “Default CDP configuration” shows the default CDP configuration.

Example 119 Default CDP configuration

```
HP Switch(config)# show cdp

Global CDP information

  Enable CDP [Yes] : Yes (Receive Only)

Port CDP
-----
1    enabled
2    enabled
3    enabled
.    .
.    .
.    .
```

Viewing the current CDP neighbors table of the switch

Devices are listed by the port on which they were detected.

Syntax:

```
show cdp neighbors
```

Lists the neighboring CDP devices the switch detects, with a subset of the information collected from the device's CDP packet.

| | |
|------------------------------|--|
| [[e] port-numb [detail]] | Lists the CDP device connected to the specified port. (Allows only one port at a time.) Using <code>detail</code> provides a longer list of details on the CDP device the switch detects on the specified port. |
| [detail [[e] port-numb]] | Provides a list of the details for all of the CDP devices the switch detects. Using <code>port-num</code> produces a list of details for the selected port. |

Example 120 “CDP neighbors table listing” displays the CDP devices that the switch has detected by receiving their CDP packets.

Example 120 CDP neighbors table listing

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

```
CDP neighbors information
```

| Port | Device ID | Platform | Capability |
|------|----------------------------|-----------------------|------------|
| 1 | Accounting (0030c1-7fcc40) | J4812A HP Switch. . . | S |
| 2 | Researç1-1 (0060b0-889e43) | J4121A HP Switch. . . | S |
| 4 | Support (0060b0_761a45) | J4121A HP Switch. . . | S |
| 7 | Marketing (0030c5_33dc59) | J4313A HP Switch. . . | S |
| 12 | Mgmt NIC(099a05-09df9b | NIC Model X666 | H |
| 12 | Mgmt NIC(099a05-09df11 | NIC Model X666 | H |

Enabling and Disabling CDP Operation

Enabling CDP operation (the default) on the switch causes the switch to add entries to its CDP Neighbors table for any CDP packets it receives from other neighboring CDP devices.

Disabling CDP operation clears the switch's CDP Neighbors table and causes the switch to drop inbound CDP packets from other devices without entering the data in the CDP Neighbors table.

Syntax:

```
[no] cdp run
```

Enables or disables CDP read-only operation on the switch.

(Default: Enabled)

Example:

To disable CDP read-only on the switch:

```
HP Switch(config)# no cdp run
```

When CDP is disabled:

- `show cdp neighbors` displays an empty CDP Neighbors table
- `show cdp` displays
Global CDP information
Enable CDP [Yes]: No

Enabling or disabling CDP operation on individual ports

In the factory-default configuration, the switch has all ports enabled to receive CDP packets.

Disabling CDP on a port causes it to drop inbound CDP packets without recording their data in the CDP Neighbors table.

Syntax:

```
[no]cdp enable <[ e ]port-list>
```

Example:

To disable CDP on port A1:

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

Configuring CDPv2 for voice transmission

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

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

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

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

The switch CDP packet includes these TLVs:

- CDP Version: 2
- CDP TTL: 180 seconds
- Checksum
- Capabilities (type 0x04): 0x0008 (is a switch)
- Native VLAN: The PVID of the port
- VOIP VLAN Reply (type 0xe): voice VLAN ID (same as advertised by 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.

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

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

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

Example 121 The show cdp output when CDP Run is disabled

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

Example 122 The show cdp output when cdp run and sdp mode are enabled

```
HP Switch(config)# show cdp

Global CDP Information

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

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

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

Example 123 The show cdp output when cdp run and cdp mode rxonly are enabled

```
HP Switch(config)# show cdp

Global CDP Information

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

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

Example 124 The show running-config when admin-status is configured

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

Running configuration:

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

Filtering CDP information

In some environments it is desirable to be able to configure a switch to handle CDP packets by filtering out the MAC address learns from untagged VLAN traffic from IP phones. This means that normal protocol processing occurs for the packets, but the addresses associated with these packets is not learned or reported by the software address management components. This enhancement also filters out the MAC address learns from LLDP and 802.1x EAPOL packets on untagged VLANs. The feature is configured per-port.

Configuring the switch to filter untagged traffic

Enter this command to configure the switch not to learn CDP, LLDP, or EAPOL traffic for a set of interfaces.

Syntax:

```
[no] ignore-untagged-mac <port-list>
```

Prevents MAC addresses from being learned on the specified ports when the VLAN is untagged and the destination MAC address is one of the following:

- 01000C-CCCCC (CDP)
- 0180c2- 00000e (LLDP)
- 0180c2-000003 (EAPOL)

Example 125 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 126 Configuration showing interfaces to ignore packet MAC address learns

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

Running configuration:

```
; J9627 Configuration Editor; Created on release XX.15.XX  
; Ver #03:03.1f.ef:f0
```

```
hostname "HP Switch"  
interface 1  
    ignore-untagged-mac  
    exit  
interface 2  
    ignore-untagged-mac  
    exit  
.  
.  
.  
vlan 1  
    name "DEFAULT_VLAN"  
    untagged 1-24  
    ip address dhcp-bootp  
    exit  
.  
.  
.
```

Filtering PVID mismatch log messages

This enhancement filters out PVID mismatch log messages on a per-port basis. PVID mismatches are logged when there is a difference in the PVID advertised by a neighboring switch and the PVID of the switch port which receives the LLDP advertisement. Logging is an LLDP feature that allows detection of possible vlan leakage between adjacent switches. However, if these events are logged too frequently, they can overwhelm the log buffer and push relevant logging data out of log memory, making it difficult to troubleshoot another issue.

Logging is disabled and enabled with the support of CLI commands.

This enhancement also includes displaying the Mac-Address in the PVID mismatch log message when the port ID is Mac-Address instead of displaying garbage characters in the peer device port ID field.

Use the following command to disable the logging of the PVID mismatch log messages:

Syntax:

```
logging filter [filter-name] [sub filter id]  
<regularexpression> deny  
Regular-expression      The regular expression should match the message  
                        which is to be filtered.
```

Syntax:

```
logging filter [filter-name] enable
```

DHCPv4 server

Introduction to DHCPv4

The Dynamic Host Configuration Protocol (DHCP) is a network protocol that enables a server to automate assignment of IP addresses to hosts. A DHCP server can be configured to provide other network information like IP addresses of TFTP servers, DNS server, boot file name and vendor specific options. Commonly there are two types of address assignments, dynamic and manual. The lease of dynamic addresses is renewed periodically; manual leases are permanently assigned to hosts. With this feature, you can configure multiple pools of IP addresses for IP address assignment and tracking.

IP pools

A DHCP server is configured with IP pools. The server is then instructed to use IP addresses falling into the specified range of IP while offering leases. Multiple IP pools are configured to not have duplicate or overlapping IP subnets. You can also configure a DHCP server with multiple IP ranges within an IP subnet; this confines the allocatable IP addresses within the configured IP pool.

An IP pool will be claimed valid only if it is either:

- Dynamic pool – Has a network address, subnet mask and IP range(s)
- Static pool – Should have a static IP-to-MAC binding.

The DHCP server will discard the invalid and incomplete pools and will only operate on the valid IP pools. The DHCP server will require at least one valid pool to start.

DHCP options

On a DHCP server, an IP pool is configured with various options. These options signify additional information about the network. Options are supported with explicit commands such as `boot-file`. Option codes that correspond to explicit commands can not be configured with a generic option command; the generic option command requires an option code and TLV.

NOTE: RFC 2132 defines various network information that a client may request when trying to get the lease.

BootP support

The DHCP server also functions as BootP server. A manual binding configured in a static IP Pool may either service a BootP client request or a DHCP client request.

Authoritative server and support for DHCP inform packets

The server message `DHCPinform` may be received when the server is already configured for static IPv4 addresses so that the server can get configuration parameters dynamically.

NOTE: From RFC 2131 states that if a client has obtained a network address through some other means (e.g., manual configuration), it may use a `DHCPinform` request message to obtain other local configuration parameters. Servers receiving a `DHCPinform` message construct a `DHCPACK` message with any local configuration parameters appropriate for the client without: allocating a new address, checking for an existing binding, filling in `yiaddr` or including lease time parameters.

Authoritative pools

To process the `DHCPINFORM` packets received from a client within the given IP pool, a DHCP server has to be configured as `authoritative` for that IP pool. The server is the sole authority for this IP pool so when a client requests an IP address lease where the server is authoritative, and the server has no record of that IP address, the server will respond with `DHCPNAK` message which indicates that the client should no longer use that IP address. Any `DHCPINFORM` packet received for a non-authoritative pool will be ignored by the DHCP server.

The `authoritative` command has no effect when configured on a static pool or an incomplete pool without a network statement. In such cases, the server intentionally not send an error message. A CLI toggle is provided under the **pool** context that will allow the `authoritative` configuration.

NOTE: The `authoritative` command requires a network statement to be configured on a pool.

Authoritative dummy pools

A dummy pool, without the range statement, can be configured and made authoritative. A dummy pool allows static-bind entries which do not have matching dynamic pools with network statements to be configured. By creating a dummy pool on a DHCP server, the support for `DHCPinform` packets will not be actively serving the client on this pool. No active leases or resource consumption will be sent to the DHCP server when this option is used.

Dummy pools help the DHCP server learn the network topology.

Example

```
dhcp-server pool dummy192
network 192.168.10.0 255.255.255.255
option 1...
option 2...
:
option n...
authoritative
exit
```

Change in server behavior

Making the server authoritative for an IP pool changes how the server processes `DHCP REQUEST` packets.

Table XX exhibits the behavior on the receiving `DHCP REQUEST` and `DHCP inform` packets from DHCP clients residing on either authoritative and non-authoritative pools.

Table 24 Authoritative and non-authoritative pools

| | Authoritative Pool | | | Non-authoritative pool | | |
|------------------------------|--------------------|--------------------------------------|--------------------------------------|------------------------|--------------------------------------|--------------------------------------|
| When a DHCP client sending.. | For Own IP | For IP belonging to different client | Unknown IP falling outside the range | For Own IP | For IP belonging to different client | Unknown IP falling outside the range |
| DHCP INFORM | send ACK | send ACK | send ACK | DROP | DROP | DROP |
| DHCP REQUEST | send ACK | send NACK | send NACK | send ACK | DROP | DROP |

DHCPv4 configuration commands

Enable/disable the DHCPv4 server

Syntax

```
[no]dhcp-server [enable | disable]
```

To enable/disable the DHCPv4 server in a switch.

- Enable the DHCPv4 server on the device. The `no` form of this command will remove all DHCPv4 server configurations.
- Disable the DHCPv4 server on the device. The `no` form of this command will remove all DHCPv4 server configurations.

The default is disabled.

Configuring the DHCP address pool name

Use the following command in the global configuration mode to configure the DHCP address pool name and enter the DHCP pool context.

Maximum of 128 pools are supported.

Syntax

```
[no]dhcp-server pool < pool-name>
```

Configure the DHCPv4 server IP address pool with either a static IP or a network IP range.

| | |
|----------------------------------|--|
| <code>pool</code> | DHCPv4 server IP address pool. |
| <code>ASCII-STR</code> | Enter an ASCII string. |
| <code>authoritative</code> | Configure the DHCP server authoritative for a pool. |
| <code>bootfile-name</code> | Specify the boot file name which is used as a boot image. |
| <code>default-router</code> | List of IP addresses of the default routers. |
| <code>dns-server</code> | List of IP addresses of the DNS servers. |
| <code>domain-name</code> | Configure the DNS (Domain Name System) domain name for translation of hostnames to IP addresses. |
| <code>lease</code> | Lease period of an IP address. |
| <code>netbios-name-server</code> | List of IP addresses of the NetBIOS (WINS) name servers. |
| <code>netbios-node-type</code> | NetBIOS node type for a Microsoft DHCPv4 client. |
| <code>network</code> | Subnet IP and mask of the DHCPv4 server address pool. |

| | |
|-------------|--|
| option | Raw DHCPv4 server options. |
| range | Range of IP addresses for the DHCPv4 server address pool. |
| static-bind | Static binding information for the DHCPv4 server address pool. |
| tftp-server | Configure a TFTP server for the DHCPv4 server address pool. |

Validations

| Validation | Error/Warning/Prompt |
|--|--|
| Configuring pool when maximum Number of pools already configured. | Maximum number of pools (128) has already been reached |
| Configuring Pool with a name that exceeds the maximum length requirement. | String %s too long. Allowed length is 32 characters. |
| Trying to delete non existing pool | The specified address pool does not exist. |
| Only alphanumeric characters, numerals and underscore is allowed in the pool name. Violating this would throw the following error message. | Invalid name. Only alphanumeric characters and hyphen are allowed. |
| Trying to delete existing pool or adding new pool when DHCP server enabled. | DHCP server should be disabled before changing the configuration. |

Authoritative

Syntax

```
[no]authoritative
```

`authoritative` Configure the DHCP server authoritative for a pool.

The DHCP server is the sole authority for the network configured under this pool. When the DHCP server is configured as authoritative, the server will respond with DHCP ACK or NACK as appropriate for all the received DHCP REQUEST and DHCP INFORM packets belonging to the subnet.

Non-authoritative DHCP INFORM packets received from the clients on a non-authoritative pool will be ignored.

Specify a boot file for the DHCP client

Syntax

```
[no]bootfile-name<filename>
```

Specify the boot file name to be used as the boot image.

Configure a default router for a DHCP client

Syntax

```
[no]default-router <IP-ADDR-STR> [IP-ADDR2 IP-ADDR8]
```

Configure the DHCP pool context to the default router for a DHCP client. List all of the IP addresses of the default routers.

Two IP addresses must be separated by a comma.

Maximum of eight default routers can be configured.

Configure the DNS IP servers

Syntax

```
[no]dns-server <IP-ADDR> [IP-ADDR2 IP-ADDR8]
```

Configure the DHCP pool context to the DNS IP servers that are available to a DHCP client. List of IP addresses of the DNS servers.

Two IP addresses must be separated by comma.

Maximum of eight DNS servers can be configured.

Configure a domain name

Syntax

```
[no]domain-name <name>
```

Configure the DNS domain name for translation of hostnames to IP addresses.

Configure lease time

Syntax

```
[no]lease [DD:HH:MM | infinite]
```

DD:HH:MM Enter lease period.

Lease Lease period of an IP address.

Configure the lease time for an IP address in the DHCP pool. Lease time is infinite for static pools.

The default lease period is one day.

Configure the NetBIOS WINS servers

Syntax

```
[no]netbios-name-server <IP-ADDR-STR> [IP-ADDR2 IP-ADDR8]
```

Configure the DHCP pool for the NetBIOS WINS servers that are available to a Microsoft DHCP client. List all IP addresses of the NetBIOS(WINS) name servers. The Windows Internet Naming Service (WINS) is a name resolution service that Microsoft DHCP clients use to correlate host names to IP addresses within a general grouping of networks.

Two IP addresses must be separated by a comma.

Maximum of 8 NetBIOS (WINS) name servers can be configured.

Configure the NetBIOS node type

Syntax

```
[no]netbios-node-type [ broadcast | hybrid | mixed |  
peer-to-peer ]
```

broadcast Broadcast node.

hybrid Hybrid node.

mixed Mixed node.

peer-to-peer Peer to peer node.

Configure the DHCP pool mode to the NetBIOS node type for a Microsoft DHCP.

The NetBIOS node type for Microsoft DHCP clients can be one of four settings:

broadcast, peer-to-peer, mixed, or hybrid.

Configure subnet and mask

Syntax

```
[no]network <ip-addr/mask-length>  
ip-addr/mask-length      Interface IP address/mask.
```

Configure the DHCPv4 server pool subnet and mask for the DHCP server address pool.
Range is configured to enable pool.

Configure DHCP server options

Syntax

```
[no]option <CODE> {ascii <ascii-string>|hex <hex-string>|ip  
<IP-ADDR-STR> [IP-ADDR2 ... IP-ADDR8] }  
ascii                    Specify ASCII string as option code value.  
hex                      Specify hexadecimal string as option code value.  
ip                       Specify one or more IP addresses as option code value.  
ip-addr-str              Specify IP address.  
ascii-str                Enter an ASCII string.  
hex-str                  Specify Hexadecimal string.  
Configure the raw DHCP server options.
```

Configure the range of IP address

Syntax

```
[no]range <IP-ADDR> [<IP-ADDR>]  
range      Range of IP addresses for the DHCPv4 server address pool.  
ip-addr    Low IP address.  
           High IP address.
```

Configure the DHCP pool to the range of IP address for the DHCP address pool.

Configure the static binding information

Syntax

```
[no]static-bind ip<IP-ADDR/MASK-LENGTH> mac <MAC-ADDR>  
ip                    Specify client IP address.  
static-bind           Static binding information for the DHCPv4 server  
                      address pool.  
ip-addr / mask-length Interface IP address or mask.  
mac                   Specify client MAC address.  
mac-addr              Enter a MAC address.
```

Configure static binding information for the DHCPv4 server address pool. Manual bindings are IP addresses that have been manually mapped to the MAC addresses of hosts that are found in the DHCP database. Manual bindings are just special address pools. There is no limit on the number of manual bindings but you can only configure one manual binding per host pool.

Configure the TFTP server domain name

Syntax

```
[no]tftp-server [server-name <server-name> | server-ip <ip-address >]
tftp-server      Configure a TFTP server for the DHCPv4 server address pool.
server-name      TFTP server name for the DHCPv4 server address pool.
```

Configure the TFTP server domain name for the DHCP address pool.

Configure the TFTP server address

Syntax

```
[no]tftp-server server-ip <ip-address>
server-ip        TFTP server IP addresses for the DHCPv4 server address pool.
ip-addr          Specify TFTP server IP address.
Configure the TFTP server address for the DHCP address pool.
```

Change the number of ping packets

Syntax

```
[no]dhcp-server ping [packets <0-10>|timeout <0-10>]
ping             Specify DHCPv4 ping parameters.
packets          Specify number of ping packets.
<0-10>           Number of ping packets (0 disables ping).
Specify, in the global configuration context, the number of ping packets the DHCP server will send to the pool address before assigning the address. The default is two packets.
```

Change the amount of time

Syntax

```
[no]dhcp-server ping timeout <1-10>
timeout          Ping timeout.
<1-10>           Ping timeout in seconds.
Amount of time the DHCPv4 server must wait before timing out a ping packet. The default is one second.
```

Configure DHCP Server to save automatic bindings

Syntax

```
[no]dhcp-server database [file ASCII-STR]
[delay<15-86400>] [timeout <0-86400>]
delay            Seconds to delay writing to the lease database file.
file             URL Format: "tftp://<ip-address>/<filename>".
database         Specifies DHCPv4 database agent and the interval between database updates and database transfers.
timeout          Seconds to wait for the transfer before failing.
ascii-str        Database URL.
<15-86400>      Delay in seconds.
```


<0-86400> Timeout in seconds.

Specifies DHCPv4 database agent and the interval between database updates and database transfers.

Configure a DHCP server to send SNMP notifications

Syntax

```
[no]snmp-server enable traps dhcp-server  
dhcp-server      Traps for DHCP-Server.
```

Configure a DHCP server to send SNMP notifications to the SNMP entity. This command enables or disables event traps sent by the switch.

Enable conflict logging on a DHCP server

Syntax

```
[no]dhcp-server conflict-logging  
conflict-logging      Enable DHCPv4 server address conflict logging.  
Enable conflict logging on a DHCP server. Default is disabled.
```

Enable the DHCP server on a VLAN

Syntax

```
[no]dhcp-server  
dhcp-server      Enable DHCPv4 server on a VLAN.
```

Enable DHCPv4 server on a VLAN. DHCPv4 client or DHCPv4 relay cannot co-exist with DHCPv4 server on a VLAN.

Clear commands

Syntax

```
clear dhcp-server conflicts [ip-addr]  
dhcp-server      Clears theDHCPv4 server information.  
ip-addr          Specify the IP address whose conflict is to be cleared.
```

Reset DHCPv4 server conflicts database. If IP address is specified, reset only that conflict.

Reset all DHCP server and BOOTP counters

Syntax

```
clear dhcp-server statistics  
statistics      Reset DHCPv4 server and BOOTP counters.  
Reset all DHCP server and BOOTP counters
```

Delete an automatic address binding

Syntax

```
clear dhcp-server binding ip-addr  
binding      Reset DHCPv4 server automatic address bindings.  
ip-addr      Specify IP address of the binding is to be cleared.
```

Delete an automatic address binding from the DHCP server database.

Show commands

Display the DHCPv4 server address bindings

Syntax

```
show dhcp-server binding
dhcp-server      Show DHCPv4 server global configuration information for the
                  device.
binding         Show DHCPv4 server IP binding information for the device.
Display the DHCPv4 server address bindings on the device.
```

Display address conflicts

Syntax

```
show dhcp-server conflicts
conflicts       Show DHCPv4 server conflicts information for the device.
Display address conflicts found by a DHCPv4 server when addresses are offered
by a client.
```

Display DHCPv4 server database agent

Syntax

```
show dhcp-server database
Database       Show DHCPv4 server database information for the device.
Display DHCPv4 server database agent information.
```

Display DHCPv4 server statistics

Syntax

```
show dhcp-server statistics
statistics     Show DHCPv4 server statistics information for the device.
Display DHCPv4 server statistics.
```

Display the DHCPv4 server IP pool information

Syntax

```
show dhcp-server pool <pool-name>
Pool         Show DHCPv4 server pool information for the device.
Display the DHCPv4 server IP pool information.
```

Display DHCPv4 server global configuration information

Syntax

```
show dhcp-server
dhcp-server   Show DHCPv4 server global configuration information for the
                  device.
Display DHCPv4 server global configuration information.
```

Event log

Event Log Messages

Table 25 Event Log Messages

| Events | Debug messages |
|--|--|
| DHCP server is enabled globally. | DHCP server is enabled globally. |
| DHCP server is enabled globally. Warnings - One or more incomplete pool configurations are found during the server startup. A dynamic pool is considered invalid, if network IP or subnet mask is not configured. A static pool is considered incomplete, if network IP, subnet mask or MAC address is not configured. | DHCP server is enabled globally. Warning -One or more incomplete pool configurations are found during the server startup. |
| DHCP server failed to start. The reason for failure is printed as the argument. | DHCP server failed to start: %s "with a manual binding. |
| DHCP server is disabled globally. | DHCP server is disabled globally. |
| The DHCP server configurations are deleted. | The DHCP server configurations are deleted |
| | |
| Decline from client when server assigns an illegal Ipv6 address. | %s: Decline offer from %x (server) of %x because the address is illegal. |
| DHCP server is enabled on a specific VLAN. | DHCP server is enabled on VLAN %d |
| DHCP server is disabled on a specific VLAN. | DHCP server is disabled on VLAN %d |
| Ping check is enabled and configured with specified retry count and timeout values | Ping-check configured with retry count = %d, timeout = %d |
| Ping check is disabled | Ping-check is disabled |
| Conflict-logging is enabled | Conflict-logging is enabled |
| Conflict-logging is disabled. | Conflict-logging is disabled. |
| A specific IP address is removed from the conflict logging database. | IP address %s is removed from the conflict-logging database. |
| All IP addresses are removed from the conflict-logging database. | "All IP addresses are removed from the conflict-logging database |
| Dynamic binding for a specific IP address is freed. | Dynamic binding for IP address %s is freed |
| All the dynamic IP bindings are freed. | All the dynamic IP bindings are freed |
| Remote binding database is configured for a specific URL. | Remote binding database is configured at %s |
| Remote biding database is disabled. | Remote binding database is disabled |
| Binding database is read from the specified URL at the specified time | Binding database read from %s at %s |
| Failed to read the remote binding from the specified URL. | Failed to read the remote binding database at %s |
| Binding database is written to the specified URL at the specified time. | Binding database written to %s at %s |

Table 25 Event Log Messages *(continued)*

| Events | Debug messages |
|--|--|
| Failed to write the binding database to the specified URL. The reason for failure is printed as argument. | Failed to write the binding database to %s. Error: %s |
| Invalid bindings are found in the database at the specified URL. | Invalid binding database at %s |
| The specified VLAN does not have a matching IP pool configured.This occurs when the DHCP-server is enabled on the specified VLAN, but no IP pool is configured with a network IP matching the VLAN network IP. | VLAN %d does not have a matching IP pool |
| Binding database is replicated to standby management module. | Binding database is replicated to standby management module |
| DHCP server is listening for DHCP packetsThis message is displayed when DHCP server is enabled globally and DHCP server is enabled on at-least one VLAN. | DHCP server is listening for DHCP packets |
| DHCP server is disabled on all the VLANs. Server is no longer listening for DHCP packets. | DHCP server is disabled on all the VLANs. Server is no longer listening for DHCP packets |
| The specified IP is not offered to the DHCP client, as it is already in use. | IP address %s is not offered, as it is already in use |
| No IP addresses available on the specified pool. | No IP addresses to offer from pool %s |
| High threshold reached for the specified pool. Count of Active bindings and Free bindings are printed as arguments. | High threshold reached for pool %s. Active bindings: %d, Free bindings: %d |
| Low threshold reached for the specified pool. Count of Active bindings and Free bindings are printed as arguments. | Low threshold reached for pool %s. Active bindings: %d, Free bindings: %d |
| No active VLAN with an IP address is available to read binding database from the configured URL. | No active Vlan with an IP address available to read binding database |

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

LACP-MAD commands

Configuration command

The following command defines whether LACP is enabled on a port, and whether it is in active or passive mode when enabled. When LACP is enabled and active, the port sends LACP packets and listens to them. When LACP is enabled and passive, the port sends LACP packets only if it is spoken to. When LACP is disabled, the port ignores LACP packets. If the command is issued without a mode parameter, 'active' is assumed. During dynamic link aggregation using LACP, ports with the same key are aggregated as a single trunk. MAD passthrough applies only to trunks and not to physical ports.

```
HP-Switch# [no] interface <port-list> lacp [mad-passthrough  
<enable|disable>|active|passive|key <key>]
```

show commands

LACP-MAD supports the following show commands:

- show LACP-MAD passthrough configuration on LACP trunks
HP-Switch# show lacp [counters [<port-list>] | local [<port-list>]
|peer [<port-list>] | distributed | mad-passthrough [counters
[<port-list>]]]
- show LACP-MAD passthrough counters on ports
HP-Switch# show lacp mad-passthrough counters [<port-list>]

clear command

Clear all LACP statistics including MAD passthrough counters. Resets LACP packets sent and received on all ports.

```
HP-Switch# clear lacp statistics
```

LACP-MAD overview

Link Aggregation Control Protocol-Multi-Active Detection (LACP-MAD) is a detection mechanism deployed by switches to recover from a breakup of the Intelligent Resilient Framework (IRF) stack due to link or other failure.

LACP-MAD is implemented by sending extended LACP data units (LACPDUs) with a type length value (TLV) that conveys the active ID of an IRF virtual device. The active ID is identical to the member ID of the master and is thus unique to the IRF virtual device. When LACP MAD detection is enabled, the members exchange their active IDs by sending extended LACPDUs.

- When the IRF virtual device operates normally, the active IDs in the extended LACPDUs sent by all members are the same, indicating that there is no multi-active collision.
- When there is a breakup in the IRF stack, the active IDs in the extended LACPDUs sent by the members in different IRF virtual devices are different, indicating that there are multi-active collisions.

LACP-MAD passthrough helps IRF-capable devices detect multi-access and take corrective action. These devices do not initiate transmission of LACP-MAD frames or participate in any MAD decision making process. These devices simply forward LACP-MAD TLVs received on one interface to the other interfaces on the trunk. LACP-MAD passthrough can be enabled for 24 LACP trunks. By default, LACP-MAD passthrough is disabled.

8 Scalability IP Address VLAN and Routing Maximum Values

The following table lists the switch scalability values for the areas of VLANs, ACLs, hardware, ARP, and routing.

| Subject | Maximum |
|--|--|
| IPv4 ACLs | |
| total named (extended or standard) | Up to 2048 (minus any IPv4 numeric standard or extended ACL assignments and any RADIUS-assigned ACLs) ¹ |
| total numbered standard | Up to 99 ¹ |
| total numbered extended | Up to 100 ¹ |
| total ACEs in all IPv4 ACLs | Up to 3072 ¹ |
| Layer-3 | |
| VLANs with at least one IP Address | 512 |
| IP addresses per system | 2048 IPv4 2048 IPv6 ² |
| IP addresses per VLAN | 32 ³ |
| Static routes (IPv4 and IPv6 combined) | 256 |
| IPv4 host hardware table | 72 K (8K internal, 64K external) |
| IPv4 BMP hardware table | 2 K |
| ARP | |
| ARP entries | 25,000 |
| Packets held for ARP resolution | 25 |
| Dynamic Routing | |
| Total routes supported | IPv4 only: 10,000 (including ARP) IPv4 and IPv6: 10 K (IPv4) and 3 K (IPv6) ⁴ IPv6 only: 5 K ⁵ |
| IPv4 Routing Protocol | |
| RIP interfaces | 128 |

| Subject | Maximum |
|------------------------------|---|
| IPv6 Routing Protocol | |
| DHCPv6 Helper Addresses | 32 unique addresses; multiple instances of same address counts as 1 towards maximum |

- ¹ Actual availability depends on combined resource usage on the switch. See [“Monitoring resources” \(page 40\)](#).
- 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.

9 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 “[Downloading switch software](#)” (page 224).

Downloading switch software

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

NOTE: This manual uses the terms *switch software* and *software image* to refer to the downloadable software files the switch uses to operate its networking features. Other terms sometimes include *Operating System*, or *OS*.

General software download rules

- Switch software that you download via the menu interface always goes to primary flash.
- After a software download, you must reboot the switch to implement the new software. Until a reboot occurs, the switch continues to run on the software it was using before the download.

NOTE: Downloading new switch software does not change the current switch configuration. The switch configuration is contained in separate files that can also be transferred. See “[Transferring switch configurations](#)” (page 241).

In most cases, if a power failure or other cause interrupts a flash image download, the switch reboots with the image previously stored in primary flash. In the unlikely event that the primary image is corrupted (which may occur if a download is interrupted by a power failure), the switch goes into boot ROM mode. In this case, use the boot ROM console to download a new image to primary flash.

Using TFTP to download software from a server

This procedure assumes that:

- A software version for the switch has been stored on a TFTP server accessible to the switch. (The software file is typically available from the HP Switch Networking website at <http://www.hp.com/networking>.)
- The switch is properly connected to your network and has already been configured with a compatible IP address and subnet mask.
- The TFTP server is accessible to the switch via IP.

Before you use the procedure, do the following:

- Obtain the IP address of the TFTP server in which the software file has been stored.
- If VLANs are configured on the switch, determine the name of the VLAN in which the TFTP server is operating.
- Determine the name of the software file stored in the TFTP server for the switch (For example, E0820.swi).

NOTE: If your TFTP server is a UNIX workstation, ensure that the case (upper or lower) that you specify for the filename is the same case as the characters in the software filenames on the server.

Downloading from a server to primary flash using TFTP (Menu)

Note that the menu interface accesses only the primary flash.

1. In the console Main Menu, select **Download OS** to display the screen in [Figure 38 \(page 225\)](#). (The term "OS" or "operating system" refers to the switch software):

Figure 38 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 39 \(page 225\)](#) appears:

Figure 39 Example: of the download OS (software) screen during a download

```
===== CONSOLE - MANAGER MODE =====
                          Download OS

Current Firmware revision : E.08.00
Method [TFTP] : TFTP
TFTP Server : 10.28.227.105

Remote File Name : K.11.00.swi

                          Received 370,000 bytes of OS download.
+-----+
|*****|
+-----+
```

A "progress" bar indicates the progress of the download. When the entire software file has been received, all activity on the switch halts and you will see **Validating and writing system software to FLASH...**

7. After the primary flash memory is updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch**).

You will see this prompt:

```
Continue reboot of system? : No
```

Press the space bar once to change **No** to **Yes**, then press **[Enter]** to begin the reboot.

NOTE: When you use the menu interface to download a switch software, the new image is always stored in primary flash. Also, using the `Reboot Switch` command in the Main Menu always reboots the switch from primary flash. Rebooting the switch from the CLI provides more options. See "Rebooting the Switch" in the *Basic Operation Guide* for your switch.


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

For troubleshooting information on download failures, see "Troubleshooting TFTP download failures" (page 226).

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 40 (page 226)).

Figure 40 Example: of message for download failure



Some of the causes of download failures include:

- Incorrect or unreachable address specified for the **TFTP Server** parameter. This may include network problems.
- Incorrect VLAN.
- Incorrect name specified for the **Remote File Name** parameter, or the specified file cannot be found on the TFTP server. This can also occur if the TFTP server is a UNIX machine and the case (upper or lower) for the filename on the server does not match the case for the filename entered for the **Remote File Name** parameter in the **Download OS** (Operating System, or software) screen.
- One or more of the switch's IP configuration parameters are incorrect.

- For a UNIX TFTP server, the file permissions for the software file do not allow the file to be copied.
- Another console session (through either a direct connection to a terminal device or through Telnet) was already running when you started the session in which the download was attempted.

To find more information on the cause of a download failure:

- Examine the messages in the switch's Event Log by executing the `show log tftp` command from the CLI.
- For descriptions of individual Event Log messages, see the latest version of the *Event Log Message Reference Guide* for your switch, available on the HP Switch website. (See "Getting Documentation From the Web".)

NOTE: If an error occurs in which normal switch operation cannot be restored, the switch automatically reboots itself, and an appropriate message is displayed after the reboot.

Downloading from a server to flash using TFTP (CLI)

Syntax:

```
copy tftp flash <ip-address> <remote-file> [ <primary | secondary> ]
```

Automatically downloads a switch software file to primary or secondary flash. If you do not specify the flash destination, the TFTP download defaults to primary flash.

Example:

To download a switch software file named `k0800.swi` from a TFTP server with the IP address of `10.28.227.103` to primary flash:

1. Execute `copy` as shown below:

Example 127 The command to download an OS (switch software)

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

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

2 Dynamic counter continually displays the number of bytes transferred.

When the switch finishes downloading the software file from the server, it displays this progress message:

Validating and Writing System Software to FLASH ...

2. When the download finishes, you must reboot the switch to implement the newly downloaded software image. To do so, use one of the following commands:

Syntax:

```
boot system flash <primary | secondary>
```

Boots from the selected flash.

Syntax:

```
reload
```

Boots from the flash image and startup-config file. A switch covered in this guide (with multiple configuration files), also uses the current startup-config file.

For more information on these commands, see "Rebooting the Switch" in the *Basic Operation Guide* for your switch.

3. To confirm that the software downloaded correctly, execute `show system` and check the **Firmware revision** line.

For information on primary and secondary flash memory and the boot commands, see "Using Primary and Secondary Flash Image Options" in the *Basic Operation Guide* for your switch.

NOTE: If you use `auto-tftp` to download a new image in a redundant management system, the active management module downloads the new image to both the active and standby modules. Rebooting after the `auto-tftp` process completes reboots the entire system.

Enabling TFTP (CLI)

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

Syntax:

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

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

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

Usage notes:

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

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

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

The `no tftp [client|server]` command does not disable auto-TFTP operation. To disable an auto-TFTP command configured on the switch, use the `no auto-tftp` command to remove the command entry from the switch's configuration.

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

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

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

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

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

Syntax:

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

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

NOTE: To enable auto-TFTP to copy a software image to primary flash memory, the version number of the downloaded software file (For example, XX_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 228).)

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

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

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

Using SCP and SFTP

For some situations you may want to use a secure method to issue commands or copy files to the switch. By opening a secure, encrypted SSH session and enabling `ip ssh file transfer`, you can then use a third-party software application to take advantage of SCP and SFTP. SCP and SFTP provide a secure alternative to TFTP for transferring information that may be sensitive (like switch configuration files) to and from the switch. Essentially, you are creating a secure SSH tunnel as a way to transfer files with SFTP and SCP channels.

Once you have configured your switch to enable secure file transfers with SCP and SFTP, files can be copied to or from the switch in a secure (encrypted) environment and TFTP is no longer necessary.

To use these commands, you must install on the administrator workstation a third-party application software client that supports the SFTP and/or SCP functions. Some examples of software that supports SFTP and SCP are PuTTY, Open SSH, WinSCP, and SSH Secure Shell. Most of these are freeware and may be downloaded without cost or licensing from the internet. There are differences in the way these clients work, so be sure you also download the documentation.

As described earlier in this chapter you can use a TFTP client on the administrator workstation to update software images. This is a plain-text mechanism that connects to a standalone TFTP server or another HP switch acting as a TFTP server to obtain the software image files. Using SCP and SFTP allows you to maintain your switches with greater security. You can also roll out new software images with automated scripts that make it easier to upgrade multiple switches simultaneously and securely.

SFTP is unrelated to FTP, although there are some functional similarities. Once you set up an SFTP session through an SSH tunnel, some of the commands are the same as FTP commands. Certain commands are not allowed by the SFTP server on the switch, such as those that create files or folders. If you try to issue commands such as `create` or `remove` using SFTP, the switch server returns an error message.

You can use SFTP just as you would TFTP to transfer files to and from the switch, but with SFTP, your file transfers are encrypted and require authentication, so they are more secure than they would be using TFTP. SFTP works only with SSH version 2 (SSH v2).

NOTE: SFTP over SSH version 1 (SSH v1) is not supported. A request from either the client or the switch (or both) using SSH v1 generates an error message. The actual text of the error message differs, depending on the client software in use. Some examples are:

```
Protocol major versions differ: 2 vs. 1
Connection closed
```

```
Protocol major versions differ: 1 vs. 2
Connection closed
```

```
Received disconnect from <ip-addr> : /usr/local/libexec/
sftp-server: command not supported
Connection closed
```

SCP is an implementation of the BSD `rcp` (Berkeley UNIX remote copy) command tunneled through an SSH connection.

SCP is used to copy files to and from the switch when security is required. SCP works with both SSH v1 and SSH v2. Be aware that the most third-party software application clients that support SCP use SSHv1.

The general process for using SCP and SFTP involves three steps:

1. Open an SSH tunnel between your computer and the switch if you have not already done so. (This step assumes that you have already set up SSH on the switch.)
2. Execute `ip ssh filetransfer` to enable secure file transfer.
3. Use a third-party client application for SCP and SFTP commands.

Enabling SCP and SFTP

For more information about secure copy and SFTP, see [“Using SCP and SFTP” \(page 229\)](#).

1. Open an SSH session as you normally would to establish a secure encrypted tunnel between your computer and the switch.

For more detailed directions on how to open an SSH session, see chapter "Configuring secure shell (SSH)" in the *Access Security Guide* for your switch. Please note that this is a one-time procedure for new switches or connections. If you have already done it once you should not need to do it a second time.

2. To enable secure file transfer on the switch (once you have an SSH session established between the switch and your computer), open a terminal window and enter the following command:

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

For information on disabling TFTP and auto-TFTP, see [“Disabling TFTP and auto-TFTP for enhanced security” \(page 230\)](#).

Disabling TFTP and auto-TFTP for enhanced security

Using the `ip ssh filetransfer` command to enable SFTP automatically disables TFTP and auto-TFTP (if either or both are enabled), as shown in [Example 128 \(page 231\)](#).

Example 128 Switch configuration with SFTP enabled

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

Running configuration:

```
; J9091A Configuration Editor; Created on release #xx.15.xx

hostname "HP Switch"
module 1 type J8702A
module 2 type J702A
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24,B1-B24
  ip address 10.28.234.176 255.255.240.0
  exit
ip ssh filetransfer 2
no tftp-enable
password manager
password operator
```

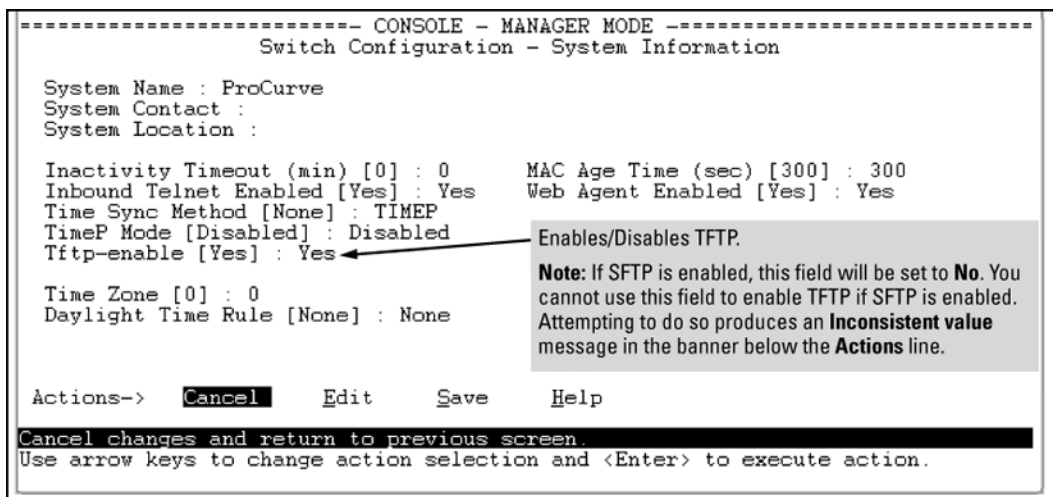
- 1 Enabling SFTP automatically disables TFTP and auto-tftp and displays this message. 2 Viewing the configuration shows that SFTP is enabled and TFTP is disabled.

If you enable SFTP and then later disable it, TFTP and auto-TFTP remain disabled unless they are explicitly re-enabled.

Operating rules are:

- The TFTP feature is enabled by default, and can be enabled or disabled through the CLI, the Menu interface (see [Figure 41 \(page 231\)](#)), or an SNMP application. Auto-TFTP is disabled by default and must be configured through the CLI.

Figure 41 Using the Menu interface to disable TFTP



- While SFTP is enabled, TFTP and auto-TFTP cannot be enabled from the CLI. Attempting to enable either non-secure TFTP option while SFTP is enabled produces one of the following messages in the CLI:

SFTP must be disabled before enabling tftp.

SFTP must be disabled before enabling auto-tftp.

Similarly, while SFTP is enabled, TFTP cannot be enabled using an SNMP management application. Attempting to do so generates an "inconsistent value" message. (An SNMP management application cannot be used to enable or disable auto-TFTP.)

- To enable SFTP by using an SNMP management application, you must first disable TFTP and, if configured, auto-TFTP on the switch. You can use either an SNMP application or the CLI to disable TFTP, but you must use the CLI to disable auto-TFTP. The following CLI commands disable TFTP and auto-TFTP on the switch.

Enabling SSH V2 (required for SFTP)

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

NOTE: As a matter of policy, administrators should *not* enable the SSH V1-only or the SSH V1-or-V2 advertisement modes. SSHv1 is supported on only some legacy switches (such as the HP Switch Series 2500 switches).

Confirming that SSH is enabled

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

Once you have confirmed that you have enabled an SSH session (with the `show ip ssh` command), enter `ip ssh filetransfer` so that SCP and/or SFTP can run. You can then open your third-party software client application to begin using the SCP or SFTP commands to safely transfer files or issue commands to the switch.

NOTE: Any attempts to use SCP or SFTP without using `ip ssh filetransfer` cause the SCP or SFTP session to fail. Depending on the client software in use, you will receive an error message on the originating console, for Example:

```
IP file transfer not enabled on the switch
```

Disabling secure file transfer

```
HP Switch(config)# no ip ssh filetransfer
```

Authentication

Switch memory allows up to ten public keys. This means the authentication and encryption keys you use for your third-party client SCP/SFTP software can differ from the keys you use for the SSH session, even though both SCP and SFTP use a secure SSH tunnel.

NOTE: SSH authentication is mutually exclusive with RADIUS servers.

Some clients, such as PSCP (PuTTY SCP), automatically compare switch host keys for you. Other clients require you to manually copy and paste keys to the `$HOME/.ssh/known_hosts` file. Whatever SCP/SFTP software tool you use, after installing the client software you must verify that the switch host keys are available to the client.

Because the third-party software utilities you may use for SCP/SFTP vary, you should refer to the documentation provided with the utility you select before performing this process.

SCP/SFTP operating notes

- Any attempts to use SCP or SFTP without using `ip ssh filetransfer` will cause the SCP or SFTP session to fail. Depending on the client software in use, you will receive an error message on the originating console, for Example:

IP file transfer not enabled on the switch

- There is a delay when SFTP is copying an image onto the switch, and although the command prompt returns in a couple of seconds, the switch may take approximately a minute and half writing the image to flash. You can keep entering the show flash command to see when the copy is complete and the flash is updated. You can also check the log for an entry similar to the following:

```
I 01/09/13 16:17:07 00150 update: Primary Image updated.
```

```
I 01/09/13 16:13:22 00636 ssh: sftp session from 15.22.22.03
```

- When an SFTP client connects, the switch provides a file system displaying all of its available files and folders. No file or directory creation is permitted by the user. Files may be only uploaded or downloaded, according to the permissions mask. All of the necessary files the switch needs are already in place on the switch. You do not need to (nor can you) create new files.
- The switch supports one SFTP session or one SCP session at a time.
- All files have read-write permission. Several SFTP commands, such as create or remove, are not allowed and return an error message. The switch displays the following files:

```
/
+---cfg
|   running-config
|   startup-config
+---log
|   crash-data
|   crash-data-a
|   crash-data-b
|   crash-data-c
|   crash-data-d
|   crash-data-e           "   "
|   crash-data-f         " "
|   crash-data-g
|   crash-data-h           "   "
|   crash-data-I         " "
|   crash-data-J         " "
|   crash-data-K         " "
|   crash-data-L         "  "
|   crash-log
|   crash-log-a
|   crash-log-b
|   crash-log-c
|   crash-log-d
|   crash-log-e         " "
|   crash-log-f         " "
|   crash-log-g
|   crash-log-h         " "
|   crash-log-I         " "
|   crash-log-J         " "
|   crash-log-K         " "
|   crash-log-L         " "
|   event log
+---os
|   primary
|   secondary
\---ssh
    +---mgr_keys
    |   authorized_keys
    \---oper_keys
    |   authorized_keys
\---core
|   port_1-24.cor      core-dump for ports 1-24 (stackable switches only)
|   port_25-48.cor    core-dump for ports 25-48 (stackable switches only)
```

Once you have configured your switch for secure file transfers with SCP and SFTP, files can be copied to or from the switch in a secure (encrypted) environment and TFTP is no longer necessary.

Troubleshooting SSH, SFTP, and SCP operations

You can verify secure file transfer operations by checking the switch's event log, or by viewing the error messages sent by the switch that most SCP and SFTP clients print out on their console.

NOTE: Messages that are sent by the switch to the client depend on the client software in use to display them on the user console.

Broken SSH connection

If an ssh connection is broken at the wrong moment (for instance, the link goes away or spanning tree brings down the link), a fatal exception occurs on the switch. If this happens, the switch gracefully exits the session and produces an Event Log message indicating the cause of failure. The following three examples show the error messages that may appear in the log, depending on the type of session that is running (SSH, SCP, or SFTP):

```
ssh: read error Bad file number, session aborted I 01/01/90
00:06:11 00636 ssh: sftp session from ::ffff:10.0.12.35 W
01/01/90 00:06:26 00641 ssh:
```

```
sftp read error Bad file number, session aborted I 01/01/90
00:09:54 00637 ssh: scp session from ::ffff:10.0.12.35 W 01/
01/90
```

```
ssh: scp read error Bad file number, session aborted
```

NOTE: The Bad file number is from the system error value and may differ depending on the cause of the failure. In the third Example:, the device file to read was closed as the device read was about to occur.

Attempt to start a session during a flash write

If you attempt to start an SCP (or SFTP) session while a flash write is in progress, the switch does not allow the SCP or SFTP session to start. Depending on the client software in use, the following error message may appear on the client console:

```
Received disconnect from 10.0.12.31: 2: Flash access in
progress

lost connection
```

Failure to exit from a previous session

This next Example: shows the error message that may appear on the client console if a new SCP (or SFTP) session is started from a client before the previous client session has been closed (the switch requires approximately ten seconds to timeout the previous session):

```
Received disconnect from 10.0.12.31: 2: Wait for previous
session to complete

lost connection
```

Attempt to start a second session

The switch supports only one SFTP session or one SCP session at a time. If a second session is initiated (For example, an SFTP session is running and then an SCP session is attempted), the following error message may appear on the client console:

```
Received disconnect from 10.0.12.31: 2: Other SCP/SFTP
session running
```

Using Xmodem to download switch software from a PC or UNIX workstation

This procedure assumes that:

- The switch is connected via the Console RS-232 port to a PC operating as a terminal. (For information on connecting a PC as a terminal and running the switch console interface, see the *Installation and Getting Started Guide* you received with the switch.)
- The switch software is stored on a disk drive in the PC.
- The terminal emulator you are using includes the Xmodem binary transfer feature. (For example, in the HyperTerminal application included with Windows NT, you would use the **Send File** option in the **Transfer** drop-down menu.)

Downloading to primary flash using Xmodem (Menu)

NOTE: The menu interface accesses only the primary flash.

1. From the console Main Menu, select
7. Download OS
2. Press **[E]** (for **Edit**).
3. Use the Space bar to select **XMODEM** in the **Method** field.
4. Press **[Enter]**, then **[X]** (for **eXecute**) to begin the software download.
The following message appears:
Press enter and then initiate Xmodem transfer from the attached computer.....
5. Press **[Enter]** and then execute the terminal emulator commands to begin Xmodem binary transfer.
For example, using HyperTerminal:
 - a. Click on **Transfer**, then **Send File**.
 - b. Enter the file path and name in the Filename field.
 - c. In the Protocol field, select **Xmodem**.
 - d. Click on the **[Send]** button.

The download then commences. It can take several minutes, depending on the baud rate set in the switch and in your terminal emulator.
6. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch**). You then see the following prompt:
Continue reboot of system? : No
Press the space bar once to change **No** to **Yes**, then press **[Enter]** to begin the reboot.
7. To confirm that the software downloaded correctly:
 - a. From the Main Menu, select
 1. **Status and Counters**
 1. **General System Information**
 - b. Check the **Firmware revision** line.

Downloading to primary or secondary flash using Xmodem and a terminal emulator (CLI)

Syntax:

```
copy xmodem flash [ <primary | secondary> ]
```

Downloads a software file to primary or secondary flash. If you do not specify the flash destination, the Xmodem download defaults to primary flash.

Example:

To download a switch software file named `E0822.swi` from a PC (running a terminal emulator program such as HyperTerminal) to primary flash:

1. Execute the following command in the CLI:

```
HP Switch# copy xmodem flash
Press 'Enter' and start XMODEM on your host...
```

2. Execute the terminal emulator commands to begin the Xmodem transfer. For example, using HyperTerminal:

- a. Click on **Transfer**, then **Send File**.
- b. Type the file path and name in the Filename field.
- c. In the Protocol field, select **Xmodem**.
- d. Click on the **[Send]** button.

The download can take several minutes, depending on the baud rate used in the transfer.

3. When the download finishes, you must reboot the switch to implement the newly downloaded software. To do so, use one of the following commands:

Syntax:

```
boot system flash <primary | secondary>
```

Reboots from the selected flash

Syntax:

```
reload
```

Reboots from the flash image currently in use

For more information on these commands, see "Rebooting the Switches" in the *Basic Operation Guide* for your switch.

4. To confirm that the software downloaded correctly:

```
HP Switch> show system
```

Check the **Firmware revision** line. It should show the software version that you downloaded in the preceding steps.

If you need information on primary/secondary flash memory and the boot commands, see "Using Primary and Secondary Flash Image Options" in the *Basic Operation Guide* for your switch.

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. 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.

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 "Using Primary and Secondary Flash Image Options" in the *Basic Operation Guide* for your switch.

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.

To copy a switch software file named `k0800.swi` from a USB device to primary flash:

1. Execute `copy` as shown below:

Example 129 The command to copy switch software from USB

```
HP Switch# copy usb flash XX.0800.swi
The Primary OS Image will be deleted, continue [y/n]? y 1
```

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

2. 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....

3. 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 "Rebooting the Switch" in the *Basic Operation Guide* for your switch.

4. 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.

Switch-to-switch download to primary flash (Menu)

Using the menu interface, you can download a switch software file from either the primary or secondary flash of one switch to the primary flash of another switch of the same series.

1. From the switch console Main Menu in the switch to receive the download, select **7. Download OS** screen.
2. Ensure that the **Method** parameter is set to **TFTP** (the default).
3. In the **TFTP Server** field, enter the IP address of the remote switch containing the software file you want to download.
4. For the **Remote File Name**, enter one of the following:
 - To download the software in the primary flash of the source switch, enter `flash` in lowercase characters.
 - To download the software in the secondary flash of the source switch, enter `/os/secondary`.
5. Press **[Enter]**, and then **[X]** (for **eXecute**) to begin the software download.

A "progress" bar indicates the progress of the download. When the entire switch software download has been received, all activity on the switch halts and the following messages appear:

Validating and writing system software to FLASH...

6. After the primary flash memory has been updated with the new software, you must reboot the switch to implement the newly downloaded software. Return to the Main Menu and press **[6]** (for **Reboot Switch**). You then see this prompt:

Continue reboot of system? : No

Press the space bar once to change `No` to `Yes`, then press **[Enter]** to begin the reboot.

7. To confirm that the software downloaded correctly:
 - a. From the Main Menu, select
Status and Counters
General System Information
 - b. Check the **Firmware revision** line.

Downloading the OS from another switch (CLI)

Where two switches in your network belong to the same series, you can download a software image between them by initiating a `copy tftp` command from the destination switch. The options for this CLI feature include:

- Copy from primary flash in the source to either primary or secondary in the destination.
- Copy from either primary or secondary flash in the source to either primary or secondary flash in the destination.

Downloading from primary only (CLI)

Syntax:

```
copy tftp flash <ip-addr> flash [ primary | secondary ]
```

When executed in the destination switch, downloads the software flash in the source switch's primary flash to either the primary or secondary flash in the destination switch.

If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

To download a software file from primary flash in a switch with an IP address of 10.29.227.103 to the primary flash in the destination switch, you would execute the following command in the destination switch's CLI:

Example 130 Switch-to-switch, from primary in source to either flash in destination

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

█ Running Total of Bytes
Downloaded

Downloading from either flash in the source switch to either flash in the destination switch (CLI)

Syntax:

```
copy tftp flash <ip-addr> </os/primary> | </os/secondary>
[ primary | secondary ]
```

This command (executed in the destination switch) gives you the most options for downloading between switches. If you do not specify either a primary or secondary flash location for the destination, the download automatically goes to primary flash.

To download a software file from secondary flash in a switch with an IP address of 10.28.227.103 to the secondary flash in a destination switch, you would execute the following command in the destination switch's CLI:

Example 131 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
```

Using IMC to update switch software

IMC includes a software update utility for updating on HP switch products. For further information, refer to the *Getting Started Guide* and the *Administrator's Guide*, provided electronically with the application.

Copying software images

NOTE: For details on how switch memory operates, including primary and secondary flash, see "Switch Memory and Configuration" in the *Basic Operation Guide* for your switch.

TFTP: Copying a software image to a remote host (CLI)

Syntax:

```
copy flash tftp <ip-addr> <filename>
```

Copies the primary flash image to a TFTP server.

Example:

To copy the primary flash to a TFTP server having an IP address of 10.28.227.105:

```
HP Switch# copy flash tftp 10.28.227.105 k0800.swi
```

where `k0800.swi` is the filename given to the flash image being copied.

Xmodem: Copying a software image from the switch to a serially connected PC or UNIX workstation (CLI)

To use this method, the switch must be connected via the serial port to a PC or UNIX workstation.

Syntax:

```
copy flash xmodem [<pc> | unix>
```

Uses Xmodem to copy a designated configuration file from the switch to a PC or UNIX workstation.

Example:

To copy the primary flash image to a serially connected PC:

1. Execute the following command:

```
HP Switch# copy xmodem flash
Press 'Enter' and start XMODEM on your host...
```

2. After you see the above prompt, press **[Enter]**.
3. Execute the terminal emulator commands to begin the file transfer.

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

Using the CLI commands described in the section beginning with [“TFTP: Copying a configuration file to a remote host \(CLI\)”](#) (page 241), 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 229).

You can also use the `include-credentials` command to save passwords, secret keys, and other security credentials in the running config file. For more information, see the section on [“Saving Security Credentials in a Config File”](#) in the *Access Security Guide* for your switch.

TFTP: Copying a configuration file to a remote host (CLI)

Syntax:

```
copy <startup-config | running-config> tftp <ip-addr>
<remote-file> [ pc | unix ]
copy config <filename> tftp <ip-addr> <remote-file> [ pc |
unix ]
```

This command can copy a designated config file in the switch to a TFTP server. For more information, see [“Multiple Configuration Files”](#) in the *Basic Operation Guide* for your switch.

Example:

To upload the current startup configuration to a file named `sw8200` in the `configs` directory on drive `"d"` in a TFTP server having an IP address of 10.28.227.105:

```
ProCurve# copy startup-config tftp 10.28.227.105
d:\configs\sw8200
```

TFTP: Copying a configuration file from a remote host (CLI)

Syntax:

```
copy tftp <startup-config | running-config> <ip-address>
<remote-file> [ pc | unix ]
copy tftp config <filename> <ip-address> <remote-file> [ pc
| unix ]
```

This command can copy a configuration from a remote host to a designated config file in the switch. For more information, see [“Multiple Configuration Files”](#) in the *Basic Operation Guide* for your switch.

For more information on flash image use, see [“Using Primary and Secondary Flash Image Options”](#) in the *Basic Operation Guide* for your switch.

Example:

To download a configuration file named `sw8200` in the `configs` directory on drive "`d`" in a remote host having an IP address of 10.28.227.105:

```
HP Switch# copy tftp startup-config 10.28.227.105
d:\configs\sw8200
```

TFTP: Copying a customized command file to a switch (CLI)

Using the `copy tftp` command with the `show-tech` option provides the ability to copy a customized command file to the switch. When the `show tech custom` command is executed, the commands in the custom file are executed instead of the hard-coded list of commands. If no custom file is found, the current hard-coded list is executed. This list contains commands to display data, such as the image stamp, running configuration, boot history, port settings, and so on.

Syntax:

```
copy tftp show-tech <ipv4 or ipv6 address> <filename>
```

Copies a customized command file to the switch (see [Example 132](#)).

Example 132 Using the copy tftp show-tech command to upload a customized command file

```
HP Switch(config)# copy tftp show-tech 10.10.10.3 commandfile1
```

Syntax:

```
show tech custom
```

Executes the commands found in a custom file instead of the hard-coded list.

NOTE: Exit the global config mode (if needed) before executing `show tech` commands.

You can include `show tech` commands in the custom file, with the exception of `show tech custom`. For example, you can include the command `show tech all`.

If no custom file is found, a message displays stating "No SHOW-TECH file found." (No custom file was uploaded with the `copy tftp show-tech` command.)

Example 133 The show tech custom command

```
HP Switch# show tech custom
No SHOW-TECH file found.
```

Xmodem: Copying a configuration file to a serially connected PC or UNIX workstation (CLI)

To use this method, the switch must be connected via the serial port to a PC or UNIX workstation. You will need to:

- Determine a filename to use.
- Know the directory path you will use to store the configuration file.

Syntax:

```
copy <startup-config | running-config> xmodem <pc | unix>
copy config <filename> xmodem <pc | unix>
```

Uses Xmodem to copy a designated configuration file from the switch to a PC or UNIX workstation. For more information, see "Multiple Configuration Files" in the *Basic Operation Guide* for your switch.

Example:

To copy a configuration file to a PC serially connected to the switch:

1. Determine the file name and directory location on the PC.
2. Execute the following command:
HP Switch# copy startup-config xmodem pc
Press 'Enter' and start XMODEM on your host...
3. After you see the above prompt, press **[Enter]**.
4. Execute the terminal emulator commands to begin the file transfer.

Xmodem: Copying a configuration file from a serially connected PC or UNIX workstation (CLI)

To use this method, the switch must be connected via the serial port to a PC or UNIX workstation on which is stored the configuration file you want to copy. To complete the copying, you need to know the name of the file to copy and the drive and directory location of the file.

Syntax:

```
copy xmodem startup-config <pc | unix>  
copy xmodem config <filename> < pc | unix>
```

Copies a configuration file from a serially connected PC or UNIX workstation to a designated configuration file on the switch.

For more information, see "Multiple Configuration Files" in the *Basic Operation Guide* for your switch.

Example:

To copy a configuration file from a PC serially connected to the switch:

1. Execute the following command:
HP Switch# copy xmodem startup-config pc
Device will be rebooted, do you want to continue [y/n]? y
Press 'Enter' and start XMODEM on your host...
2. After you see the above prompt, press **[Enter]**.
3. Execute the terminal emulator commands to begin the file transfer.
4. When the download finishes, you must reboot the switch to implement the newly downloaded software. To do so, use one of the following commands:

Syntax:

```
boot system flash [ primary | secondary ]  
boot system flash [config <filename>]
```

Switches boot from the designated configuration file. For more information, see "Multiple Configuration Files" in the *Basic Operation Guide* for your switch.

Syntax:

```
reload
```

Reboots from the flash image currently in use.

(For more on these commands, see "Rebooting the Switch" in the *Basic Operation Guide* for your switch.)

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 "Multiple Configuration Files" in 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

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

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

Syntax:

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

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

| | |
|-----------------------------------|---|
| <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> |
| <code><unix pc></code> | The type of workstation used for serial, Telnet, or SSH access to the switch CLI |

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

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

Example:

Suppose you:

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

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

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

The switch displays this message:

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

To continue with the upload, press the **[Y]** key. To abort the upload, press the **[N]** key. Note that if the switch detects an illegal (non-ACL) command in the file, it bypasses the illegal command, displays a notice (as shown in [Example 134 \(page 246\)](#)), and continues to implement the remaining ACL commands in the file.

Example 134 Using the `copy` command to download and configure an ACL

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

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

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

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

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

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

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

Syntax:

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

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

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

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

Example:

Suppose you:

1. Created an ACL command file named `vlan10_in.txt` to update an existing ACL.
2. Copied the file to a 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 Example 134 \(page 246\)](#)), and continues to implement the remaining ACL commands in the file.

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**.

Copying command output to a destination device (CLI)

Syntax:

```
copy command-output <"cli-command"> tftp <ip-address>
<filepath-filename>
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, or to a serially connected PC or UNIX workstation.

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.

Example:

To use Xmodem to copy the output of `show config` to a serially connected PC:

Figure 42 Sending command output to a file on an attached PC



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

Copying Event Log output to a destination device (CLI)

Syntax:

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


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

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

Example:

To copy the event log to a PC connected to the switch:

Figure 43 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>  
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.

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

Example 135 Copying switch crash data content to a PC

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

Transfer complete

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

Flight Data Recorder (FDR)

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.

- | | |
|------------|--|
| all | Copies all the log files from both management modules and all slots. |
| mn-active | Copies the active management module's log. |
| mn-standby | Copies the standby management module's log. |
| slot | Retrieves the crash log from the module in the identified slots. |

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 xx.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 [“Using USB to transfer files to and from the switch” \(page 236\)](#).

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” \(page 254\)](#)).
 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 [“Downloading switch software” \(page 224\)](#).
 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” \(page 252\)](#).

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 \(CLI\) \(page 253\)](#)):

- 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: IMC 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 IMC 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” \(page 303\)](#).

Configuring autorun on the switch (CLI)

To enable/disable the autorun feature on the switch, the following commands can be executed from configuration mode in the 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” \(page 253\)](#).

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 chapter "Username and Password Security" in the *Access Security Guide* for your switch.

Viewing autorun configuration information

The `show autorun` command displays autorun configuration status information, as shown in the following Example:

Example 136 The `show autorun` command

```
HP Switch(config)# show autorun

  Autorun configuration status

  Enabled           : Yes
  Secure-mode       : Disabled
  Encryption-key    :
```

10 Monitoring and Analyzing Switch Operation

Overview

The switches have several built-in tools for monitoring, analyzing, and troubleshooting switch and network operation:

- **Status:** Includes options for displaying general switch information, management address data, port status, port and trunk group statistics, MAC addresses detected on each port or VLAN, and STP, IGMP, and VLAN data (“[Status and counters data](#)” (page 255)).
- **Counters:** Display details of traffic volume on individual ports (“[Accessing port and trunk statistics \(Menu\)](#)” (page 262)).
- **Event Log:** Lists switch operating events (“[Using the Event Log for troubleshooting switch problems](#)” (page 303)).
- **Alert Log:** Lists network occurrences detected by the switch—in the System > Logging screen of the WebAgent.
- **Configurable trap receivers:** Uses SNMP to enable management stations on your network to receive SNMP traps from the switch.
- **Port monitoring (mirroring):** Copy all traffic from the specified ports to a designated monitoring port.

NOTE: Link test and ping test—analysis tools in troubleshooting situations—are described in “[Troubleshooting](#)” (page 276). See “[Diagnostic tools](#)” (page 336).

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 44 The Status and Counters menu

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

1. General System Information
2. Switch Management Address Information
3. Module Information
4. Port Status
5. Port Counters
6. Vlan Address Table
7. Port Address Table
8. Spanning Tree Information
0. Return to Main Menu...

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

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

General system information

Accessing system information (Menu)

From the console Main Menu, select:

1. Status and Counters
1. General System Information

Figure 45 Example: of general switch information

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

System Contact      :
System Location     :

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

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

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

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

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

Accessing system information (CLI)

Syntax:

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

Displays global system information and operational parameters for the switch.

| | |
|---------------|--|
| chassislocate | Shows the chassisLocator LED status. Possible values are On, Off, or Blink. When the status is On or Blink, the number of minutes that the Locator LED will continue to be on or to blink is displayed. (See Example 137 (page 257)) |
| information | Shows global system information and operational parameters for the switch. (See Example 139 (page 257) .) |
| power-supply | Shows chassis power supply and settings. |
| temperature | Shows system temperature and settings. |
| fans | Shows system fan status. (See Example 138 (page 257) .) |

Example 137 Command results for show system chassislocate command

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

Example 138 System fan status

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

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

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

Example 139 Switch system information

```
HP Switch(config)# show system

Status and Counters - General System Information

System Name       : HP Switch
System Contact    :
System Location   :

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None

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

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

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

Collecting processor data with the task monitor (CLI)

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

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

Syntax:

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

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

Example 140 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

```

Task usage reporting

The task usage reporting feature provides the ability to collect and display CPU usage data (with a refresh rate of 5 seconds) of running tasks on the switch. It includes the following commands:

- `process-tracking`: This command is used to enable/disable the task-usage collecting capability for a specific module on the switch.
- `show cpu process`: This command is used to display task-usage statistics for a specific module.

Syntax:

```
[no] process-tracking [slot[SLOT-LIST] [<time>]] [<time>]
```

Enables/disables module process-tracking functionality.

```
process-tracking <tab>
```

| | |
|---------|---|
| slot | Enables/disables process-tracking for a module. |
| INTEGER | Specifies time to track value between 1 second to 30 seconds. |
| <cr> | |

```
process-tracking slot <tab>
```

| | |
|---------------|---|
| SLOT-ID-RANGE | Enter an alphabetic device slot identifier or slot range. |
|---------------|---|

```
process-tracking slot A
```

| | |
|---------|---|
| INTEGER | Specifies time to track value between 1 second to 30 seconds. |
| <cr> | |

```
process-tracking slot A 10 <tab>
```

| | |
|------|--|
| <cr> | |
|------|--|

```
process-tracking 10 <tab>
```

| | |
|------|--|
| <cr> | |
|------|--|

Syntax:

```
show cpu  
[<CHASSIS_MIN_CPU_UTIL_INDEX-CHASSIS_MAX_CPU_UTIL_INDEX>]  
[slot <SLOT-LIST>  
[<CHASSIS_MIN_CPU_UTIL_INDEX-CHASSIS_MODULE_MAX_CPU_UTIL_INDEX>]]  
[process [[slot <SLOT-LIST>] [refresh <iterations>]]  
[refresh <iterations>]
```

Shows average CPU utilization over the last 1, 5, and 60 seconds, or the number of seconds specified.

Use the `slot` option to display CPU utilization for the specified modules, rather than the chassis CPU.

Use the `process` option to display module process usages.

Syntax:

```
show cpu process [slot [SLOT-LIST] [refresh <iterations>]]  
[refresh <iterations>]
```

Displays module process usage.

```
show cpu <tab>
```

| | |
|---------|--|
| process | Displays process usage. |
| slot | Displays module CPU statistics. |
| <1-300> | Time (in seconds) over which to average CPU utilization. |
| <cr> | |

```
show cpu process <tab>
```

| | |
|---------|---|
| refresh | Number of times to refresh process usage display. |
| slot | Displays module process usage. |
| <cr> | |

```
show cpu process refresh <tab>
```

| | |
|---------|--------------------------|
| INTEGER | Enter an integer number. |
|---------|--------------------------|

```
show cpu process refresh 10 <tab>
```

| | |
|------|--|
| <cr> | |
|------|--|

```
show cpu process slot <tab>
```

| | |
|---------------|---|
| SLOT-ID-RANGE | Enter an alphabetic device slot identifier or slot range. |
|---------------|---|

show cpu process slot A <tab>

| | |
|---------|---|
| refresh | Number of times to refresh process usage display. |
| <cr> | |

show cpu process slot A refresh <tab>

| | |
|---------|--------------------------|
| INTEGER | Enter an integer number. |
|---------|--------------------------|

show cpu process slot A refresh 10 <tab>

| | |
|------|--|
| <cr> | |
|------|--|

Example 141 Output for the show cpu process command

HP Switch# show cpu process

| Process Name | Priority | Recent Time | % CPU | Time Since Last Ran | Times Ran | Max Time |
|-------------------|----------|-------------|-------|---------------------|-----------|----------|
| Idle-1 | 226 | 10 s | 41 | 57 us | 380986 | 69 us |
| Idle-3 | 1 | 5 s | 20 | 52 us | 761665 | 55 us |
| Idle-0 | 226 | 8 s | 33 | 19 us | 380867 | 66 us |
| Sessions & I/O-24 | 171 | 926 ms | 3 | 1 ms | 150 | 335 ms |

Example 142 Output for the show cpu process slot <slot-list> command

HP Switch# show cpu process slot A
slot a:

| Process Name | Priority | Recent Time | % CPU | Time Since Last Ran | Times Ran | Max Time |
|-------------------|----------|-------------|-------|---------------------|-----------|----------|
| System Services-2 | 156 | 253 ms | 2 | 767 ms | 12 | 35 ms |
| Idle-3 | 1 | 3 s | 28 | 13 ms | 101309 | 150 us |
| Hardware Mgmt-2 | 192 | 282 ms | 2 | 303 us | 44 | 12 ms |
| Idle-1 | 226 | 6 s | 55 | 13 ms | 50793 | 233 us |
| Idle-0 | 226 | 1 s | 9 | 14 ms | 50633 | 106 us |

Switch management address information

Accessing switch management address information (Menu)

From the Main Menu, select:

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

Figure 46 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 46 \(page 261\)](#), all VLANs on the switches use the same **MAC address**. (This includes both the statically configured VLANs and any dynamic VLANs existing on the switch as a result of GVRP operation.)

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

Accessing switch management address information (CLI)

Syntax:

```
show management
```

Port Status

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

Viewing port status (CLI)

Syntax:

```
show interfaces brief
```

Viewing port status (Menu)

From the Main Menu, select:

- 1. Status and Counters ...**
- 4. Port Status**

Figure 47 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.
-----
```

Viewing port and trunk group statistics (WebAgent)

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

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

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

Port and trunk group statistics and flow control status

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

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

You can also reset the counters for a specific port.

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

As mentioned above, rebooting or resetting the switch resets the counters to zero. You can also reset the counters to zero for the current session. This is useful for troubleshooting. See the Note, below.

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

Accessing port and trunk statistics (Menu)

From the Main Menu, select:

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

Figure 48 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 49 (page 263), below.

Figure 49 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.

NOTE: Once cleared, statistics cannot be reintroduced.

Accessing port and trunk group statistics (CLI)

Viewing the port counter summary report

Syntax:

```
show interfaces
```

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

Viewing a detailed traffic summary for specific ports

Syntax:

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

Provides traffic details for the ports you specify.

Displaying trunk load balancing statistics

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

Syntax:

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

Displays the trunk counter information since the trunk was formed.

Example 143 Output for the show trunk-statistics command

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

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

Totals

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

Rates (5 minute weighted average):

```
Trunk utilization Rx : 30.2 %  
Trunk utilization Tx : 78.2 %
```

Traffic Spread past 5 minutes

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

Clearing trunk load balancing statistics

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

Syntax:

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

Clears statistics for all trunks if no trunks identified.

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

Resetting the port counters

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

Syntax:

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

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

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

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

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

Viewing the switch's MAC address tables

Accessing MAC address views and searches (CLI)

Syntax:

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

[Listing all learned MAC addresses on the switch, with the port number on which each MAC address was learned](#)

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

[Listing all learned MAC addresses on one or more ports, with their corresponding port numbers](#)

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

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

[Listing all learned MAC addresses on a VLAN, with their port numbers](#)

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

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

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

[Finding the port on which the switch learned a specific MAC address](#)

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

```
| select VLAN : DEFAULT_VLAN |
```

Accessing MAC address views and searches (Menu)

[Viewing and searching per-VLAN MAC-addresses](#)

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

From the Main Menu, select:

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

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

```

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

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

Actions->  Back  Search  Next page  Prev page  Help

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

```

2. Use the Space bar to select the VLAN you want, and then press **[Enter]**.
The switch then displays the MAC address table for that VLAN (Figure 50 (page 266)).

Figure 50 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 50 (page 266), 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 51 (page 266)). If the switch does not find the MAC address on the currently selected VLAN, it leaves the MAC address listing empty.

Figure 51 Example: of menu indicating located MAC address

```

----- CONSOLE - MANAGER MODE -----
                          Status and Counters Menu

  1. General System Information
  2. Switch Management Address Information
  3. Module Information
  4. Port Status
  5. Port Counters
  6. Vlan Address Table
  7. Port Address Table
  8. Spanning Tree Information
  0. Return to Main Menu...

Select port : A1

Type port number or press <Space> to scroll ports. Press <Enter> to select.
To select menu item, press item number, or highlight item and press <Enter>.

```

3. Press **[P]** (for **Prev page**) to return to the full address table listing.

Viewing and searching port-level MAC addresses

This feature displays and searches for MAC addresses on the specified port instead of for all ports on the switch.

1. From the Main Menu, select:
 1. **Status and Counters ...**
 7. **Port Address Table**
2. Use the Space bar to select the port you want to list or search for MAC addresses, then press **[Enter]** to list the MAC addresses detected on that port.

Determining whether a specific device is connected to the selected port

Proceeding from [step 2 \(page 267\)](#), above:

1. Press **[S]** (for **Search**), to display the following prompt:
Enter MAC address: _
2. Enter the MAC address you want to locate and press **[Enter]**.
The address is highlighted if found. If the switch does not find the address, it leaves the MAC address listing empty.
3. Press **[P]** (for **Prev page**) to return to the previous per-port listing.

Accessing MSTP Data (CLI)

Syntax:

```
show spanning-tree
```

Displays the switch's global and regional spanning-tree status, plus the per-port spanning-tree operation at the regional level.

Values for the following parameters appear only for ports connected to active devices: Designated Bridge, Hello Time, PtP, and Edge.

Example:

Figure 52 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 | Prio rity | 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)

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

| Show command | Output |
|------------------------|---|
| show ip igmp | Global command listing IGMP status for all VLANs configured in the switch: <ul style="list-style-type: none"> VLAN ID (VID) and name Querier address Active group addresses per VLAN Number of report and query packets per group Querier access port per VLAN |
| show ip igmp config | Displays the IGMP configuration information, including VLAN ID, VLAN name, status, forwarding, and Querier information. |
| show ip igmp <vlan-id> | Per-VLAN command listing above, IGMP status for specified VLAN (VID) |

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

Example 144 Output from show ip igmp config command

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

```
IGMP Service
```

| VLAN ID | VLAN Name | IGMP Enabled | Forward with High Priority | Querier Allowed | Querier Internal |
|---------|--------------|--------------|----------------------------|-----------------|------------------|
| 1 | DEFAULT_VLAN | No | No | Yes | 125 |
| 2 | VLAN2 | Yes | No | Yes | 125 |
| 12 | New_VLAN | No | No | Yes | 125 |

Example 145 IGMP statistical information

```
HP Switch(vlan-2)# show ip igmp statistics
```

```
IGMP Service Statistics
```

```
Total VLANs with IGMP enabled           : 1
Current count of multicast groups joined  : 1
```

```
IGMP Joined Groups Statistics
```

| VLAN ID | VLAN Name | Filtered | Flood |
|---------|-----------|----------|-------|
| 2 | VLAN2 | 2 | 1 |

Viewing VLAN information (CLI)

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

Example:

Suppose that your switch has the following VLANs:

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

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

Example 146 Listing the VLAN ID (vid) and status for specific ports

```
HP Switch# show vlan ports A1-A2
```

```
Status and Counters = VLAN Information - for ports A1,A2
```

```
802.1Q VLAN ID Name          Status
-----
1          DEFAULT_VLAN Static
33         VLAN-33      Static
```

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

Example 147 VLAN listing for the entire switch

```
HP Switch# show vlan
```

```
Status and Counters = VLAN Information
```

```
VLAN support : Yes
Maximum VLANs to support : 9
Primary VLAN: DEFAULT_VLAN
```

```
802.1Q VLAN ID Name          Status
-----
1          DEFAULT_VLAN Static
33         VLAN-33      Static
44         VLAN-44      Static
```

Example 148 Port listing for an individual VLAN

```
HP Switch(config)# show vlan 1
```

```
Status and Counters - VLAN Information - VLAN 1
```

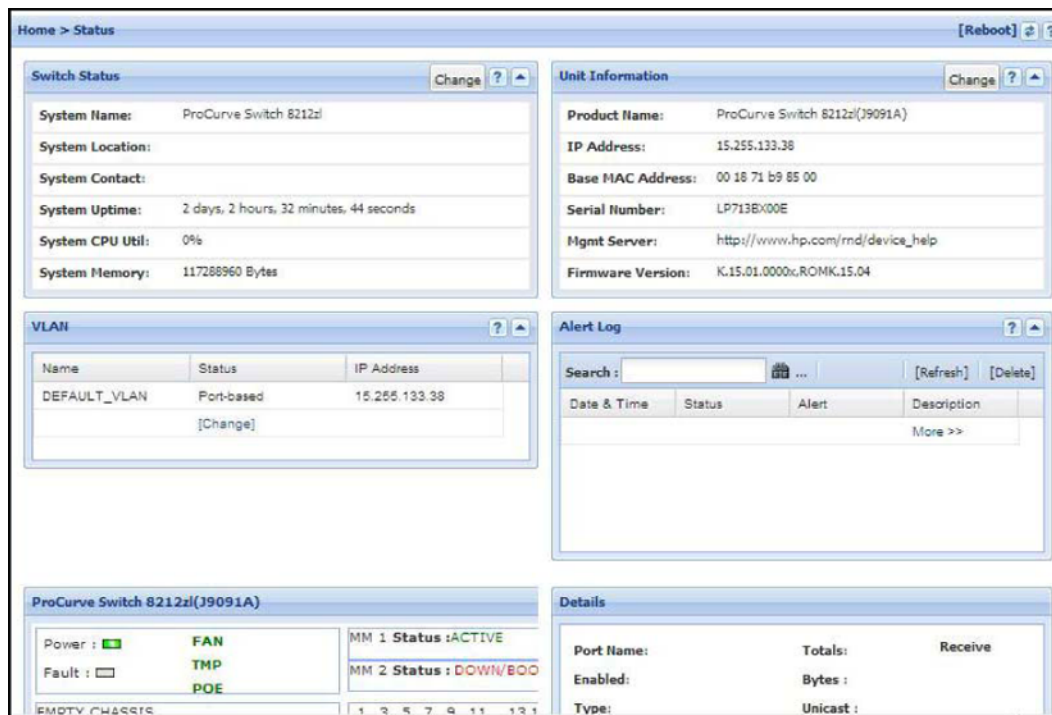
```
VLAN ID : 1
Name : DEFAULT_VLAN
Status : Static
Voice : Yes
Jumbo : No
```

```
Port Information Mode      Unknown VLAN Status
-----
A1          Untagged Learn          Up
A2          Untagged Learn          Up
A3          Untagged Learn          Up
A4          Untagged Learn          Down
A5          Untagged Learn          Up
A6          Untagged Learn          Up
A7          Untagged Learn          Up
```

WebAgent status information

The WebAgent Status screen provides an overview of the status of the switch. Scroll down to view more details. For information about this screen, click on **?** in the upper right corner of the WebAgent screen. For an Example: of a status screen, see [Figure 53 \(page 272\)](#).

Figure 53 Example: of a WebAgent status screen



Interface monitoring features

You can designate monitoring of inbound and outbound traffic on:

- **Ports and static trunks:** Allows monitoring of individual ports, groups of contiguous ports, and static port trunks.
- **Static VLANs:** Allows traffic monitoring on one static VLAN.

The switch monitors network activity by copying all traffic inbound and outbound on the specified interfaces to the designated monitoring port, to which a network analyzer can be attached.

If a tagged packet arrives on a monitored port, the packet will remain tagged when it goes out a monitored port even if that port is configured as untagged. If the packet is untagged, it will remain untagged going out the monitor port. The monitor port state (tagged or untagged) does not affect the tagging of the packet. However, egress mirroring does not reflect the tagged or untagged characteristic to the mirror port, instead it reflects the tagged or untagged characteristic of the mirror port.

NOTE: When both inbound and outbound monitoring is done, and IGMP is enabled on any VLAN, you may get two copies of IGMP packets on the monitored port.

NOTE: VLANs and port trunks cannot be used as a monitoring port.

The switch can monitor static LACP trunks, but not dynamic LACP trunks.

It is possible, when monitoring multiple interfaces in networks with high traffic levels, to copy more traffic to a monitor port than the link can support. In this case, some packets may not be copied to the monitor port.

Configuring port and static trunk monitoring (Menu)

This procedure describes configuring the switch for monitoring when monitoring is disabled. (If monitoring has already been enabled, the screens will appear differently than shown in this procedure.)

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

3. Network Monitoring Port

2. In the Actions menu, press **[E]** (for Edit).
3. If monitoring is currently disabled (the default) then enable it by pressing the Space bar (or **[Y]**) to select Yes.
4. Press the down arrow key to display a screen similar to the following and move the cursor to the **Monitoring Port** parameter.
5. Use the Space bar to select the port to use for monitoring.
6. Highlight the Monitor field and use the Space bar to select the interfaces to monitor:
 - Ports:** Use for monitoring ports or static trunks.
 - VLAN:** Use for monitoring a VLAN.
7. Do one of the following:
 - If you are monitoring ports or static trunks go to step 8.
 - If you are monitoring a VLAN:
 - a. i. Press **[Tab]** or the down arrow key to move to the **VLAN** field.
 - b. Use the Space bar to select the VLAN you want to monitor.
 - c. Go to step 10.
8. Use the down arrow key to move the cursor to the **Action** column for the individual ports and position the cursor at a port you want to monitor.
9. Press the Space bar to select **Monitor** for each port and trunk that you want monitored. (Use the down arrow key to move from one interface to the next in the **Action** column.)
10. When you finish selecting ports to monitor, press **[Enter]**, then press **[S]** (for **Save**) to save your changes and exit from the screen.
11. Return to the Main Menu.

Configuring port and static trunk monitoring (CLI)

You must use the following configuration sequence to configure port and static trunk monitoring in the CLI:

1. Assign a monitoring (mirror) port.
2. Designate the port(s) and/or static trunk(s) to monitor.

Displaying the monitoring configuration

Syntax:

```
show monitor
```

This command lists the port assigned to receive monitored traffic and the ports and/or trunks being monitored.

For example, if you assign port 5 as the monitoring port and configure the switch to monitor ports 2-4, `show monitor` displays the following:

Example 149 Monitored port listing

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

```
Network Monitoring Port
```

```
Mirror Port: 5 1
```

```
Monitoring sources 2
```

```
-----
```

```
2
```

```
3
```

```
4
```

- 1** Port receiving monitored traffic. **2** Monitored Ports
-

Configuring the monitor port

Syntax:

```
[no] mirror-port [<port-num>]
```

This command assigns or removes a monitoring port, and must be executed from the global configuration level. Removing the monitor port disables port monitoring and resets the monitoring parameters to their factory-default settings.

For example, to assign port 6 as the monitoring port:

```
HP Switch(config)# mirror-port 6
```

To turn off monitoring:

```
HP Switch(config)# no mirror-port
```

Selecting or removing monitoring source interfaces

After you configure a monitor port you can use either the global configuration level or the interface context level to select ports, static trunks, or VLANs as monitoring sources. You can also use either level to remove monitoring sources.

Syntax:

```
[no] interface <monitor-list> monitor
```

| | |
|----------------|---|
| <monitor-list> | Includes port numbers and static trunk names such as 4, 7, 5-8, trk1 . |
|----------------|---|

NOTE: Individual ports and static trunks can be monitored at the same time. However, if you configure the switch to monitor a VLAN, all other interfaces are removed from monitoring. Also, you can configure only one VLAN at a time for monitoring.

Elements in the monitor list can include port numbers and static trunk names at the same time.

For example, with a port such as port 5 configured as the monitoring (mirror) port, you would use either of the following commands to select these interfaces for monitoring:

- Ports 6-9, and 14
- Trunk 2

Example 150 Selecting ports and static trunks as monitoring sources

```
HP Switch(config)# int 6-9, 14 trk2, monitor
```

To monitor a VLAN:

Example 151 Configuring VLAN monitoring

```
HP Switch(config)# vlan 20 monitor
```

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

```
Network Monitoring Port
```

```
  Mirror Port: 5
```

```
  Monitoring sources
```

```
  -----
```

```
  VLAN_20
```

Example 152 Disabling monitoring at the interface context and the global config level

```
HP Switch(eth-1-3, 5)# no int 5 monitor 1
```

```
HP Switch(eth-1-3, 5)# no monitor
```

```
HP Switch(config)# no int 5 monitor 2
```

```
HP Switch(config)# no int 1-3, 5 monitor
```

- | | |
|--|---|
| <p>1 These two commands show how to disable monitoring at the interface context level for a single port or all ports in an interface context level.</p> | <p>2 These two commands show how to disable monitoring at the global config level for a single port or a group of ports.</p> |
|--|---|
-

11 Troubleshooting

Overview

This appendix addresses performance-related network problems that can be caused by topology, switch configuration, and the effects of other devices or their configurations on switch operation. (For switch-specific information on hardware problems indicated by LED behavior, cabling requirements, and other potential hardware-related problems, see the *Installation Guide* you received with the switch.)

NOTE: HP periodically places switch software updates on the HP Switch Networking website. HP Switch recommends that you check this website for software updates that may have fixed a problem you are experiencing.

For information on support and warranty provisions, see the Support and Warranty booklet shipped with the switch.

Troubleshooting approaches

Use these approaches to diagnose switch problems:

- Check the HP website for software updates that may have solved your problem: www.hp.com/networking
- Check the switch LEDs for indications of proper switch operation:
 - Each switch port has a Link LED that should light whenever an active network device is connected to the port.
 - Problems with the switch hardware and software are indicated by flashing the Fault and other switch LEDs.
For a description of the LED behavior and information on using the LEDs for troubleshooting, see the *Installation Guide* shipped with the switch.
- Check the network topology/installation. For topology information, see the *Installation Guide* shipped with the switch.
- Check cables for damage, correct type, and proper connections. You should also use a cable tester to check your cables for compliance to the relevant IEEE 802.3 specification. For correct cable types and connector pin-outs, see the *Installation Guide* shipped with the switch.
- Use HP PCM+ to help isolate problems and recommend solutions.
- Use the Port Utilization Graph and Alert Log in the WebAgent included in the switch to help isolate problems. These tools are available through the WebAgent:
 - Port Utilization Graph
 - Alert log

- Port Status and Port Counters screens
- Diagnostic tools (Link test, Ping test, configuration file browser)
- For help in isolating problems, use the easy-to-access switch console built into the switch or Telnet to the switch console. For operating information on the Menu and CLI interfaces included in the console, see chapters 3 and 4. These tools are available through the switch console:
 - Status and Counters screens
 - Event Log
 - Diagnostics tools (Link test, Ping test, configuration file browser, and advanced user commands)

Browser or Telnet access problems

Cannot access the WebAgent

- Access may be disabled by the Web Agent Enabled parameter in the switch console. Check the setting on this parameter by selecting:
 - 2. Switch Configuration**
 - 1. System Information**
- The switch may not have the correct IP address, subnet mask, or gateway. Verify by connecting a console to the switch's Console port and selecting:
 - 2. Switch Configuration**
 - 5. IP Configuration**

Note: If DHCP/Bootp is used to configure the switch, the IP addressing can be verified by selecting:

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

Also check the DHCP/Bootp server configuration to verify correct IP addressing.

- If you are using DHCP to acquire the IP address for the switch, the IP address "lease time" may have expired so that the IP address has changed. For more information on how to "reserve" an IP address, see the documentation for the DHCP application that you are using.
- If one or more IP-authorized managers are configured, the switch allows inbound telnet access only to a device having an authorized IP address. For more information on IP Authorized managers, see the Access Security Guide for your switch.
- Java™ applets may not be running on the web browser. They are required for the switch WebAgent to operate correctly. Refer to the online Help on your web browser for how to run the Java applets.

Cannot Telnet into the switch console from a station on the network

- Off-subnet management stations can lose Telnet access if you enable routing without first configuring a static (default) route. That is, the switch uses the IP default gateway only while operating as a Layer 2 device. While routing is enabled on the switch, the IP default gateway is not used. You can avoid this problem by using the ip route command to configure a static (default) route before enabling routing. For more information, see chapter "IP Routing Features" in the *Multicast and Routing Guide* for your switch.
- Telnet access may be disabled by the Inbound Telnet Enabled parameter in the System Information screen of the menu interface:
 - 2. Switch Configuration**

1. System Information

- The switch may not have the correct IP address, subnet mask, or gateway. Verify by connecting a console to the switch's Console port and selecting:

2. Switch Configuration

5. IP Configuration

- If you are using DHCP to acquire the IP address for the switch, the IP address "lease time" may have expired so that the IP address has changed. For more information on how to "reserve" an IP address, see the documentation for the DHCP application that you are using.
- If one or more IP-authorized managers are configured, the switch allows inbound telnet access only to a device having an authorized IP address. For more information on IP Authorized managers, see the *Access Security Guide* for your switch.

Unusual network activity

Network activity that fails to meet accepted norms may indicate a hardware problem with one or more of the network components, possibly including the switch. Such problems can also be caused by a network loop or simply too much traffic for the network as it is currently designed and implemented. Unusual network activity is usually indicated by the LEDs on the front of the switch or measured with the switchconsole interface or with a network management tool such as HP PCM+. For information on using LEDs to identify unusual network activity, see the *Installation Guide* you received with the switch.

A topology loop can also cause excessive network activity. The Event Log "FFI" messages can be indicative of this type of problem.

General problems

The network runs slow; processes fail; users cannot access servers or other devices

Broadcast storms may be occurring in the network. These may be caused by redundant links between nodes.

- If you are configuring a port trunk, finish configuring the ports in the trunk before connecting the related cables. Otherwise you may inadvertently create a number of redundant links (that is, topology loops) that will cause broadcast storms.
- Turn on STP to block redundant links
- Check for FFI messages in the Event Log

Duplicate IP addresses

This is indicated by this Event Log message:

```
ip: Invalid ARP source: IP address on IP address
```

where both instances of *IP address* are the same address, indicating that the switch's IP address has been duplicated somewhere on the network.

Duplicate IP addresses in a DHCP network

If you use a DHCP server to assign IP addresses in your network, and you find a device with a valid IP address that does not appear to communicate properly with the server or other devices, a duplicate IP address may have been issued by the server. This can occur if a client has not released a DHCP-assigned IP address after the intended expiration time and the server "leases" the address to another device. This can also happen, For example, if the server is first configured to issue IP addresses with an unlimited duration, and then is subsequently configured to issue IP addresses that will expire after a limited duration. One solution is to configure "reservations" in

the DHCP server for specific IP addresses to be assigned to devices having specific MAC addresses. For more information, see the documentation for the DHCP server.

One indication of a duplicate IP address in a DHCP network is this Event Log message:

```
ip: Invalid ARP source: <IP-address> on <IP-address>
```

where both instances of *IP-address* are the same address, indicating that the IP address has been duplicated somewhere on the network.

The switch has been configured for DHCP/Bootp operation, but has not received a DHCP or Bootp reply

When the switch is first configured for DHCP/Bootp operation, or if it is rebooted with this configuration, it immediately begins sending request packets on the network. If the switch does not receive a reply to its DHCP/Bootp requests, it continues to periodically send request packets, but with decreasing frequency. Thus, if a DHCP or Bootp server is not available or accessible to the switch when DHCP/Bootp is first configured, the switch may not immediately receive the desired configuration.

After verifying that the server has become accessible to the switch, reboot the switch to re-start the process.

802.1Q Prioritization problems

Ports configured for non-default prioritization (level 1 to 7) are not performing the specified action

If the ports were placed in a trunk group after being configured for non-default prioritization, the priority setting was automatically reset to zero (the default). Ports in a trunk group operate only at the default priority setting.

Addressing ACL problems

ACLs are properly configured and assigned to VLANs, but the switch is not using the ACLs to filter IP layer 3 packets

1. The switch may be running with IP routing disabled. To ensure that IP routing is enabled, execute `show running` and look for the IP routing statement in the resulting listing. For Example:

Example 153 Indication that routing is enabled

```
HP Switch(config)# show running
Running configuration:
; J9091A Configuration Editor; Created on release #XX.15.06
hostname " HPswitch "
ip default-gateway 10.33.248.1
ip routing 1
logging 10.28.227.2
snmp-server community "public" Unrestricted
ip access-list extended "Controls for VLAN 20"
permit tcp 0.0.0.0 255.255.255.255 10.10.20.98 0.0.0.0 eq 80
permit tcp 0.0.0.0 255.255.255.255 10.10.20.21 0.0.0.0 eq 80
deny tcp 0.0.0.0 255.255.255.255 10.10.20.1 0.0.0.255 eq 80
deny tcp 10.10.20.1? 0.0.0.0 10.10.10.100 0.0.0.0 eq 20 log
deny tcp 10.10.20.20 0.0.0.0 10.10.10.100 0.0.0.0 eq 20 log
deny tcp 10.10.20.43 0.0.0.0 10.10.10.100 0.0.0.0 eq 20 log
permit ip 10.10.20.1 0.0.0.255 10.10.10.100 0.0.0.0
deny ip 10.10.30.1 0.0.0.255 10.10.10.100 0.0.0.0
permit ip 10.10.30.1 0.0.0.255 10.10.10.1 0.0.0.255
exit
```

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

NOTE: If an ACL assigned to a VLAN includes an ACE referencing an IP address on the switch itself as a packet source or destination, the ACE screens traffic to or from this switch address regardless of whether IP routing is enabled. This is a security measure designed to help protect the switch from unauthorized management access.

If you need to configure IP routing, execute the `ip routing` command.

2. ACL filtering on the switches applies only to routed packets and packets having a destination IP address (DA) on the switch itself.

Also, the switch applies assigned ACLs only at the point where traffic enters or leaves the switch on a VLAN. Ensure that you have correctly applied your ACLs ("in" and/or "out") to the appropriate VLANs.

The switch does not allow management access from a device on the same VLAN

The implicit `deny any` function that the switch automatically applies as the last entry in any ACL always blocks packets having the same DA as the switch's IP address on the same VLAN. That is, bridged packets with the switch itself as the destination are blocked as a security measure.

To preempt this action, edit the ACL to include an ACE that permits access to the switch's DA on that VLAN from the management device.

Error (Invalid input) when entering an IP address

When using the "host" option in the Command syntax, ensure that you are not including a mask in either dotted decimal or CIDR format. Using the "host" option implies a specific host device and therefore does not permit any mask entry.

Example 154 Correctly and incorrectly specifying a single host

```
Switch(config)# access-list 6 permit host 10.28.100.100 1
```

```
Switch(config)# access-list 6 permit host 10.28.100.100 255.255.255.255 2  
Invalid input: 255.255.255.255
```

```
Switch(config)# access-list 6 permit host 10.28.100.100/32 3  
Invalid input: 10.28.100.100/32
```

1 Correct.

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

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

Apparent failure to log all "deny" matches

Where the `log` statement is included in multiple ACEs configured with a "deny" option, a large volume of "deny" matches generating logging messages in a short period of time can impact switch performance. If it appears that the switch is not consistently logging all "deny" matches, try reducing the number of logging actions by removing the `log` statement from some ACEs configured with the "deny" action.

The switch does not allow any routed access from a specific host, group of hosts, or subnet

The implicit `deny any` function that the switch automatically applies as the last entry in any ACL may be blocking all access by devices not specifically permitted by an entry in an ACL affecting those sources. If you are using the ACL to block specific hosts, a group of hosts, or a subnet, but want to allow any access not specifically permitted, insert `permit any` as the last explicit entry in the ACL.

The switch is not performing routing functions on a VLAN

Two possible causes of this problem are:

- Routing is not enabled. If `show running` indicates that routing is not enabled, use the `ip routing` command to enable routing.
- An ACL may be blocking access to the VLAN (on a switch covered in this guide). Ensure that the switch's IP address on the VLAN is not blocked by one of the ACE entries in an ACL applied to that VLAN. A common mistake is to either not explicitly permit the switch's IP address as a DA or to use a wildcard ACL mask in a `deny` statement that happens to include the switch's IP address. For an Example: of this problem, see section "General ACL Operating Notes" in the "Access Control Lists (ACLs)" chapter of the latest *Access Security Guide* for your switch.

Routing through a gateway on the switch fails

Configuring a "deny" ACE that includes a gateway address can block traffic attempting to use the gateway as a next-hop.

Remote gateway case

Configuring ACL "101" (Example 155 (page 282)) and applying it outbound on VLAN 1 in Figure 54 (page 282) includes the router gateway (10.0.8.1) needed by devices on other networks. This can prevent the switch from sending ARP and other routing messages to the gateway router to support traffic from authorized remote networks.

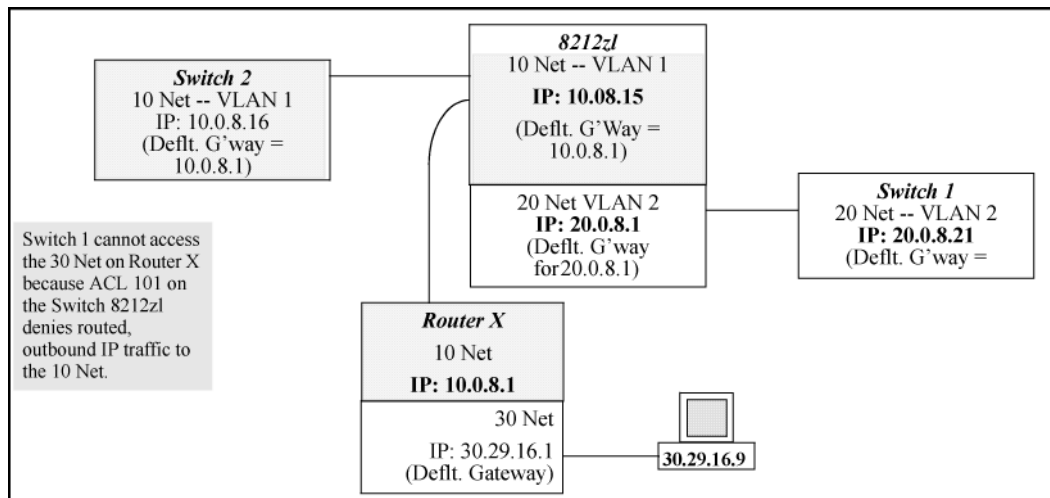
In Figure 54 (page 282), this ACE (see data in bold below) denies access to the 10 Net's 10.0.8.1 router gateway needed by the 20 Net (Subnet mask is 255.255.255.0).

Example 155 ACE blocking an entire subnet

```
HP Switch(config)# access-list config

ip access-list extended "101"
deny ip 0.0.0.0 255.255.255.255 10.0.8.30 0.0.0.255
permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
exit
```

Figure 54 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 282).

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.
-

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 285).

The switch does not receive a response to RADIUS authentication requests

In this case, the switch attempts authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use `ping` to ensure that the switch has access to the configured RADIUS servers.
- Verify that the switch is using the correct encryption key (RADIUS secret key) for each server.
- Verify that the switch has the correct IP address for each RADIUS server.
- Ensure that the `radius-server timeout` period is long enough for network conditions.

The switch does not authenticate a client even though the RADIUS server is properly configured and providing a response to the authentication request

If the RADIUS server configuration for authenticating the client includes a VLAN assignment, ensure that the VLAN exists as a static VLAN on the switch. See "How 802.1X Authentication Affects VLAN Operation" in the *Access Security Guide* for your switch.

During RADIUS-authenticated client sessions, access to a VLAN on the port used for the client sessions is lost

If the affected VLAN is configured as untagged on the port, it may be temporarily blocked on that port during an 802.1X session. This is because the switch has temporarily assigned another VLAN as untagged on the port to support the client access, as specified in the response from the RADIUS server. See "How 802.1X Authentication Affects VLAN Operation" in the *Access Security Guide* for your switch.

The switch appears to be properly configured as a supplicant, but cannot gain access to the intended authenticator port on the switch to which it is connected

If `aaa authentication port-access` is configured for Local, ensure that you have entered the local *login* (operator-level) username and password of the authenticator switch into the `identity` and `secret` parameters of the supplicant configuration. If instead, you enter the `enable` (manager-level) username and password, access will be denied.

The supplicant statistics listing shows multiple ports with the same authenticator MAC address

The link to the authenticator may have been moved from one port to another without the supplicant statistics having been cleared from the first port. See "Note on Supplicant Statistics" in the chapter on Port-Based and User-Based Access Control in the *Access Security Guide* for your switch.

The `show port-access authenticator <port-list>` command shows one or more ports remain open after they have been configured with `control unauthorized`

802.1X is not active on the switch. After you execute `aaa port-access authenticator active`, all ports configured with `control unauthorized` should be listed as Closed.

Example 156 Authenticator ports remain "open" until activated

```
HP Switch(config)# show port-access authenticator e 9
Port Access Authenticator Status
  Port-access authenticator activated [No] : No
                Access Authenticator Authenticator
Port Status Control State Backend State
-----
9   Open 1   FU           Force Auth   Idle

Switch(config)# show port-access authenticator active
Switch(config)# show port-access authenticator e 9
Port Access Authenticator Status
  Port-access authenticator activated [No] : Yes
                Access Authenticator Authenticator
Port Status Control State Backend State
-----
9   Closed FU           Force Unauth Idle
```

- 1** Port A9 shows an "Open" status even though Access Control is set to Unauthorized (Force Auth). This is because the `port-access authenticator` has not yet been activated.
-

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch

Use `show radius` to verify that the encryption key (RADIUS secret key) the switch is using is correct for the server being contacted. If the switch has only a global key configured, it either must

match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, it overrides the global key and must match the server key.

Example 157 Displaying encryption keys

```
HP Switch(config)# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key : My-Global-Key
  Dynamic Authorization UDP Port : 3799

Server IP Addr      Auth Acct DM/ Time
-----
10.33.18.119      1812 1813          119-only-key
```

Also, ensure that the switch port used to access the RADIUS server is not blocked by an 802.1X configuration on that port. For example, `show port-access authenticator <port-list>` gives you the status for the specified ports. Also, ensure that other factors, such as port security or any 802.1X configuration on the RADIUS server are not blocking the link.

The authorized MAC address on a port that is configured for both 802.1X and port security either changes or is re-acquired after execution of `aaa port-access authenticator <port-list> initialize`

If the port is force-authorized with `aaa port-access authenticator <port-list> control authorized` command and port security is enabled on the port, then executing `initialize` causes the port to clear the learned address and learn a new address from the first packet it receives after you execute `initialize`.

A trunked port configured for 802.1X is blocked

If you are using RADIUS authentication and the RADIUS server specifies a VLAN for the port, the switch allows authentication, but blocks the port. To eliminate this problem, either remove the port from the trunk or reconfigure the RADIUS server to avoid specifying a VLAN.

QoS-related problems

Loss of communication when using VLAN-tagged traffic

If you cannot communicate with a device in a tagged VLAN environment, ensure that the device either supports VLAN tagged traffic or is connected to a VLAN port that is configured as `Untagged`.

Radius-related problems

The switch does not receive a response to RADIUS authentication requests

In this case, the switch attempts authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use `ping` to ensure that the switch has access to the configured RADIUS server.
- Verify that the switch is using the correct encryption key for the designated server.
- Verify that the switch has the correct IP address for the RADIUS server.

- Ensure that the `radius-server` timeout period is long enough for network conditions.
- Verify that the switch is using the same UDP port number as the server.

NOTE: Because of an inconsistency between the Windows XP 802.1x supplicant timeout value and the switch default timeout value, which is 5, when adding a backup RADIUS server, set the switch `radius-server` timeout value to 4. Otherwise, the switch may not failover properly to the backup RADIUS server.

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch

Use `show radius` to verify that the encryption key the switch is using is correct for the server being contacted. If the switch has only a global key configured, it either must match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, it overrides the global key and must match the server key.

Example 158 Global and unique encryption keys

```
Switch(config)# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key : My-Global-Key 1
  Dynamic Authorization UDP Port : 3799

      Auth Acct DM/ Time
Server IP Addr  Port Port  CoA Window Encryption Key
-----
10.33.18.119   1812 1813          119-only-key 2
```

- 1 Global RADIUS Encryption Key 2 Unique RADIUS Encryption Key for the RADIUS server at 10.33.18.119

MSTP and fast-uplink problems

- ⚠ **CAUTION:** If you enable MSTP, HP recommends that you leave the remainder of the MSTP parameter settings at their default values until you have had an opportunity to evaluate MSTP performance in your network. Because incorrect MSTP settings can adversely affect network performance, you should avoid making changes without having a strong understanding of how MSTP operates. To learn the details of MSTP operation, see the IEEE802.1s standard.

Broadcast storms appearing in the network

This can occur when there are physical loops (redundant links) in the topology. Where this exists, you should enable MSTP on all bridging devices in the topology to detect the loop.

STP blocks a link in a VLAN even though there are no redundant links in that VLAN

In 802.1Q-compliant switches, MSTP blocks redundant physical links even if they are in separate VLANs. A solution is to use only one, multiple-VLAN (tagged) link between the devices. Also, if ports are available, you can improve the bandwidth in this situation by using a port trunk. See "Spanning Tree Operation with VLANs" in chapter "Static Virtual LANs (VLANs)" in the *Advanced Traffic Management Guide* for your switch.

Fast-uplink troubleshooting

Some of the problems that can result from incorrect use of fast-uplink MSTP include temporary loops and generation of duplicate packets.

Problem sources can include:

- Fast-uplink is configured on a switch that is the MSTP root device.
- Either the `Hello Time` or the `Max Age` setting (or both) is too long on one or more switches. Return the `Hello Time` and `Max Age` settings to their default values (2 seconds and 20 seconds, respectively, on a switch).
- A "downlink" port is connected to a switch that is further away (in hop count) from the root device than the switch port on which fast-uplink MSTP is configured.
- Two edge switches are directly linked to each other with a fast-uplink (`Mode = Uplink`) connection.
- Fast uplink is configured on both ends of a link.
- A switch serving as a backup MSTP root switch has ports configured for fast-uplink MSTP and has become the root device because of a failure in the original root device.

SSH-related problems

Switch access refused to a client

Even though you have placed the client's public key in a text file and copied the file (using the `copy tftp pub-key-file` command) into the switch, the switch refuses to allow the client to have access. If the source SSH client is an SSHv2 application, the public key may be in the PEM format, which the switch (SSHv1) does not interpret. Check the SSH client application for a utility that can convert the PEM-formatted key into an ASCII-formatted key.

Executing IP SSH does not enable SSH on the switch

The switch does not have a host key. Verify by executing `show ip host-public-key`. If you see the message

```
ssh cannot be enabled until a host key is configured (use 'crypto'
command) .
```

you need to generate an SSH key pair for the switch. To do so, execute `crypto key generate` (see "Generating the switch's public and private key pair" in the SSH chapter of the *Access Security Guide* for your switch.)

Switch does not detect a client's public key that does appear in the switch's public key file (`show ip client-public-key`)

The client's public key entry in the public key file may be preceded by another entry that does not terminate with a new line (CR). In this case, the switch interprets the next sequential key entry as simply a comment attached to the preceding key entry. Where a public key file has more than one entry, ensure that all entries terminate with a new line (CR). While this is optional for the last entry in the file, not adding a new line to the last entry creates an error potential if you either add another key to the file at a later time or change the order of the keys in the file.

An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages

```
Download failed: overlength key in key file.
Download failed: too many keys in key file.
Download failed: one or more keys is not a valid RSA public key.
```


The public key file you are trying to download has one of the following problems:

- A key in the file is too long. The maximum key length is 1024 characters, including spaces. This could also mean that two or more keys are merged together instead of being separated by a <CR> <LF>.
- There are more than ten public keys in the key file.
- One or more keys in the file is corrupted or is not a valid rsa public key.

Client ceases to respond ("hangs") during connection phase

The switch does not support data compression in an SSH session. Clients often have compression turned on by default, but then disable it during the negotiation phase. A client that does not recognize the compression-request FAILURE response may fail when attempting to connect. Ensure that compression is turned *off* before attempting a connection to prevent this problem.

TACACS-related problems

Event Log

When troubleshooting TACACS+ operation, check the switch's Event Log for indications of problem areas.

All users are locked out of access to the switch

If the switch is functioning properly, but no username/password pairs result in console or Telnet access to the switch, the problem may be caused by how the TACACS+ server and/or the switch are configured. Use one of the following methods to recover:

- Access the TACACS+ server application and adjust or remove the configuration parameters controlling access to the switch.
- If the above method does not work, try eliminating configuration changes in the switch that have not been saved to flash (boot-up configuration) by causing the switch to reboot from the boot-up configuration (which includes only the configuration changes made prior to the last `write memory` command.) If you did not use `write memory` to save the authentication configuration to flash, pressing the `Reset` button reboots the switch with the boot-up configuration.
- Disconnect the switch from network access to any TACACS+ servers and then log in to the switch using either Telnet or direct console port access. Because the switch cannot access a TACACS+ server, it defaults to local authentication. You can then use the switch's local Operator or Manager username/password pair to log on.
- As a last resort, use the `Clear/Reset` button combination to reset the switch to its factory default boot-up configuration. Taking this step means you will have to reconfigure the switch to return it to operation in your network.

No communication between the switch and the TACACS+ server application

If the switch can access the server device (that is, it can ping the server), a configuration error may be the problem. Some possibilities include:

- The server IP address configured with the switch's `tacacs-server host` command may not be correct. (Use the switch's `show tacacs-server` command to list the TACACS+ server IP address.)
- The encryption key configured in the server does not match the encryption key configured in the switch (by using the `tacacs-server key` command). Verify the key in the server and compare it to the key configured in the switch. (Use `show tacacs-server` to list the global key. Use `show config` or `show config running` to list any server-specific keys.)
- The accessible TACACS+ servers are not configured to provide service to the switch.

Access is denied even though the username/password pair is correct

Some reasons for denial include the following parameters controlled by your TACACS+ server application:

- The account has expired.
- The access attempt is through a port that is not allowed for the account.
- The time quota for the account has been exhausted.
- The time credit for the account has expired.
- The access attempt is outside of the time frame allowed for the account.
- The allowed number of concurrent logins for the account has been exceeded.

For more help, see the documentation provided with your TACACS+ server application.

Unknown users allowed to login to the switch

Your TACACS+ application may be configured to allow access to unknown users by assigning them the privileges included in a *default user* profile. See the documentation provided with your TACACS+ server application.

System allows fewer login attempts than specified in the switch configuration

Your TACACS+ server application may be configured to allow fewer login attempts than you have configured in the switch with the `aaa authentication num-attempts` command.

TimeP, SNTP, or Gateway problems

The switch cannot find the time server or the configured gateway

TimeP, SNTP, and Gateway access are through the primary VLAN, which in the default configuration is the DEFAULT_VLAN. If the primary VLAN has been moved to another VLAN, it may be disabled or does not have ports assigned to it.

VLAN-related problems

Monitor port

When using the monitor port in a multiple-VLAN environment, the switch handles broadcast, multicast, and unicast traffic output from the monitor port as follows:

- If the monitor port is configured for tagged VLAN operation on the same VLAN as the traffic from monitored ports, the traffic output from the monitor port carries the same VLAN tag.
- If the monitor port is configured for untagged VLAN operation on the same VLAN as the traffic from the monitored ports, the traffic output from the monitor port is untagged.
- If the monitor port is not a member of the same VLAN as the traffic from the monitored ports, traffic from the monitored ports does not go out the monitor port.

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

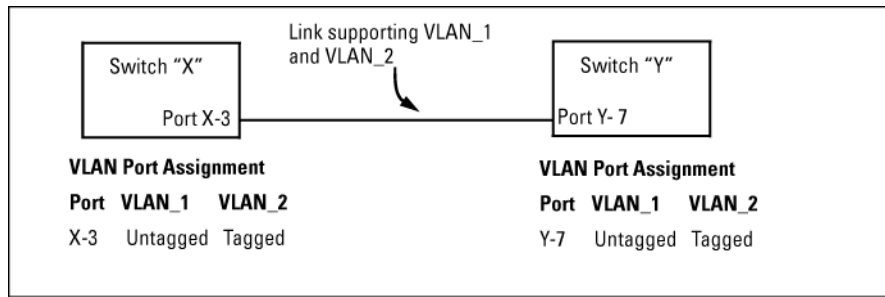
If multiple VLANs are being used on ports connecting 802.1Q-compliant devices, inconsistent VLAN IDs may have been assigned to one or more VLANs. For a given VLAN, the same VLAN ID must be used on all connected 802.1Q-compliant devices.

Link configured for multiple VLANs does not support traffic for one or more VLANs

One or more VLANs may not be properly configured as "Tagged" or "Untagged." A VLAN assigned to a port connecting two 802.1Q-compliant devices must be configured the same on both ports.

For example, VLAN_1 and VLAN_2 use the same link between switch "X" and switch "Y," as shown in Figure 55 (page 290).

Figure 55 Example: of correct VLAN port assignments on a link



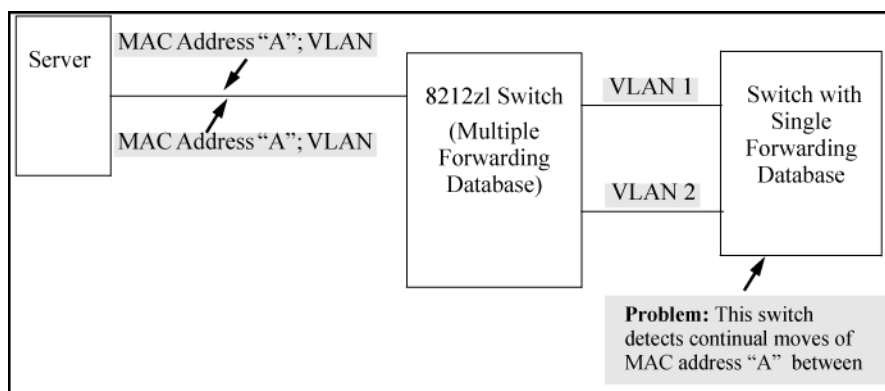
- If VLAN_1 (VID=1) is configured as "Untagged" on port 3 on switch "X," it must also be configured as "Untagged" on port 7 on switch "Y." Make sure that the VLAN ID (VID) is the same on both switches.
- Similarly, if VLAN_2 (VID=2) is configured as "Tagged" on the link port on switch "A," it must also 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 56 Example: of duplicate MAC address



Disabled overlapping subnet configuration

Previous software versions allowed configuration of VLAN IP addresses in overlapping subnets which can cause incorrect routing of packets and result in IP communication failure. As of software version WB.15.09, overlapping subnet configurations are no longer allowed. An overlapping

subnet is determined by the configuration order. The subnet that is configured first is valid, but any subsequent IP addresses that overlap are not allowed.

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

- The event log provides an error message in the format:
`ip: VLANx : IP initialization failed for vlan x.`
For a multinetted VLAN (multiple IP addresses assigned to the VLAN), only the IP addresses that are overlapping subnets are removed. The other IP addresses on the VLAN are retained and function correctly. The error message can be somewhat misleading; the IP addresses on the VLAN that are not overlapping are initialized correctly.
- The output of the `show ip` command correctly indicates that the overlapping IP address does not exist on the VLANs that have error messages in the event log.
- The output of the `show running-config` command incorrectly indicates that the overlapping IP address is configured. In [Example 159 “An IP address that is not actually configured on the VLAN”](#), the IP address shown in VLAN6 is not actually configured on the VLAN; it has been removed.

Example 159 An IP address that is not actually configured on the VLAN

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

The information is retained in the config file to allow you to boot up the switch and have it function as it did when it was configured with earlier software that allows overlapping subnets.

If you attempt to remove the overlapping subnet from the VLAN, the switch displays an error message similar to:

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

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

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

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

Fan failure

Whenever a fan failure occurs, the Fan/Fault LEDs blink amber and a log entry is recorded. During a fan failure, all operational fans are automatically set to the maximum operating speed until the fan failure has been resolved. At that time, the fan speed is reset to the minimum operating speed.

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 link-flap 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.
```

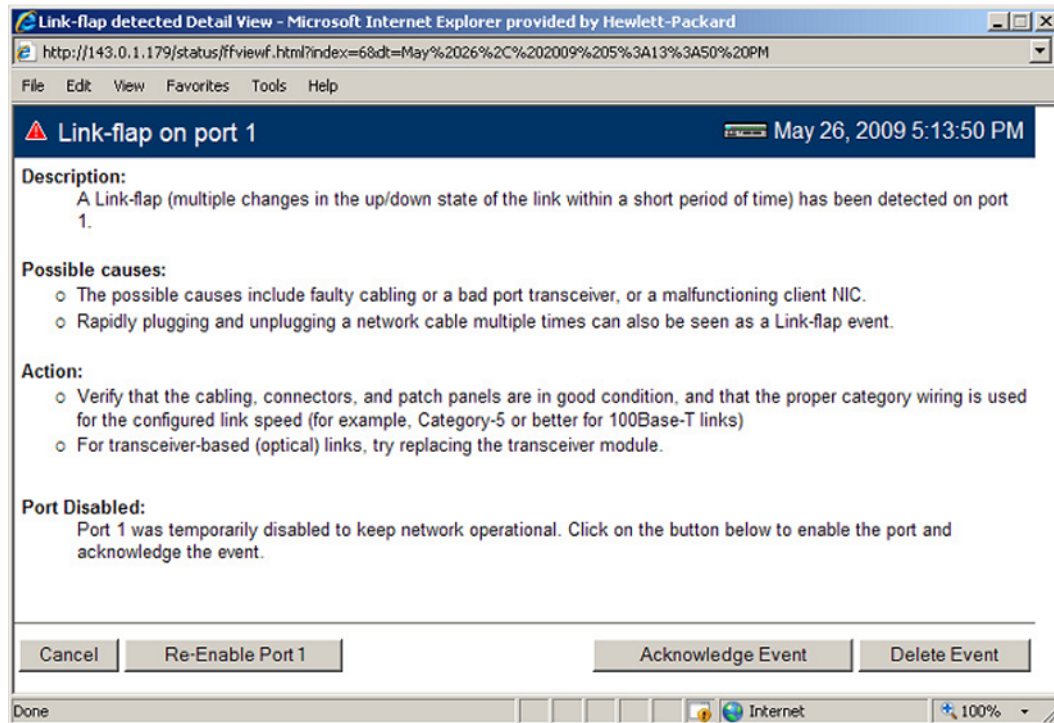
```
02674 FFI: port <number>-Administrator action required to re-enable.
```

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 57 \(page 293\)](#). 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 57 Link-flap on port 1 event detail dialog box



Fault finder thresholds

HP switches feature automatic fault detection, which helps protect against network loops and defective equipment. The fault detection sensitivity setting determines the types of alerts reported to the Alert Log based on their level of severity or sensitivity. The sensitivity levels are:

- **High Sensitivity.** This policy directs the switch to send all alerts to the Alert Log. This setting is most effective on networks that have none or few problems.
- **Medium Sensitivity.** This policy directs the switch to send alerts related to network problems to the Alert Log. If you want to be notified of problems which cause a noticeable slowdown on the network, use this setting.
- **Low Sensitivity.** This policy directs the switch to send only the most severe alerts to the Alert Log. This policy is most effective on a network where there are normally a lot of problems and you want to be informed of only the most severe ones
- **Disabled.** Disables the Alert Log and transmission of alerts (traps) to the management server (in cases where a network management tool such as ProCurve Manager is in use). Use this option when you don't want to use the Alert Log.

Enabling fault finder using the CLI

Enter this CLI command to enable fault detection:

Syntax:

```
[no] fault-finder [fault] [sensitivity <low|medium|high>] [action <warn|warn-and-disable>]
```

Enables or disables Fault Finder and sets sensitivity.

When the warn-and-disable action option is configured, Fault Finder may also shut down a bad port in addition to sending an alert to the Alert Log.

Default setting: `fault-finder sensitivity medium action warn`

[fault]: Supported values are:

- `all`: All fault types
- `bad-driver`: Too many undersized/giant packets
- `bad-transceiver`: Excessive jabbering
- `bad-cable`: Excessive CRC/alignment errors
- `too-long-cable`: Excessive late collisions
- `over-bandwidth`: High collision or drop rate
- `broadcast-storm`: Excessive broadcasts
- `duplex-mismatch-HDx`: Duplex mismatch. Reconfigure to Full Duplex
- `duplex-mismatch-FDx`: Duplex mismatch. Reconfigure port to Auto
- `link-flap`: Rapid detection of link faults and recoveries
- `loss-of-link`: Link loss detected. (Sensitivity not applicable)

Examples:

To set Fault Finder with a high sensitivity to issue a warning and then disable a port on which there is a high collision or drop rate, you could configure these options:

```
HPswitch(config)# fault-finder over-bandwidth sensitivity high action warn-and-disable
```

To set Fault Finder with a medium sensitivity to issue a warning about excessive CRC or alignment errors on a port, you could configure these options:

```
HPswitch(config)# fault-finder bad-cable sensitivity medium action warn
```

To set Fault Finder with a low sensitivity to issue a warning about rapid detection of link faults, you could configure these options:

```
HPswitch(config)# fault-finder link-flap sensitivity low action warn
```

To disable Fault Finder, enter this command:

```
HPswitch(config)# no fault-finder all
```

Table 26 Fault finder sensitivities for supported conditions

| Condition triggering fault finder | Sensitivities | | | Units (in packets) | Time period | Fault finder reacts: |
|--|---------------|--------|-----|--------------------|-------------|---|
| | High | Medium | Low | | | |
| Bad driver — Too many under-sized packets or too | 6 | 21 | 36 | 1/10,000 Incoming | 20 secs | If (undersized/total) >= (sensitivity/10,000) |

Table 26 Fault finder sensitivities for supported conditions *(continued)*

| Condition triggering fault finder | Sensitivities | | | Units (in packets) | Time period | Fault finder reacts: |
|---|---------------|-----------|-----------|------------------------------------|--------------------|---|
| many giant packets | | | | | | Or If (giant/total) >= (sensitivity/10,000) |
| Bad transceiver — Excessive jabbering -Jabbers: (Jabbers are packets longer than the MTU) -Fragments: (packets shorter than they should be) | 6 5 | 21 10 | 36 14 | 1/10,000 Incoming One Fragments | 20 secs 20 secs | If (jabbers/total) >= (sensitivity/10,000) Or If fragment count in the last 20 seconds >= sensitivity |
| Bad cable — Excessive CRC/alignment errors | 6 | 21 | 36 | 1/10,000 Incoming | 20 secs | If (CRC and alignment errors/ total) >= (sensitivity/10,000) |
| Too Long Cable — Excessive late collisions (a late collision error occurs after the first 512 bit times) | 6 | 21 | 36 | 1/10,000 Outgoing | 20 secs | If (late collisions/total) >= (sensitivity/10,000) |
| Over bandwidth -High collision rate -High drop rate | 6 65 | 21 257 | 36 449 | 1/10,000 Outgoing One Packet | 5 mins 5 mins | If (excessive collisions/total) >= (sensitivity/10,000) The count of dropped packets >= sensitivity during the last 5 minutes. |
| Broadcast storm — Excessive broadcasts | 2750 | 9200 | 15600 | One Broadcast Packet | 1 sec | If the average per second of broadcast packets in the last 20 seconds >= sensitivity |
| Duplex mismatch HDx | 6 | 21 | 36 | 1/10,000 Outgoing | 20 sec | If (late collisions/total) >= (sensitivity/10,000) |
| Duplex mismatch FDx | 6 | 21 | 36 | 1/10,000 Incoming | 20 sec | If (CRC and alignment errors/ total) >= (sensitivity/10,000) |
| Link flap — Excessive | 4 | 7 | 11 | One Transitions | 10 secs | If the Transition count in the last |

Table 26 Fault finder sensitivities for supported conditions *(continued)*

| Condition triggering fault finder | Sensitivities | | | Units (in packets) | Time period | Fault finder reacts: |
|---|---------------|--|--|--------------------|-------------|----------------------|
| transitions between link-up and link-down states. | | | | | | 10s >= sensitivity. |

Example: of sensitivity calculation:

If a sensitivity is set to High, and a bad cable is causing 15 CRC errors out of a total of 3500 packets transmitted in a 20 second period:

1. CRC errors/total must be >= (sensitivity/10,000) to trigger an alert.
2. CRC errors/total = 15/3500 = .00043
3. Sensitivity/10,000 = 6/10,000 = .0006
4. .00043 is not greater than or equal to .0006, so an alert is not triggered.

Viewing transceiver information

This feature provides the ability to view diagnostic monitoring information for transceivers with Diagnostic Optical Monitoring (DOM) support. The following table indicates the support level for specific transceivers:

| Product # | Description | Support* |
|-----------|------------------------------|----------|
| J8436A | 10GbE X2-SC SR Optic | V |
| J8437A | 10GbE X2-SC LR Optic | V |
| J8440B | 10GbE X2-CX4 Xcver | NA |
| J8440C | 10GbE X2-CX4 Xcver | NA |
| J4858A | Gigabit-SX-LC Mini-GBIC | V |
| J4858B | Gigabit-SX-LC Mini-GBIC | V |
| J4858C | Gigabit-SX-LC Mini-GBIC | V (some) |
| J9054B | 100-FX SFP-LC Transceiver | N |
| J8177C | Gigabit 1000Base-T Mini-GBIC | NA |
| J9150A | 10GbE SFP+ SR Transceiver | D |
| J9151A | 10GbE SFP+ LR Transceiver | D |
| J9152A | 10GbE SFP+ LRM Transceiver | D |
| J9153A | 10GbE SFP+ ER Transceiver | D |
| J9144A | 10GbE X2-SC LRM Transceiver | D |
| J8438A | 10GbE X2-SC ER Transceiver | D |

* Support indicators:

- V - Validated to respond to DOM requests
- N - No support of DOM
- D - Documented by the component suppliers as supporting DOM
- NA - Not applicable to the transceiver (copper transceiver)

NOTE: Not all transceivers support Digital Optical Monitoring. If DOM appears in the Diagnostic Support field of the `show interfaces transceiver detail` command, or the `hpicfTransceiverMIB hpicfXcvrDiagnostics` MIB object, DOM is supported for that transceiver.

Viewing information about transceivers (CLI)

Syntax:

```
show interfaces transceiver [port-list] [detail]
```

Displays information about the transceivers. If a port is specified, displays information for the transceiver in that port.

| | |
|----------|--|
| [detail] | Displays detailed transceiver information. |
|----------|--|

MIB support

The `hpicfTransceiver` MIB is available for displaying transceiver information.

Viewing transceiver information

The transceiver information displayed depends on the `show` command executed.

The output for `show interfaces transceiver [port-list]` is shown below. You can specify multiple ports, separated by commas, and the information for each transceiver will display.

Example 160 Output for a specified transceiver

```
HP Switch(config)# show interfaces transceiver 21
```

```
Transceiver Technical information:
```

| Port | Type | Product Number | Serial Number | Part Number |
|------|--------|----------------|---------------|-------------|
| 21 | 1000SX | J4858C | MY050VM9WB | 1990-3657 |

If there is no transceiver in the port specified in the command, the output displays as shown below.

Example 161 Output when no transceiver is present in specified interface

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

```
No Transceiver found on interface 22
```

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

Example 162 Output when no ports are specified

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

```
Transceiver Technical information:
```

| Port | Type | Product Number | Serial Number | Part Number |
|------|--------|----------------|---------------|-------------|
| 21 | 1000SX | J4858C | MY050VM9WB | 1990-3657 |
| 22 | 1000SX | J4858B | P834DIP2 | |

You can specify all for port-list as shown below.

Example 163 Output when "all" is specified

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

```
No Transceiver found on interface 1
```

```
No Transceiver found on interface 2
```

```
.  
. .
```

```
No Transceiver found on interface 24
```

```
Transceiver Technical information:
```

| Port | Type | Product Number | Serial Number | Part Number |
|------|--------|----------------|---------------|-------------|
| 21 | 1000SX | J4858C | MY050VM9WB | 1990-3657 |
| 22 | 1000SX | J4858B | P834DIP2 | |

Information displayed with the detail parameter

When the show interfaces transceiver [port-list] detail command is executed, the following information displays.

Table 27 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 28 \(page 299\)](#), [Table 29 \(page 299\)](#), and [Table 30 \(page 299\)](#) is only displayed when the transceiver supports DOM.

Table 28 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 29 \(page 299\)](#).

Table 29 Alarm and error information (GBIC/SFP transceivers only)

| Alarm | Description |
|-------------------|-----------------------------------|
| RX loss of signal | Incoming (RX) signal is lost |
| RX power high | Incoming (RX) power level is high |
| RX power low | Incoming (RX) power level is low |
| TX fault | Transmit (TX) fault |
| TX bias high | TX bias current is high |
| TX bias low | TX bias current is low |
| TX power high | TX power is high |
| TX power low | TX power is low |
| Temp high | Temperature is high |
| Temp low | Temperature is low |
| Voltage High | Voltage is high |
| Voltage Low | Voltage is low |

The alarm information for XENPAK transceivers is shown in [Table 30 \(page 299\)](#).

Table 30 Alarm and error information (XENPAK transceivers)

| Alarm | Description |
|------------------------------|---|
| WIS local fault | WAN Interface Sublayer local fault |
| Receive optical power fault | Receive optical power fault |
| PMA/PMD receiver local fault | Physical Medium Attachment/Physical Medium Dependent receiver local fault |
| PCS receiver local fault | Physical Coding Sublayer receiver local fault |
| PHY XS receive local fault | PHY Extended Sublayer receive local fault |
| RX power high | RX power is high |
| RX power low | RX power is low |
| Laser bias current fault | Laser bias current fault |
| Laser temperature fault | Laser temperature fault |
| Laser output power fault | Laser output power fault |

Table 30 Alarm and error information (XENPAK transceivers) (continued)

| Alarm | Description |
|---------------------------------|---------------------------------|
| TX fault | TX fault |
| PMA/PMD transmitter local fault | PMA/PMD transmitter local fault |
| PCS Transmit local fault | PCS transmit local fault |
| PHY XS transmit local fault | PHY SX transmit local fault |
| TX bias high | TX bias current is high |
| TX bias low | TX bias current is low |
| TX power high | TX power is high |
| TX power low | TX power is low |
| Temp high | Temperature is high |
| Temp low | Temperature is low |

An Example: of the output for the show interfaces transceiver [port-list] detail for a 1000SX transceiver is shown below.

Example 164 Detailed information for a 1000SX Mini-GBIC transceiver

```
HP Switch(config)# show interfaces transceiver 21 detail
```

```
Transceiver in 21
Interface index      : 21
Type                : 1000SX
Model               : J4858C
Connector type      : LC
Wavelength          : 850nm
Transfer distance   : 300m (50um), 150m (62.5um),
Diagnostic support  : DOM
Serial number       : MY050VM9WB
```

```
Status
Temperature : 50.111C
Voltage      : 3.1234V
TX Bias      : 6mA
TX Power     : 0.2650mW, -5.768dBm
RX Power     : 0.3892mW, -4.098dBm
```

```
Time stamp      : Mon Mar 7 14:22:13 2011
```

An Example: of the output for a 10GbE-LR transceiver is shown below.

Example 165 Detailed information for a 10GbE-LR transceiver

```
HP Switch(config)# show interfaces transceiver 23 detail

Transceiver in 23
Interface Index   : 24
Type              : 10GbE-LR
Model             : J8437A
Connector type    : SC
Wavelength        : Channel #0: 1310nm, #1:0nm, #2:0nm, #3:0nm
Transfer distance : 10000m (SM)
Diagnostic support: DOM
Serial number     : ED456SS987

Status
Temperature      : 32.754C
TX Bias          : 42.700mA
TX Power         : 0.5192mW, -2.847dBm
RX Power         : 0.0040mW, -23.979dBm

Recent Alarms:

Rx power low alarm
Rx power low warning

Recent errors:
Receive optical power fault
PMA/PMD receiver local fault
PMA/PMD transmitter local fault
PCS receive local fault
PHY XS transmit local fault

Time stamp : Mon Mar 7 16:26:06 2013
```

Viewing transceiver information for copper transceivers with VCT support

This feature provides the ability to view diagnostic monitoring information for copper transceivers with Virtual Cable Test (VCT) support. The cable quality of the copper cables connected between transceivers can be ascertained using the transceiver cable diagnostics. Results of the diagnostics are displayed with the appropriate CLI show commands and with SNMP using the `hpicfTransceiver` MIB.

The J8177C 1000Base-T Mini-GBIC is supported.

Testing the Cable

Enter the `test cable-diagnostics` command in any context to begin cable diagnostics for the transceiver. The diagnostic attempts to identify cable faults. The tests may take a few seconds to complete for each interface. There is the potential of link loss during the diagnostic.

Syntax:

```
test cable-diagnostics [port-list]
```

Invokes cable diagnostics and displays the results.

Example 166 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 167 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 31 General transceiver information

| Parameter | Description |
|-------------------|--|
| Interface Index | The switch interface number |
| Transceiver-type | Pluggable transceiver type |
| Transceiver model | Pluggable transceiver model |
| Connector-type | Type of connector of the transceiver |
| Wavelength | For an optical transceiver: the central wavelength of the laser sent, in nm. If the transceiver supports multiple wavelengths, the values will be separated by a comma. An electrical transceiver value is displayed as N/A. |

Table 31 General transceiver information *(continued)*

| Parameter | Description |
|--------------------|--|
| Transfer Distance | Link-length supported by the transceiver in meters. The corresponding transfer medium is shown in brackets following the transfer distance value, For example, 50um multimode fiber. If the transceiver supports multiple transfer media, the values are separated by a comma. |
| Diagnostic Support | Shows whether the transceiver supports diagnostics: None Supported DOM Supported VCT Supported |
| Serial Number | Serial number of the transceiver |
| Link Status | Link up or down |
| Speed | Speed of transceiver in Mbps |
| Duplex | Type of duplexing |
| Cable Status | Values are OK, Open, Short, or Impedance |
| Distance to Fault | The distance in meters to a cable fault (accuracy is +/- 2 meters); displays 0 (zero) if there is no fault |
| Pair Skew | Difference in propagation between the fastest and slowest wire pairs |
| Pair Polarity | Signals on a wire pair are polarized, with one wire carrying the positive signal and one carrying the negative signal. |
| MDI Mode | The MDI crossover status of the two wire pairs (1&2, 3&6, 4&5, 7&8), will be either MDI or MDIX |

Using the Event Log for troubleshooting switch problems

The Event Log records operating events in single- or double-line entries and serves as a tool to isolate and troubleshoot problems.

Once the log has received 2000 entries, it discards the oldest message each time a new message is received. The Event Log window contains 14 log entry lines. You can scroll through it to view any part of the log.

Once the log has received 2000 entries, it discards the oldest message each time a new message is received. The Event Log window contains 14 log-entry lines. You can scroll through it to view any part of the log.


NOTE: The Event Log is *erased* if power to the switch is interrupted or if you enter the `boot` system command. The contents of the Event Log are *not* erased if you:

- Reboot the switch by choosing the **Reboot Switch** option from the menu interface.
- Enter the `reload` command from the CLI.

Event Log entries

As shown in [Figure 58 \(page 304\)](#), each Event Log entry is composed of six or seven fields, depending on whether numbering is turned on or not:

Figure 58 Format of an event log entry



| Item | Description |
|----------------------|---|
| Severity | One of the following codes (from highest to lowest severity): M —(major) indicates that a fatal switch error has occurred. E —(error) indicates that an error condition occurred on the switch. W —(warning) indicates that a switch service has behaved unexpectedly. I —(information) provides information on normal switch operation. D —(debug) is reserved for HP internal diagnostic information. |
| Date | The date in the format <i>mm/dd/yy</i> when an entry is recorded in the log. |
| Time | The time in the format <i>hh:mm:ss</i> when an entry is recorded in the log. |
| Event number | The number assigned to an event. You can turn event numbering on and off with the [no] <code>log-number</code> command. |
| System module | The internal module (such as "ports:" for port manager) that generated a log entry. If VLANs are configured, a VLAN name also appears for an event that is specific to an individual VLAN. |
| Event message | A brief description of the operating event. |

Table 32 Event Log system modules

| System module | Description | Documented in HP Switch hardware/software guide |
|---------------|---|---|
| 802.1x | 802.1X authentication: Provides access control on a per-client or per-port basis: <ul style="list-style-type: none"> Client-level security that allows LAN access to 802.1X clients (up to 32 per port) with valid user credentials Port-level security that allows LAN access only on ports on which a single 802.1X-capable client (supplicant) has entered valid RADIUS user credentials | <i>Access Security Guide</i> |
| acl | ACLs: Filter layer-3 IP traffic to or from a host to block unwanted IP traffic and block or limit other protocol traffic such as TCP, UDP, IGMP, and ICMP. ACEs specify the filter criteria and an action (permit or deny) to take on a packet if it meets the criteria. | <i>Advanced Traffic Management Guide</i> |
| addrmgr | Address Table Manager: Manages MAC addresses that the switch has learned and are stored in the switch's address table. | <i>Management and Configuration Guide</i> |
| arp-protect | Dynamic ARP Protection: Protects the network from ARP cache poisoning. Only valid ARP requests and responses are relayed or used to update the local ARP cache. ARP | <i>Access Security Guide</i> |

Table 32 Event Log system modules *(continued)*

| System module | Description | Documented in HP Switch hardware/software guide |
|---------------|---|--|
| | packets with invalid IP-to-MAC address bindings advertised in the source protocol address and source physical address fields are discarded. | |
| auth | Authorization: A connected client must receive authorization through web, AMC, RADIUS-based, TACACS+-based, or 802.1X authentication before it can send traffic to the switch. | <i>Access Security Guide</i> |
| cdp | Cisco Discovery Protocol: Supports reading CDP packets received from neighbor devices, enabling a switch to learn about adjacent CDP devices. HP does not support the transmission of CDP packets to neighbor devices. | <i>Management and Configuration Guide</i> |
| chassis | Hardware operation, including modules and ports, power supply, fans, transceivers, CPU interrupt errors, switch temperature, and so on. Chassis messages include events on Power Over Ethernet (POE) operation. | <i>Installations Guides</i> <i>Management and Configuration Guide</i> |
| connfilt | Connection-rate filtering: Used on the network edge to protect the network from attack by worm-like malicious code by detecting hosts that are generating IP traffic that exhibits this behavior and (optionally) either throttling or dropping all IP traffic from the offending hosts. Connection-rate filtering messages include events on virus throttling. Virus throttling uses connection-rate filtering to stop the propagation of malicious agents. | <i>Access Security Guide</i> |
| console | Console interface used to monitor switch and port status, reconfigure the switch, and read the event log through an in-band Telnet or out-of-band connection. | <i>Installation and Getting Started Guide</i> |
| cos | Class of Service (CoS): Provides priority handling of packets traversing the switch, based on the IEEE 802.1p priority carried by each packet. CoS messages also include QoS events. The QoS feature classifies and prioritizes traffic throughout a network, establishing an end-to-end traffic priority policy to manage available bandwidth and improve throughput of important data. | <i>Advanced Traffic Management Guide</i> |
| dca | Dynamic Configuration Arbiter (DCA) determines the client-specific parameters that are assigned in an authentication session. | <i>Access Security Guide</i> |

Table 32 Event Log system modules *(continued)*

| System module | Description | Documented in HP Switch hardware/software guide |
|----------------------|---|--|
| dhcp | Dynamic Host Configuration Protocol (DHCP) server configuration: Switch is automatically configured from a DHCP (Bootp) server, including IP address, subnet mask, default gateway, Timep Server address, and TFTP server address. | <i>Management and Configuration Guide</i> |
| dhcp v6c | DHCP for IPv6 prefix assignment | <i>IPv6 Configuration Guide</i> |
| dhcpr | DHCP relay: Forwards client-originated DHCP packets to a DHCP network server. | <i>Advanced Traffic Management Guide</i> |
| download | Download operation for copying a software version or files to the switch. | <i>Management and Configuration Guide</i> |
| dhcp-snoop | DHCP snooping: Protects your network from common DHCP attacks, such as address spoofing and repeated address requests. | <i>Access Security Guide</i> |
| dma | Direct Access Memory (DMA): Transmits and receives packets between the CPU and the switch. | — |
| fault | Fault Detection facility, including response policy and the sensitivity level at which a network problem should generate an alert. | <i>Management and Configuration Guide</i> |
| fdr-log | FDR collects information that is “interesting” at the time of the crash, as well as when the switch is misbehaving, 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. | <i>Management and Configuration Guide</i> |
| ffi | Find, Fix, and Inform: Event or alert log messages indicating a possible topology loop that causes excessive network activity and results in the network running slow. FFI messages include events on transceiver connections with other network devices. | <i>Installation and Getting Started Guide</i> <i>Management and Configuration Guide</i> |
| garp | Generic Attribute Registration Protocol (GARP), defined in the IEEE 802.1D-1998 standard. | <i>Advanced Traffic Management Guide</i> |
| gvrp | GARP VLAN Registration Protocol (GVRP): Manages dynamic 802.1Q VLAN operations, in which the switch creates temporary VLAN membership on a port to provide a link to another port in the same VLAN on another device. | <i>Advanced Traffic Management Guide</i> |
| hpesp | Management module that maintains communication between switch ports. | <i>Installation and Getting Started Guide</i> |

Table 32 Event Log system modules *(continued)*

| System module | Description | Documented in HP Switch hardware/software guide |
|---------------|--|---|
| 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</i> <i>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> |
| kms | Key Management System: Configures and maintains security information (keys) for all routing protocols, including a timing mechanism for activating and deactivating an individual protocol. | <i>Access Security Guide</i> |
| lACP | LACP trunks: The switch can either automatically establish an 802.3ad-compliant trunk group or provide a manually configured, static LACP trunk. | <i>Management and Configuration Guide</i> |
| ldbal | Load balancing in LACP port trunks or 802.1s Multiple Spanning Tree protocol (MSTP) that uses VLANs in a network to improve network resource utilization and maintain a loop-free environment. Load-balancing messages also include switch meshing events. The switch meshing feature provides redundant links, improved bandwidth use, and | <i>Management and Configuration Guide</i> <i>Advanced Traffic Management Guide</i> |

Table 32 Event Log system modules *(continued)*

| System module | Description | Documented in HP Switch hardware/software guide |
|---------------|---|---|
| | support for different port types and speeds. | |
| 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 | Web and MAC authentication: Port-based security employed on the network edge to protect private networks and the switch itself from unauthorized access using one of the following interfaces: <ul style="list-style-type: none"> • Web page login to authenticate users for access to the network • RADIUS server that uses a device's MAC address for authentication | <i>Access Security Guide</i> |
| maclock | MAC lockdown and MAC lockout <ul style="list-style-type: none"> • MAC lockdown prevents station movement and MAC address "hijacking" by requiring a MAC address to be used only on an assigned port on the switch. MAC Lockdown also restricts the client device to a specific VLAN. • MAC lockout blocks a specific MAC address so that the switch drops all traffic to or from the specified address. | <i>Access Security Guide</i> |
| mgr | HP PCM and PCM+: Windows-based network management solutions for managing and monitoring performance of HP switches. PCM messages also include events for configuration operations. | <i>Management and Configuration Guide</i> |
| mld | Multicast Listener Discovery (MLD): IPv6 protocol used by a router to discover the presence of multicast listeners. MLD can also optimize IPv6 multicast traffic flow with the snooping feature. | <i>Multicast and Routing Guide</i> |
| mtm | Multicast Traffic Manager (MTM): Controls and coordinates L3 multicast traffic for upper layer protocols. | <i>Multicast and Routing Guide</i> |

Table 32 Event Log system modules *(continued)*

| System module | Description | Documented in HP Switch hardware/software guide |
|---------------|--|--|
| netinet | Network Internet: Monitors the creation of a route or an Address Resolution Protocol (ARP) entry and sends a log message in case of failure. | <i>Advanced Traffic Management Guide</i> |
| pagp | Ports Aggregation Protocol (PAgP): Obsolete. Replaced by LACP (802.3ad). | — |
| ports | Port status and port configuration features, including mode (speed and duplex), flow control, broadcast limit, jumbo packets, and security settings. Port messages include events on POE operation and transceiver connections with other network devices. | <i>Installation and Getting Started Guide</i> <i>Management and Configuration Guide</i> <i>Access Security Guide</i> |
| radius | RADIUS (Remote Authentication Dial-In User Service) authentication and accounting: A network server is used to authenticate user-connection requests on the switch and collect accounting information to track network resource usage. | <i>Access Security Guide</i> |
| 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> |

Table 32 Event Log system modules *(continued)*

| System module | Description | Documented in HP Switch hardware/software guide |
|---------------|--|--|
| 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 switches in the same IP subnet (broadcast domain), resulting in a reduced number of IP addresses and simplified management of small workgroups for scaling your network to handle increased bandwidth demand. | <i>Advanced Traffic Management Guide</i> |
| stp | Multiple-instance spanning tree protocol/MSTP (802.1s): Ensures that only one active path exists between any two nodes in a group of VLANs in the network. MSTP operation is designed to avoid loops and broadcast storms of duplicate messages that can bring down the network. | <i>Advanced Traffic Management Guide</i> |
| system | Switch management, including system configuration, switch bootup, activation of boot ROM image, memory buffers, traffic and security filters. System messages also include events from management interfaces (menu, CLI, and HP PCM+) used to reconfigure the switch and monitor switch status and performance. | <i>Basic Operation 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>Basic Operation Guide</i> |
| tftp | Trivial File Transfer Protocol: Supports the download of files to the switch from a TFTP network server. | <i>Basic Operation Guide</i> |
| timep | Time Protocol: Synchronizes and ensures a uniform time among interoperating devices. | <i>Management and Configuration Guide</i> |

Table 32 Event Log system modules *(continued)*

| System module | Description | Documented in HP Switch hardware/software guide |
|---------------|--|---|
| 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> |
| update | Updates (TFTP or serial) to HP switch software and updates to running-config and start-up config files | <i>Basic Operation Guide</i> |
| usb | Auxiliary port that allows you to connect external devices to the switch. | <i>Installation and Getting Started Guide</i> |
| vlan | <p>Static 802.1Q VLAN operations, including port-and protocol-based configurations that group users by logical function instead of physical location</p> <ul style="list-style-type: none"> • A port-based VLAN creates a layer-2 broadcast domain comprising member ports that bridge IPv4 traffic among themselves. • A protocol-based VLAN creates a layer-3 broadcast domain for traffic of a particular routing protocol, and comprises member ports that bridge traffic of the specified protocol type among themselves. <p>VLAN messages include events from management interfaces (menu, CLI, and HP PCM+) used to reconfigure the switch and monitor switch status and performance.</p> | <i>Advanced Traffic Management Guide</i> |
| xmodem | Xmodem: Binary transfer feature that supports the download of software files from a PC or UNIX workstation. | <i>Basic Operation Guide</i> |

Using the Menu

To display the Event Log from the Main Menu, select `Event Log`. [Example 168 \(page 312\)](#) shows a sample event log display.

Example 168 An event log display

```
HP Switch 5406z1                               25-Oct-2013  18:02:52
=====CONSOLE - MANAGER MODE -
=====
M 10/25/13 16:30:02 sys: 'Operator cold reboot from CONSOLE session.'
I 10/25/13 17:42:51 00061 system: -----
-
I 10/25/13 17:42:51 00063 system: System went down : 10/25/13 16:30:02
I 10/25/13 17:42:51 00064 system: Operator cold reboot from CONSOLE session.
W 10/25/13 17:42:51 00374 chassis: WARNING: SSC is out of Date: Load 8.2 or
newer
I 10/25/13 17:42:51 00068 chassis: Slot D Inserted
I 10/25/13 17:42:51 00068 chassis: Slot E Inserted
I 10/25/13 17:42:51 00068 chassis: Slot F Inserted
I 10/25/13 17:42:51 00690 udpf: DHCP relay agent feature enabled
I 10/25/13 17:42:51 00433 ssh: Ssh server enabled
I 10/25/13 17:42:51 00400 stack: Stack Protocol disabled
I 10/25/13 17:42:51 00128 tftp: Enable succeeded
I 10/25/13 17:42:51 00417 cdp: CDP enabled

----  Log events stored in memory 1-751. Log events on screen 690-704.

  Actions->   Back      Next page   Prev page   End      Help
```

Return to previous screen.

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

The *log status line* below the recorded entries states the total number of events stored in the event log and which logged events are currently displayed.

To scroll to other entries in the Event Log, either preceding or following the currently visible portion, press the keys indicated at the bottom of the display (Back, Nextpage, Prev page, or End) or the keys described in [Table 3-3 \(page 312\)](#).

Table 33 Event Log control keys

| Key | Action |
|-----|---|
| [N] | Advances the display by one page (next page). |
| [P] | Rolls back the display by one page (previous page). |
| [v] | Advances display by one event (down one line). |
| [^] | Rolls back display by one event (up one line). |
| [E] | Advances to the end of the log. |
| [H] | Displays Help for the Event Log. |

Using the CLI

Syntax:

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


By default, the `show logging` command displays the log messages recorded since the last reboot in chronological order:

| | |
|--------------|--|
| -a | Displays all recorded log messages, including those before the last reboot. |
| -b | Displays log events as the time since the last reboot instead of in a date/time format. |
| -r | Displays all recorded log messages, with the most recent entries listed first (reverse order). |
| -s | Displays the active management module (AM) and standby management module (SM) log events. |
| -t | Displays the log events with a granularity of 10 milliseconds. |
| -m | Displays only major log events. |
| -p | Displays only performance log events. |
| -w | Displays only warning log events. |
| -i | Displays only informational log events. |
| -d | Displays only debug log events. |
| <option-str> | Displays all Event Log entries that contain the specified text. Use an <option-str> value with -a or -r to further filter show logging command output. |

Example:

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

```
HP Switch# show logging -a system
```

To display all Event Log messages recorded since the last reboot that have the word "system" in the message text or module name, enter:

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

Clearing Event Log entries

Syntax:

```
clear logging
```

Removes all entries from the event log display output.

Use the `clear logging` command to hide, but not erase, Event Log entries displayed in `show logging` command output. Only new entries generated after you enter the command will be displayed.

To redisplay all hidden entries, including Event Log entries recorded prior to the last reboot, enter the `show logging -a` command.

Turning event numbering on

Syntax:

```
[no] log-numbers
```

Turns event numbering on and off

Using log throttling to reduce duplicate Event Log and SNMP messages

A recurring event can generate a series of duplicate Event Log messages and SNMP traps in a relatively short time. As a result, the Event Log and any configured SNMP trap receivers may be

flooded with excessive, exactly identical messages. To help reduce this problem, the switch uses *log throttle periods* to regulate (throttle) duplicate messages for recurring events, and maintains a counter to record how many times it detects duplicates of a particular event since the last system reboot.

When the first instance of a particular event or condition generates a message, the switch initiates a log throttle period that applies to all recurrences of that event. If the logged event recurs during the log throttle period, the switch increments the counter initiated by the first instance of the event, but does not generate a new message.

If the logged event repeats again after the log throttle period expires, the switch generates a duplicate of the first message, increments the counter, and starts a new log throttle period during which any additional instances of the event are counted, but not logged. Thus, for a particular recurring event, the switch displays only one message in the Event Log for each log throttle period in which the event reoccurs. Also, each logged instance of the event message includes counter data showing how many times the event has occurred since the last reboot. The switch manages messages to SNMP trap receivers in the same way.

Log throttle periods

The length of the log throttle period differs according to an event's severity level:

| Severity level | Log throttle period |
|-----------------|---------------------|
| I (Information) | 6000 Seconds |
| W (Warning) | 600 Seconds |
| D (Debug) | 60 Seconds |
| M (Major) | 6 Seconds |

Example:

Suppose that you configure VLAN 100 on the switch to support PIM operation, but do not configure an IP address. If PIM attempts to use VLAN 100, the switch generates the first instance of the following Event Log message and counter.

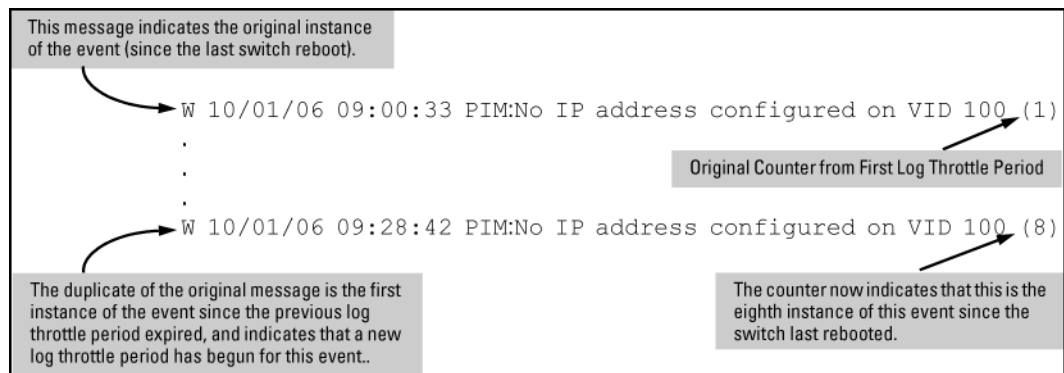
NOTE: In [Example 169 “The first instance of an event message and counter”](#) the counter (1) indicates that this is the first instance of this event since the switch last rebooted.

Example 169 The first instance of an event message and counter

```
W 10/01/12 09:00:33 PIM:No IP address configured on VID 100 (1)
```


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

Figure 59 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 60 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 \(page 316\)](#). (The same operation applies for messages sent to any configured SNMP trap receivers.)

Table 34 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* |
|--|--|--|----------------------------|
| 3 | | | 1 |
| | 5 | | 4 |
| | | 4 | 9 |

* This value always comprises the first instance of the duplicate message in the current log throttle period plus all previous occurrences of the duplicate message occurring since the switch last rebooted.

Reporting information about changes to the running configuration

Syslog can be used for sending notifications to a remote syslog server about changes made to the running configuration. The notifications in the syslog messages are sent in ASCII format and contain this information:

- Notice-Type: Describes the syslog notification as a “running config change”.
- Event-ID: Identifier for the running config change event that occurred on the switch.
- Config-Method: The source for the running config change.
- Device-Name: The managed device.
- User-Name: User who made the running config change.
- Remote-IP-Address: IP address of a remote host from which the user is connected.

Syntax:

```
[no] logging notify <running-config-change>  
[transmission-interval <0-4294967295>
```

Enables sending the running configuration change notifications to the syslog server. The `no` form of the command disables sending the running configuration changes to the syslog server.

Default: Disabled

| | |
|---|--|
| <code><running-config-change></code> | Mandatory option for the notify parameter. Specifies the type of notification to send. |
| <code>transmission-interval <0-4294967295></code> | Specifies the time interval (in seconds) between the transmission of two consecutive notifications. Running config changes occurring within the specified interval will not generate syslog notifications. |

A value of zero means there is no limit; a notification is sent for every running config change.

Default: Zero

Example 170 Sending running config changes to the syslog server

```
HP Switch(config)# logging notify running-config-change  
transmission-interval 10
```

Debug/syslog operation

While the Event Log records switch-level progress, status, and warning messages on the switch, the debug/system logging (*syslog*) feature provides a way to record Event Log and debug messages on a remote device. For example, you can send messages about routing misconfigurations and

other network protocol details to an external device, and later use them to debug network-level problems.

Debug/syslog messaging

The debug/syslog feature allows you to specify the types of Event Log and debug messages that you want to send to an external device. You can perform the following operations:

- Use the `debug` command to configure messaging reports for the following event types:

- ACL "deny" matches
- Dynamic ARP protection events
- DHCP snooping events
- DIPLD events
- Events recorded in the switch's Event Log
- IP routing events (IPv4 and IPv6)
- LACP events
- LLDP events
- SNMP events
- SSH events

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

- Severity level
- System module

Hostname in syslog messages

The syslog now messages the sender identified by hostname.

The hostname field identifies the switch that originally sends the syslog message. Configurable through the CLI and SNMP, the format of the hostname field supports the following formats:

- `ip-address`: The IP address of the sending interface will be used as the message origin identifier. This is the default format for the origin identifier. The IP address of the sending interface (in dotted decimal notation) is the default format.
- `hostname`: The hostname of the sending switch will be used as the message origin identifier.
- `none`: No origin identifier will be embedded in the syslog message. Nil value is used as defined by "-".

This configuration is system-wide, not per syslog server.

NOTE: There is no support in this feature for menu interface, WebUI or a fully qualified domain name. There are no changes in this feature to PCM or IDM. There are no new log events added in this feature.

Logging origin-id

Use the `logging origin-id` command to specify the content for the hostname field.

Syntax:

```
logging origin-id [ip-address|hostname|none]
[no]logging origin-id [ip-address|hostname|none]
```

To reset the hostname field content back to default (IP-address), use the `no` form of the command.

| | |
|----------------|--|
| filter | Creates a filter to restrict which events are logged. |
| IP-ADDR | Adds an IPv4 address to the list of receiving syslog servers. |
| IPV6-ADDR | Adds an IPv6 address to the list of receiving syslog servers. |
| origin-id | Sends the Syslog messages with the specified origin-id. |
| notify | Notifies the specified type sent to the syslog server(s). |
| priority-descr | A text string associated with the values of facility, severity, and system-module. |
| severity | Event messages of the specified severity or higher sent to the syslog server. |
| system-module | Event messages of the specified system module (subsystem) sent to the syslog server. |
| hostname | Sets the hostname of the device as the origin-id. |
| none | Disables origin-id in the syslog message. |

Add an IP address to the list of receiving syslog servers.

Use of `no` without an IP address specified will remove all IP addresses from the list of syslog receivers. If an IP address is specified, that receiver will be removed. Both link-local with zone ID and global IPv6 addresses are supported.

- Specify syslog server facility with the option `<facility>`. The command `no logging <facility>` sets the facility back to defaults.
- Specify filtering rules.
- Specify severity for event messages to be filtered to the syslog server with the option `<severity>`. The command `no logging <severity>` sets the severity back to default.
- Event messages of specified system module will be sent to the syslog server. Using `no` sends messages from all system modules. Messages are first filtered by selected severity.
- Specify syslog server transport layer with options `[udp] | [tcp] | [tls]`.
- Specify syslog server port number with options `[udp PORT-NUM] | [tcp PORT-NUM] | [tls PORT-NUM]`.
- Specify notification types to be sent to the syslog server.
- Use the option `transmission-interval` to control the egress rate limit for transmitting notifications, 0 value means there is no rate limit. The values are in seconds. Only one syslog message is allowed for transmission within specified time interval.
- Specify the origin information for the syslog messages with the option `origin-id`.

NOTE: When the syslog server receives messages from the switch, the IPv6 address of the switch is partly displayed.

Example:

Configured Host Ipv6 Address: 2001::1

Expected Syslog message:

```
Syslog message: USER.INFO: Oct 11 02:40:02 2001::1 00025 ip:
ST1CMDR: VLAN60: ip address 30.1.1.1/24 configured on vlan 60
```

Actual Truncated syslog message:

```
Syslog message: USER.INFO: Oct 11 02:40:02 2001:: 00025 ip: ST1CMDR:
VLAN60: ip address 30.1.1.1/24 configured on vlan 60
```

Use the command in [Example 171](#) to set the origin-id to the hostname.

Example 171 Setting the origin-id to the hostname

```
HP Switch(config)# logging origin-id hostname
```

The following syslog message will occur:

```
<14> Jan 1 00:15:35 HP-2910a1-24G 00076 ports: port 2 is now on-line
```

Use the command in [Example 172](#) to set the origin-id to none (nilvalue).

Example 172 Setting the origin-id to none (nilvalue)

```
HP Switch(config)# logging origin-id none
```

The following syslog message will occur:

```
<14> Jan 1 00:15:35 - 00076 ports: port 2 is now on-line
```

Use any of the commands in [Example 173](#) to set the origin-id to ip-address (default).

Example 173 Setting the origin-id to ip-address (default)

```
HP Switch(config)# logging origin-id ip-address
```

```
HP Switch(config)# no logging origin-id hostname
```

```
HP Switch(config)# no logging origin-id none
```

The following syslog message will occur:

```
<14> Jan 1 00:15:35 169.254.230.236 00076 ports: port 2 is now
on-line
```

Viewing the identification of the syslog message sender

Use the commands `show debug` or `show running-config` to display the identification of the syslog message sender. The default option for `origin-id` is `ip-address`. The command `show running-config` will not display the configured option when `origin-id` is set to the default value of `ip address`.

When `hostname` or `none` is configured using `logging origin-id`, the same displays as part of the `show running-config` command.

Syntax:

```
show debug
Default option is ip-address.
```

[Example 174](#) shows the output of the `show debug` command when configured without `login origin-id`.

Example 174 Output of the show debug command when configured without login origin-id

```
Debug Logging
  Origin identifier: Outgoing Interface IP
  Destination:      None

Enabled debug types:
  None are enabled.
```

The command `logging origin-id hostname` will produce the syslog message shown in [Example 175](#).

Example 175 Syslog message for logging origin-id hostname

```
Debug Logging
  Origin identifier: Hostname
  Destination:      None

Enabled debug types:
  None are enabled.
```

The command `logging origin-id none` will produce the syslog message shown in [Example 176](#).

Example 176 Syslog message for logging origin-id none

```
Debug Logging
  Origin identifier: none
  Destination:      None

Enabled debug types:
  None are enabled.
```

Syntax:

```
show running-config
```

[Example 177](#) shows the output of the `show running-config` command.

Example 177 Output of the show running-config command

```
The command logging origin-id hostname will display the
following:
logging origin-id hostname
```

The command `logging origin-id none` will display as the following:

```
logging origin-id none
```

SNMP MIB

SNMP support will be provided through the following MIB objects.

HpicfSyslogOriginId = textual-convention

| | |
|-----------------|--|
| Description | This textual convention enumerates the origin identifier of syslog message. |
| Syntax: integer | <ul style="list-style-type: none"> • ip-address • hostname • none |
| Status | <ul style="list-style-type: none"> • current |

hpicfSyslogOriginId OBJECT-TYPE

| | |
|-------------|--|
| Description | Specifies the content of a Hostname field in the header of a syslog message. |
| Syntax: | <ul style="list-style-type: none"> • HpicfSyslogOriginId |
| Max-access | <ul style="list-style-type: none"> • read-write |
| Status | <ul style="list-style-type: none"> • current |
| Default | <ul style="list-style-type: none"> • ip-address |

Debug/syslog destination devices

To use debug/syslog messaging, you must configure an external device as the logging destination by using the `logging` and `debug destination` commands. For more information, see [“Debug destinations” \(page 329\)](#) and [“Configuring a syslog server” \(page 330\)](#).

A debug/syslog destination device can be a syslog server and/or a console session. You can configure debug and logging messages to be sent to:

- Up to six syslog servers
- A CLI session through a direct RS-232 console connection, or a Telnet or SSH session

Debug/syslog configuration commands

| | | |
|----------------------------|--|--|
| Event notification logging | — | Automatically sends switch-level event messages to the switch's Event Log. Debug and syslog do not affect this operation, but add the capability of directing Event Log messaging to an external device. |
| logging command | <code><syslog-ip-addr></code> | Enables syslog messaging to be sent to the specified IP address. IPv4 and IPv6 are supported. |
| | <code>facility</code> | (Optional) The logging facility command specifies the destination (facility) subsystem used on a syslog server for debug reports. |
| | <code>priority-desc</code> | A text string associated with the values of facility, severity, and system-module. |
| | <code>neighbor-adjacency [detail]</code> | Enables or disables OSPFv3 (IPv6) adjacency logging. Must be executed in OSPFv3 context. The <code>detail</code> option displays all the adjacency state transitions and adjacency-related errors. |
| | <code>severity</code> | Sends Event Log messages of equal or greater severity than the specified value to configured debug destinations. (The default setting is to send Event Log messages from all severity levels.) |

| | | |
|---------------|---------------|---|
| | system-module | <p>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.</p> <p>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.</p> |
| 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 | <p>logging: Disables or re-enables syslog logging on one or more syslog servers configured with the logging syslog-ip-addr command.</p> <p>session: Assigns or re-assigns destination status to the terminal device that was most recently used to request debug output.</p> <p>buffer: Enables syslog logging to send the debug message types specified by the debug <debug-type> command to a buffer in switch memory.</p> |
| | 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</p> |

| | | |
|--|----------|--|
| | | <p>run</p> <p>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> |
| | lACP | <p>event: Sends messages related to change events.</p> <p>packet: Sends messages when BPDUs are exchanged.</p> |
| | lldp | Sends LLDP debug messages to the debug destinations. |
| | security | Sends security messages to the debug destination. |
| | services | Displays debug messages on the services module. |
| | snmp | Sends snmp messages to the debug destination. |

Using the Debug/Syslog feature, you can perform the following operations:

- Configure the switch to send Event Log messages to one or more Syslog servers. In addition, you can configure the messages to be sent to the User log facility (default) or to another log facility on configured Syslog servers.
- Configure the switch to send Event Log messages to the current management- access session (serial-connect CLI, Telnet CLI, or SSH).

- Disable all Syslog debug logging while retaining the Syslog addresses from the switch configuration. This allows you to configure Syslog messaging and then disable and re-enable it as needed.
- Display the current debug configuration. If Syslog logging is currently active, the list of configured Syslog servers is displayed.
- Display the current Syslog server list when Syslog logging is disabled.

Configuring debug/syslog operation

1. To use a syslog server as the destination device for debug messaging, follow these steps:
 - a. Enter the `logging <syslog-ip-addr>` command at the global configuration level to configure the syslog server IP address and enable syslog logging. Optionally, you may also specify the destination subsystem to be used on the syslog server by entering the `logging facility` command.

If no other syslog server IP addresses are configured, entering the `logging` command enables both debug messaging to a syslog server and the event debug message type. As a result, the switch automatically sends Event Log messages to the syslog server, regardless of other debug types that may be configured.
 - b. Re-enter the `logging` command in step "a (page 324)" 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 324)" 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 32 \(page 304\)](#).

5. If you configure system-module, severity-level values, or both to filter Event Log messages, when you finish troubleshooting, you may want to reset these values to their default settings so that the switch sends all Event Log messages to configured debug destinations (syslog servers, CLI session, or both).

To remove a configured setting and restore the default values that send all Event Log messages, enter one or both of the following commands:

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

△ CAUTION: If you configure a severity-level, system-module, logging destination, or logging facility value and save the settings to the startup configuration (For example, by entering the `write memory` command), the debug settings are saved after a system reboot (power cycle or reboot) and re-activated on the switch. As a result, after switch startup, one of the following situations may occur:

- Only a partial set of Event Log messages may be sent to configured debug destinations.
 - Messages may be sent to a previously configured syslog server used in an earlier debugging session.
-

Viewing a debug/syslog configuration

Use the `show debug` command to display the currently configured settings for:

- Debug message types and Event Log message filters (severity level and system module) sent to debug destinations
- Debug destinations (syslog servers or CLI session) and syslog server facility to be used

Syntax:

```
show debug
```

Displays the currently configured debug logging destinations and message types selected for debugging purposes. (If no syslog server address is configured with the `logging <syslog-ip-addr>` command, no `show debug` command output is displayed.)

Example 178 Output of the show debug command

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

```
Debug Logging
Destination:
Logging --
 10.28.38.164
Facility=kern
Severity=warning
System module=all-pass
Enabled debug types:
event
```

Example:

In the following Example:, no syslog servers are configured on the switch (default setting). When you configure a syslog server, debug logging is enabled to send Event Log messages to the server. To limit the Event Log messages sent to the syslog server, specify a set of messages by entering the `logging severity` and `logging system-module` commands.

Figure 61 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 61 \(page 326\)](#), 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 62 \(page 327\)](#).

Figure 62 Debug/syslog configuration for multiple debug types and multiple destinations

```

HP Switch# config
HP Switch(config)# logging 10.38.64.164
HP Switch(config)# show debug
Debug Logging
Destination:
Logging --
  10.38.64.164
  Facility=user
  Severity=debug
  System module=all-pass
Enabled debug types:
event
HP Switch(config)# no debug event
HP Switch(config)# debug acl
HP Switch(config)# debug ip ospf packet
HP Switch(config)# debug destination session
HP Switch(config)# show debug
Debug Logging
Destination:
Logging --
  10.38.64.164
  Facility=user
  Severity=debug
  System module=all-pass
Session
Enabled debug types:
acl log
ip ospf packet
  
```

Configure a Syslog server IP address. (No other Syslog servers are configured on the switch.) The server address serves as an active debug destination for any configured debug types.)

Display the new debug configuration. (Default debug settings - facility, severity, system module, and debug types- are displayed.)

Remove the unwanted event message logging to debug destinations.

Configure the debug messages types that you want to send to the Syslog server and CLI session.

Configure the CLI session as a debug destination.

Display the final debug and Syslog server configuration.

Debug command

At the manager level, use the `debug` command to perform two main functions:

- Specify the types of event messages to be sent to an external destination.
- Specify the destinations to which selected message types are sent.

By default, no debug destination is enabled and only Event Log messages are enabled to be sent.

NOTE: To configure a syslog server, use the `logging <syslog-ip-addr>` command. For more information, see “Configuring a syslog server” (page 330).

Debug messages

Syntax:

```
[no] debug <debug-type>
```

| | |
|-----|--|
| acl | When a match occurs on an ACL "deny" ACE (with <code>log</code> configured), the switch sends an ACL message to configured debug destinations. |
|-----|--|

| | |
|--|--|
| | <p>For information on ACLs, see the "Access Control Lists (ACLs)" chapter in the latest version of the following guides:</p> <ul style="list-style-type: none"> • IPv4 ACLs: <i>Access Security Guide</i> • IPv6 ACLs: <i>IPv6 Configuration Guide</i> <p>NOTE: ACE matches (hits) for permit and deny entries can be tracked using the <code>show statistics <aclv4 aclv6></code> command. (Default: Disabled—ACL messages for traffic that matches "deny" entries are not sent.)</p> |
| all | <p>Configures the switch to send all debug message types to configured debug destinations.</p> <p>(Default: Disabled—No debug messages are sent.)</p> |
| cdp | <p>Sends CDP information to configured debug destinations.</p> |
| destination | <p>logging—Disables or re-enables syslog logging on one or more syslog servers configured with the <code>logging <syslog-ip-addr></code> command.</p> <p>session—Assigns or re-assigns destination status to the terminal device that was most recently used to request debug output.</p> <p>buffer—Enables syslog logging to send the debug message types specified by the <code>debug <debug-type></code> command to a buffer in switch memory.</p> <p>For more information on these options, see "Debug destinations" (page 329).</p> |
| event | <p>Configures the switch to send Event Log messages to configured debug destinations.</p> <p>NOTE: This value does not affect the reception of event notification messages in the Event Log on the switch.</p> <p>Event Log messages are automatically enabled to be sent to debug destinations in these conditions:</p> <ul style="list-style-type: none"> • If no syslog server address is configured and you enter the <code>logging <syslog-ip-addr></code> command to configure a destination address. • If at least one syslog server address is configured in the startup configuration, and the switch is rebooted or reset. <p>Event log messages are the default type of debug message sent to configured debug destinations.</p> |
| ip [fib forwarding packet rip] | <p>Sends IP messages to configured destinations.</p> |
| ip [fib[events]] | <p>For the configured debug destinations:</p> <p>events—Sends IP forwarding information base events.</p> |
| ip [packet] | <p>Enables the specified PIM message type.</p> |
| ip [rip[database event trigger]] | <p>rip <database event trigger> —Enables the specified RIP message type for the configured destination(s).</p> <p>database—Displays database changes.</p> <p>event—Displays RIP events.</p> <p>trigger—Displays trigger messages.</p> |
| ipv6 [dhcpv6-client nd packet] | <p>NOTE: See the "IPv6 Diagnostic and Troubleshooting" chapter in the <i>IPv6 Configuration Guide</i> for your switch for more detailed IPv6 debug options.</p> |

| | |
|--|--|
| | <p>When no debug options are included, displays debug messages for all IPv6 debug options.</p> <p><code>dhcpv6-client [events packet]</code> —Displays DHCPv6 client event and packet data.</p> <p><code>nd</code>—Displays debug messages for IPv6 neighbor discovery.</p> <p><code>packet</code>—Displays IPv6 packet messages.</p> |
| <code>lldp</code> | Enables all LLDP message types for the configured destinations. |
| <code>security [arp-protect dhcp-snooping dynamic-ip-lockdown port-access port-security radius-server ssh tacacs-server user-profile-mib]</code> | <p><code>arp-protect</code>— Sends dynamic ARP protection debug messages to configured debug destinations.</p> <p><code>dhcp-snooping</code>—Sends DHCP snooping debug messages to configured debug destinations.</p> <p><code>agent</code>—Displays DHCP snooping agent messages.</p> <p><code>event</code>—Displays DHCP snooping event messages.</p> <p><code>packet</code>—Displays DHCP snooping packet messages.</p> <p><code>dynamic-ip-lockdown</code>—Sends dynamic IP lockdown debug messages to the debug destination.</p> <p><code>port-access</code>—Sends port-access debug messages to the debug destination.</p> <p><code>radius-server</code>—Sends RADIUS debug messages to the debug destination.</p> <p><code>ssh</code>—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</code> <code><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 <pdu></code> | <p>Displays the SNMP debug messages.</p> <p><code>pdu</code>—Displays SNMP pdu debug messages.</p> |

Debug destinations

Use the `debug destination` command to enable (and disable) syslog messaging on a syslog server or to a CLI session for specified types of debug and Event Log messages.

Syntax:

```
[no] debug destination <logging | session | buffer>
```

| | |
|----------------------|--|
| <code>logging</code> | <p>Enables syslog logging to configured syslog servers so that the debug message types specified by the <code>debug <debug-type></code> command (see “Debug messages” (page 327)) are sent.</p> <p>(Default: Logging disabled)</p> <p>To configure a syslog server IP address, see “Configuring a syslog server” (page 330).</p> |
|----------------------|--|

| | |
|---------|--|
| | <p>NOTE: Debug messages from the switches covered in this guide have a debug severity level. Because the default configuration of some syslog servers ignores syslog messages with the debug severity level, ensure that the syslog servers you want to use to receive debug messages are configured to accept the debug level. For more information, see “Operating notes for debug and Syslog” (page 335).</p> |
| session | <p>Enables transmission of event notification messages to the CLI session that most recently executed this command. The session can be on any one terminal emulation device with serial, Telnet, or SSH access to the CLI at the Manager level prompt (HP Switch#_).</p> <p>If more than one terminal device has a console session with the CLI, you can redirect the destination from the current device to another device. Do so by executing <code>debug destination session</code> in the CLI on the terminal device on which you now want to display event messages.</p> <p>Event message types received on the selected CLI session are configured with the <code>debug <debug-type></code> command.</p> |
| buffer | <p>Enables syslog logging to send the debug message types specified by the <code>debug <debug-type></code> command to a buffer in switch memory.</p> <p>To view the debug messages stored in the switch buffer, enter the <code>show debug buffer</code> command.</p> |

Logging command

At the global configuration level, the `logging` command allows you to enable debug logging on specified syslog servers and select a subset of Event Log messages to send for debugging purposes according to:

- Severity level
- System module

By specifying both a severity level and system module, you can use both configured settings to filter the Event Log messages you want to use to troubleshoot switch or network error conditions.

CAUTION: After you configure a syslog server and a severity level and/or system module to filter the Event Log messages that are sent, if you save these settings to the startup configuration file by entering the `write memory` command, these debug and logging settings are automatically re-activated after a switch reboot or power recycle. The debug settings and destinations configured in your previous troubleshooting session will then be applied to the current session, which may not be desirable.

After a reboot, messages remain in the Event Log and are not deleted. However, after a power recycle, all Event Log messages are deleted.

If you configure a severity level, system module, or both to temporarily filter Event Log messages, be sure to reset the values to their default settings by entering the `no` form of the following commands to ensure that Event Log messages of all severity levels and from all system modules are sent to configured syslog servers:

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

Configuring a syslog server

Syslog is a client-server logging tool that allows a client switch to send event notification messages to a networked device operating with syslog server software. Messages sent to a syslog server can be stored to a file for later debugging analysis.

To use the syslog feature, you must install and configure a syslog server application on a networked host accessible to the switch. For instructions, see the documentation for the syslog server application.

To configure a syslog service, use the `logging <syslog-ip-addr>` command as shown below. When you configure a syslog server, Event Log messages are automatically enabled to be sent to the server. To reconfigure this setting, use the following commands:

- `debug`
Specifies additional debug message types (see [“Debug messages” \(page 327\)](#)).
- `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” \(page 334\)](#) and [“Configuring the system module used to select the Event Log messages sent to a syslog server” \(page 335\)](#).)

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” \(page 321\)](#)).

Syntax:

```
[no] logging <syslog-ip-addr>
```

Enables or disables syslog messaging to the specified IP address. You can configure up to six addresses. If you configure an address when none are already configured, this command enables destination logging (syslog) and the Event debug type.

Therefore, at a minimum, the switch begins sending Event Log messages to configured syslog servers. The ACL, IP-OSPF, and/or IP-RIP message types are also sent to the syslog servers if they are currently enabled as debug types. (See [“Debug messages” \(page 327\)](#).)

| | |
|---|---|
| <code>no logging</code> | Removes all currently configured syslog logging destinations from the running configuration. Using this form of the command to delete the only remaining syslog server address disables debug destination logging on the switch, but the default Event debug type does not change. |
| <code>no logging <syslog-ip-address></code> | Removes only the specified syslog logging destination from the running configuration. Removing all configured syslog destinations with the <code>no logging</code> command (or a specified syslog server destination with the <code>no logging <syslog-ip-address></code> command) does not delete the syslog server IP addresses stored in the startup configuration. |

Deleting syslog addresses in the startup configuration

Enter a `no logging` command followed by the `write memory` command.

Verifying the deletion of a syslog server address

Display the startup configuration by entering the `show config` command.

Blocking the messages sent to configured syslog servers from the currently configured debug message type

Enter the `no debug <debug-type>` command. (See [“Debug messages” \(page 327\)](#).)

Disabling syslog logging on the switch without deleting configured server addresses

Enter the `no debug destination logging` command. Note that, unlike the case in which no syslog servers are configured, if one or more syslog servers are already configured and syslog messaging is disabled, configuring a new server address does not re-enable syslog messaging. To re-enable syslog messaging, you must enter the `debug destination logging` command.

Sending logging messages using TCP

Syntax:

```
[no] logging <ip-addr> [ udp 1024-49151 | tcp 1024-49151 ]
```

Allows the configuration of the UDP or TCP transport protocol for the transmission of logging messages to a syslog server.

Specifying a destination port with UDP or TCP is optional.

Default ports: UDP port is 514

TCP port is 1470

Default Transport Protocol: UDP

Because TCP is a connection-oriented protocol, a connection must be present before the logging information is sent. This helps ensure that the logging message will reach the syslog server. Each configured syslog server needs its own connection. You can configure the destination port that is used for the transmission of the logging messages.

Example 179 Configuring TCP for logging message transmission using the default port

```
HP Switch(config)# logging 192.123.4.5 tcp  
(Default TCP port 1470 is used.)
```

Example 180 Configuring TCP for logging message transmission using a specified port

```
HP Switch(config)# logging 192.123.4.5 9514  
(TCP port 9514 is used.)
```

Example 181 Configuring UDP for logging message transmission using the default port

```
HP Switch(config)# logging 192.123.4.5 udp  
(Default UDP port 514 is used.)
```

Example 182 Configuring UDP for logging message transmission using a specified port

```
HP Switch(config)# logging 192.123.4.5 9512  
(UDP port 9512 is used.)
```

Syntax:

```
[no] logging facility <facility-name>
```

The logging facility specifies the destination subsystem used in a configured syslog server. (All configured syslog servers must use the same subsystem.) HP recommends

the default (user) subsystem unless your application specifically requires another subsystem. Options include:

| | |
|-------------------|---|
| user | (default) Random user-level messages |
| kern | Kernel messages |
| mail | Mail system |
| daemon | System daemons |
| auth | Security/authorization messages |
| syslog | Messages generated internally by syslog |
| lpr | Line-printer subsystem |
| news | Netnews subsystem |
| uucp | uucp subsystem |
| cron | cron/at subsystem |
| sys9 | cron/at subsystem |
| sys10 - sys14 | Reserved for system use |
| local10 - local17 | Reserved for system use |

Use the `no` form of the command to remove the configured facility and reconfigure the default (user) value.

Adding a description for a Syslog server

You can associate a user-friendly description with each of the IP addresses (IPv4 only) configured for syslog using the CLI or SNMP.

CAUTION: Entering the `no logging` command removes ALL the syslog server addresses without a verification prompt.

NOTE: The HP enterprise MIB `hpicfSyslog.mib` allows the configuration and monitoring of syslog for SNMP (RFC 3164 supported).

The CLI command is:

Syntax:

```
logging <ip-addr> [control-descr <text_string>]
no logging <ip-addr> [control-descr]
```

An optional user-friendly description that can be associated with a server IP address. If no description is entered, this is blank. If `<text_string>` contains white space, use quotes around the string. IPv4 addresses only.

Use the `no` form of the command to remove the description. Limit: 255 characters

NOTE: To remove the description using SNMP, set the description to an empty string.

Example 183 The logging command with a control description

```
HP Switch(config)# logging 10.10.10.2 control-descr syslog_one
```

Adding a priority description

This description can be added with the CLI or SNMP. The CLI command is:

Syntax:

```
logging priority-descr <text_string>  
no logging priority-descr
```

Provides a user-friendly description for the combined filter values of `severity` and `system module`. If no description is entered, this is blank.

If `text_string` contains white space, use quotes around the string.

Use the `no` form of the command to remove the description.

Limit: 255 characters

Example 184 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” \(page 330\)](#).

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 330).)

See [Table 32](#) (page 304) 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 disabled. |
| IP (debug type) | Disabled. |

- Debugcommands do not affect normal message output to the Event Log.
Using the `debug event` command, you can specify that Event Log messages are sent to the debug destinations you configure (CLI session, syslog servers, or both) in addition to the Event Log.

- Ensure that your syslog servers accept debug messages.
All syslog messages resulting from a debug operation have a "debug" severity level. If you configure the switch to send debug messages to a syslog server, ensure that the server's syslog application is configured to accept the "debug" severity level. (The default configuration for some syslog applications ignores the "debug" severity level.)
- Duplicate IP addresses are not stored in the list of syslog servers.
- If the default severity value is in effect, all messages that have severities greater than the default value are passed to syslog. For example, if the default severity is "debug," all messages that have severities greater than debug are passed to syslog.
- There is a limit of six syslog servers. All syslog servers are sent the same messages using the same filter parameters. An error is generated for an attempt to add more than six syslog servers.

Diagnostic tools

Port auto-negotiation

When a link LED does not light (indicating loss of link between two devices), the most common reason is a failure of port auto-negotiation between the connecting ports. If a link LED fails to light when you connect the switch to a port on another device, do the following:

1. Ensure that the switch port and the port on the attached end-node are both set to `Auto` mode.
2. If the attached end-node does not have an `Auto` mode setting, you must manually configure the switch port to the same setting as the end-node port. See [“Port Status and Configuration” \(page 44\)](#).

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” \(page 352\)](#).

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 63 \(page 337\)](#).

Figure 63 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. |

| | |
|---------------------|--|
| data-size <0-65471> | Size of packet sent. Default: 0 (zero) |
| data-fill <0-1024> | The data pattern in the packet. Default: Zero length string |
| ip-option | Specify an IP option, such as loose or strict source routing, or an include-timestamp option: 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 include-timestamp-and-address: Records the intermediate router's timestamp and IP address. Default: 4 include-timestamp-from: Records the timestamp of the specified router addresses. loose-source-route <IP-addr>: The loose-source-route 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. 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. 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 strict-source-route <IP-addr>: Restricts the ping packet to only those IP addresses that have been specified and no other addresses. |
| tos <0-255> | Specifies the type of service to be entered in the header packet. Default: 0 (zero) |

Example 185 Ping tests

```
HP Switch# ping 10.10.10.10
10.10.10.10 is alive, time = 15 ms

HP Switch# ping 10.10.10.10 repetitions 3
10.10.10.10 is alive, iteration 1, time = 15 ms
10.10.10.10 is alive, iteration 1, time = 15 ms
10.10.10.10 is alive, iteration 1, time = 15 ms

HP Switch# ping 10.10.10.10 timeout 2
10.10.10.10 is alive, time = 10 ms

HP Switch# ping 10.11.12.13
The destination address is unreachable.
```

Halting a ping test

To halt a ping test before it concludes, press **[Ctrl] [C]**.

NOTE: To use the ping (or traceroute) command with host names or fully qualified domain names, see [“DNS resolver” \(page 352\)](#).

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 64 Link tests

| | |
|---|---|
| Basic Link Test | HP Switch# link 0030c1-7fcc40 Link-test passed. |
| Link Test with Repetitions | HP Switch# link 0030c1-7fcc40 repetitions 3 802.2 TEST packets sent: 3, responses received: 3 |
| Link Test with Repetitions and Timeout | HP Switch# link 0030c1-7fcc40 repetitions 3 timeout 1 802.2 TEST packets sent: 3, responses received: 3 |
| Link Test Over a Specific VLAN | HP Switch# link 0030c1-7fcc40 repetitions 3 timeout 1 vlan 1 802.2 TEST packets sent: 3, responses received: 3 |
| Link Test Over a Specific VLAN; Test Fail | HP Switch# link 0030c1-7fcc40 repetitions 3 timeout 1 vlan 222 802.2 TEST packets sent: 3, responses received: 0 |

Tracing the route from the switch to a host address

The traceroute command enables you to trace the route from the switch to a host address.

This command outputs information for each (router) hop between the switch and the destination address. Note that every time you execute `traceroute`, it uses the same default settings unless you specify otherwise for that instance of the command.

Syntax:

```
traceroute <ip-address | hostname> [maxttl <1-255>][minttl <1-255>][probes <1-5>][source <ip-address | source-vlan <vid> | loopback <0-7> ] [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.

| | |
|---|--|
| <code><ip-address hostname></code> | <p>The IP address or hostname of the device to which to send the traceroute.</p> |
| <code>[minttl <1-255>]</code> | <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> |
| <code>[maxttl <1-255>]</code> | <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> |
| <code>[probes <1-5>]</code> | <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> |
| <code>[source <ip-addr vid loopback <0-7>>]</code> | <p>The source IPv4 address, VLAN ID, or Loopback address.</p> |
| <code>[dstport <1-34000>]</code> | <p>Destination port.</p> |
| <code>[srcport <1-34000>]</code> | <p>Source port.</p> |
| <code>[ip-option]</code> | <p>Specify an IP option, such as loose or strict source routing, or an include-timestamp option:</p> <p><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</p> <p><code>[include-timestamp-and-address]</code>: Records the intermediate router's timestamp and IP address. Default: 4</p> <p><code>[loose-source-route <IP-addr>]</code>: Prompts for the IP address of each source IP on the path.</p> <p>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><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.</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.</p> <p>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.</p> <p>For any instance of traceroute, if you want a timeout value other than the default, you must specify that value.</p> <p>(Default: 5 seconds)</p> |
|--|---|

NOTE: For information about traceroute6, see the *IPv6 Configuration Guide* for your switch.

Halting an ongoing traceroute search

Press the **[Ctrl] [C]** keys.

A low maxttl causes traceroute to halt before reaching the destination address

Executing traceroute with its default values for a destination IP address that is four hops away produces a result similar to this:

Figure 65 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 65 (page 341)), executing traceroute with an insufficient maxttl for the actual hop count produces an output similar to this:

Figure 66 Incomplete traceroute because of low maxttl setting

```

Traceroute does not reach destination IP address because of low maxttl setting.

Switch# traceroute 125.25.24.35 (maxttl 3)
traceroute to 125.25.24.35 ,
      1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2      0 ms      0 ms      0 ms
 2 10.71.217.2      0 ms      0 ms      0 ms
 3 10.243.170.1     0 ms *      0 ms

```

The asterisk indicates there was a timeout on the second probe to the third hop.

If a network condition prevents traceroute from reaching the destination

Common reasons for traceroute failing to reach a destination include:

- Timeouts (indicated by one asterisk per probe, per hop)
- Unreachable hosts
- Unreachable networks

- Interference from firewalls
- Hosts configured to avoid responding

Executing `tracert` 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 67 Traceroute failing to reach the destination address

```

HP Switch# traceroute 125.25.24.35 maxttl 7
traceroute to 107.64.197.100 ,
          1 hop min, 7 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2          0 ms          0 ms          0 ms
 2 10.71.217.2          0 ms          0 ms          0 ms
 3 * 10.243.170.1 ----- 0 ms *
 4 * * * *
 5 * * * *
 6 * * * *
 7 * * * *

```

At hop 3, the first and third probes timed out but the second probe reached the router.
 All further probes within the `maxttl` timed-out without finding a router or the destination IP address.

An asterisk indicates a timeout without finding the next hop.

Viewing switch configuration and operation

In some troubleshooting scenarios, you may need to view the switch configuration to diagnose a problem. The complete switch configuration is contained in a file that you can browse from the CLI using the commands described in this section.

Viewing the startup or running configuration file

Syntax:

```
write terminal
```

Displays the running configuration.

| | |
|----------------------------------|-------------------------------------|
| <code>show config</code> | Displays the startup configuration. |
| <code>show running-config</code> | Displays the running-config file. |

For more information and examples of how to use these commands, see “Switch Memory and Configuration” in the *Basic Operation Guide*.

Viewing the configuration file (WebAgent)

To display the running configuration using the WebAgent:

1. In the navigation pane, click **Troubleshooting**.
2. Click **Configuration Report**.
3. Use the right-side scroll bar to scroll through the configuration listing.

Viewing a summary of switch operational data

Syntax:

```
show tech
```

By default, the `show tech` command displays a single output of switch operating and running-configuration data from several internal switch sources, including:

- Image stamp (software version data)
- Running configuration

- Event Log listing
- Boot history
- Port settings
- Status and counters — port status
- IP routes
- Status and counters — VLAN information
- GVRP support
- Load balancing (trunk and LACP)

[Example 186](#) shows sample output from the `show tech` command.

Example 186 The show tech command

```
HP Switch# show tech

show system

Status and Counters - General System Information

System Name       : Switch
System Contact    :
System Location   :

MAC Age Time (sec) : 300

Time Zone         : 0
Daylight Time Rule : None

Software revision : XX.14.xx      Base MAC Addr  : 001871-c42f00

ROM Version       : XX.12.12      Serial Number  : SG641SU00L

Up Time          : 23 hours      Memory - Total :
CPU Util (%)     : 10             Free          :

IP Mgmt - Pkts Rx : 759          Packet - Total : 6750
                Pkts Tx : 2          Buffers Free  : 5086
                                   Lowest           : 4961
                                   Missed           : 0

show flash
Image      Size(Bytes)  Date  Version
-----  -

```

To specify the data displayed by the `show tech` command, use the `copy show tech` command as described in [“Customizing show tech command output”](#) (page 344).

Saving show tech command output to a text file


When you enter the `show tech` command, a summary of switch operational data is sent to your terminal emulator. You can use your terminal emulator's text capture features to save the `show tech` data to a text file for viewing, printing, or sending to an associate to diagnose a problem.

For example, if your terminal emulator is the Hyperterminal application available with Microsoft® Windows® software, you can copy the `show tech` output to a file and then use either Microsoft Word or Notepad to display the data. (In this case, Microsoft Word provides the data in an easier-to-read format.)

The following Example: uses the Microsoft Windows terminal emulator. If you are using a different terminal emulator application, see the documentation provided with the application.


1. In Hyperterminal, click on Transfer | Capture Text...(see [Figure 68 \(page 344\)](#)).

Figure 68 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 69 \(page 344\)](#).

Figure 69 Entering a path and filename for saving show tech output



3. Click **[Start]** to create and open the text file.
4. From the global configuration context, enter the show tech command:

```
HP Switch# show tech
```

The show tech command output is copied into the text file and displayed on the terminal emulator screen. When the command output stops and displays -- MORE --, press the Space bar to display and copy more information. The CLI prompt appears when the command output finishes.

5. Click on Transfer | Capture Text | Stop in HyperTerminal to stop copying data and save the text file.
If you do not stop HyperTerminal from copying command output into the text file, additional unwanted data can be copied from the HyperTerminal screen.
6. To access the file, open it in Microsoft Word, Notepad, or a similar text editor.

Customizing show tech command output

Use the copy show tech command to customize the detailed switch information displayed with the show tech command to suit your troubleshooting needs.

To customize the information displayed with the show tech command:

1. Determine the information that you want to gather to troubleshoot a problem in switch operation.
2. Enter the copy show tech command to specify the data files that contain the information you want to view.

Syntax:

```
copy <source> show-tech
```


Specifies the operational and configuration data from one or more source files to be displayed by the `show tech` command. Enter the command once for each data file that you want to include in the display.

Default: Displays data from all source files, where `<source>` can be any one of the following values:

| | |
|--|--|
| <pre>command-output "<command>"</pre> | <p>Includes the output of a specified command in <code>show-tech</code> command output.</p> <p>Enter the command name between double-quotation marks, For example, <code>copy "show system" show-tech</code>.</p> |
| <pre>crash-data [slot-id master]</pre> | <p>Includes the crash data from all management and interface modules in <code>show tech</code> command output.</p> <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 "File Transfers" (page 224).</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> |

| | |
|---|---|
| | For more information on using <code>copy usb</code> commands, see "File Transfers" (page 224) . |
| <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 show <code>tech</code> command output, where:</p> <p><code>startup-config</code>: Specifies the name of the startup configuration file on the connected device.</p> <p><code>config <filename></code>: Specifies the pathname of a configuration file on the connected device.</p> <p><code>command-file <acl-filename.txt></code>: Specifies the pathname of an ACL command file on the connected device.</p> <p><code>pc unix</code>: Specifies whether the connected device is a DOS-based PC or UNIX workstation.</p> <p>For more information on using <code>copy xmodem</code> commands, see "File Transfers" (page 224).</p> |

Viewing more information on switch operation

Use the following commands to display additional information on switch operation for troubleshooting purposes.

Syntax:

```
show boot-history
```

Displays the crash information saved for each management module on the switch.

```
show history
```

Displays the current command history. This command output is used for reference or when you want to repeat a command (See ["Displaying the information you need to diagnose problems" \(page 349\)](#)).

```
show system-information
```

Displays globally configured parameters and information on switch operation.

```
show version
```

Displays the software version currently running on the switch and the flash image from which the switch booted (primary or secondary). For more information, see ["Displaying Management Information" in the "Redundancy \(Switch 8212zl\)" chapter](#).

```
show interfaces
```

Displays information on the activity on all switch ports (see ["Viewing Port Status and Configuring Port Parameters" in the "Port Status and Configuration" chapter](#)).

```
show interfaces-display
```

Displays the same information as the `show interfaces` command and dynamically updates the output every three seconds. Press **Ctrl + C** to stop the dynamic updates of system information. Use the Arrow keys to view information that is off the screen.

Searching for text using pattern matching with show command

Selected portions of the output are displayed, depending on the parameters chosen.

Syntax:

```
show <command option> | <include | exclude | begin>  
<regular expression>
```

Uses matching pattern searches to display selected portions of the output from a `show` command. There is no limit to the number of characters that can be matched. Only regular expressions are permitted; symbols such as the asterisk cannot be substituted to perform more general matching.

| | |
|---------|--|
| include | Only the lines that contain the matching pattern are displayed in the output. |
| exclude | Only the lines that contain the matching pattern are <i>not</i> displayed in the output. |
| begin | The display of the output begins with the line that contains the matching pattern. |

NOTE: Pattern matching is case-sensitive.

Below are examples of what portions of the running config file display depending on the option chosen.

Example 187 Pattern matching with include option

```
HP Switch(config)# show run | include ipv6 1
  ipv6 enable
  ipv6 enable
ipv6 access-list "EH-01"
HP Switch(config)#
```

- 1** Displays only lines that contain "ipv6".
-

Example 188 Pattern matching with exclude option

```
HP Switch(config)# show run | exclude ipv6 1
```

Running configuration:

```
; J9299A Configuration Editor; Created on release #WB.15.XX

hostname "HP Switch"
snmp-server community "notpublic" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24,B1-B20
  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
```

- 1** Displays all lines that do not contain "ipv6".
-

Example 189 Pattern matching with begin option

```
HP Switch(config)# show run | begin ipv6 1
  ipv6 enable
  no untagged 21-24
  exit
vlan 20
  name "VLAN20"
  untagged 21-24
  ipv6 enable
  no ip address
  exit
policy qos "michael"
  exit
ipv6 access-list "EH-01"
  sequence 10 deny tcp 2001:db8:255::/48 2001:db8:125::/48
  exit
no autorun
password manager
```

- 1 Displays the running config beginning at the first line that contains "ipv6".
-

Example 190 (page 349) is an Example: of the `show arp` command output, and then the output displayed when the `include` option has the IP address of `15.255.128.1` as the regular expression.

Example 190 The show arp command and pattern matching with the include option

```
HP Switch(config)# show arp

IP ARP table

  IP Address          MAC Address          Type    Port
  -----
  15.255.128.1        00000c-07ac00        dynamic B1
  15.255.131.19       00a0c9-b1503d        dynamic
  15.255.133.150      000bcd-3cbeec        dynamic B1

HP Switch(config)# show arp | include 15.255.128.1
15.255.128.1 00000c-07ac00 dynamic B1
```

Displaying the information you need to diagnose problems

Use the following commands in a troubleshooting session to more accurately display the information you need to diagnose a problem.

Syntax:

```
alias
```

Creates a shortcut alias name for commonly used commands and command options.

Syntax:

```
kill
```

Terminates a currently running, remote troubleshooting session. Use the `show ip ssh` command to list the current management sessions.

Syntax:

```
[no] page
```

Toggles the paging mode for `show` commands between continuous listing and per-page listing.

Syntax:

```
repeat
```

Repeatedly executes one or more commands so that you can see the results of multiple commands displayed over a period of time. To halt the command execution, press any key on the keyboard.

Syntax:

```
setup
```

Displays the Switch Setup screen from the menu interface.

Restoring the factory-default configuration

As part of your troubleshooting process, it may become necessary to return the switch configuration to the factory default settings. This process:

- Momentarily interrupts the switch operation
- Clears any passwords
- Clears the console Event Log
- Resets the network counters to zero
- Performs a complete self test
- Reboots the switch into its factory default configuration, including deleting an IP address

There are two methods for resetting to the factory-default configuration:

- CLI
- Clear/Reset button combination

NOTE: HP recommends that you save your configuration to a TFTP server before resetting the switch to its factory-default configuration. You can also save your configuration via Xmodem to a directly connected PC.

Resetting to the factory-default configuration

Using the CLI

This command operates at any level *except* the Operator level.

Syntax:

```
erase startup-configuration
```

Deletes the startup-config file in flash so that the switch will reboot with its factory-default configuration.

NOTE: The `erase startup-config` command does not clear passwords unless `include-credentials` has been set, at which time this command does erase username/password information and any other credentials stored in the config file. For more information, see the section on "Saving Security Credentials in a Config File" in the *Access Security Guide* for your switch.

Using Clear/Reset

1. Using pointed objects, simultaneously press both the `Reset` and `Clear` buttons on the front of the switch.
2. Continue to press the `Clear` button while releasing the `Reset` button.
3. When the Self Test LED begins to flash, release the `Clear` button.
The switch then completes its self test and begins operating with the configuration restored to the factory default settings.

Restoring a flash image

The switch can lose its operating system if either the primary or secondary flash image location is empty or contains a corrupted OS file and an operator uses the `erase flash` command to erase a good OS image file from the opposite flash location.

Recovering from an empty or corrupted flash state

Use the switch's console serial port to connect to a workstation or laptop computer that has the following:

- A terminal emulator program with Xmodem capability, such as the HyperTerminal program included in Windows PC software.
- A copy of a good OS image file for the switch.

NOTE: The following procedure requires the use of Xmodem and copies an OS image into primary flash only.

This procedure assumes you are using HyperTerminal as your terminal emulator. If you use a different terminal emulator, you may need to adapt this procedure to the operation of your particular emulator.

-
1. Start the terminal emulator program.
Ensure that the terminal program is configured as follows:
 - Baud rate: 9600
 - No parity
 - 8 Bits
 - 1 stop bit
 - No flow control
 2. Use the `Reset` button to reset the switch.
The following prompt should then appear in the terminal emulator:

```
Enter h or ? for help.  
  
=>
```
 3. Because the OS file is large, you can increase the speed of the download by changing the switch console and terminal emulator baud rates to a high speed. For Example:
 - a. Change the switch baud rate to 115,200 Bps.

=> sp 115200

- b. Change the terminal emulator baud rate to match the switch speed:
 - i. In HyperTerminal, select **Call | Disconnect**.
 - ii. Select **File | Properties**.
 - iii. Click on **Configure**.
 - iv. Change the baud rate to **115200**.
 - v. Click on **[OK]**, then in the next window, click on **[OK]** again.
 - vi. Select **Call | Connect**.
 - vii. Press **[Enter]** one or more times to display the => prompt.
4. Start the Console Download utility by entering `do` at the => prompt and pressing **[Enter]**:
=> do

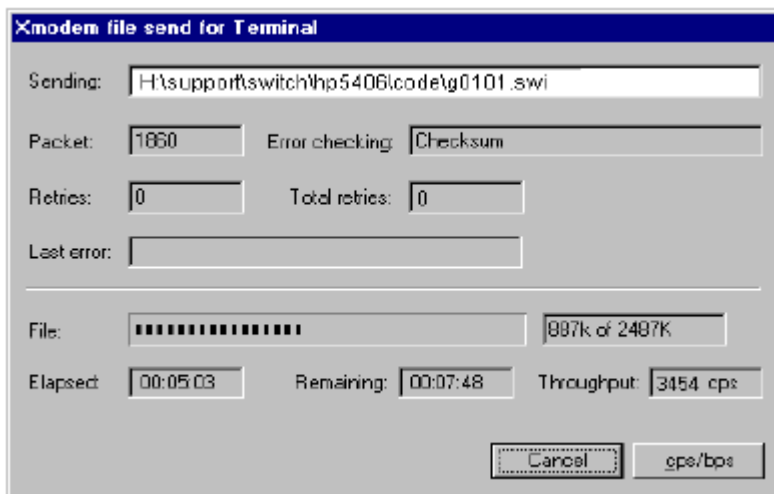
5. You then see this prompt:

```
You have invoked the console download utility.  
Do you wish to continue? (Y/N)>_
```

6. At the above prompt:
 - a. Enter **y** (for Yes)
 - b. Select **Transfer | File** in HyperTerminal.
 - c. Enter the appropriate filename and path for the OS image.
 - d. Select the **Xmodem** protocol (and not the 1k Xmodem protocol).
 - e. Click on **[Send]**.

If you are using HyperTerminal, you will see a screen similar to the following to indicate that the download is in progress:

Figure 70 Example: of Xmodem download in progress



When the download completes, the switch reboots from primary flash using the OS image you downloaded in the preceding steps, plus the most recent startup-config file.

DNS resolver

The domain name system (DNS) resolver is designed for use in local network domains, where it enables the use of a host name or fully qualified domain name with DNS-compatible switch CLI commands.

DNS operation supports both IPv4 and IPv6 DNS resolution and multiple, prioritized DNS servers. (For information on IPv6 DNS resolution, see the latest *IPv6 Configuration Guide* for your switch.)


Basic operation

- When the switch is configured with only the IP address of a DNS server available to the switch, a DNS-compatible command, executed with a fully qualified domain name, can reach a device found in any domain accessible through the configured DNS server.
- When the switch is configured with both of the following:
 - The IP address of a DNS server available to the switch
 - The domain suffix of a domain available to the configured DNS serverthen:
 - A DNS-compatible command that includes the host name of a device in the same domain as the configured domain suffix can reach that device.
 - A DNS-compatible command that includes a fully qualified domain name can reach a device in any domain that is available to the configured DNS server.

Example:

Suppose the switch is configured with the domain suffix [mygroup.HP Switch.net](#) and the IP address for an accessible DNS server. If an operator wants to use the switch to ping a target host in this domain by using the DNS name "leader" (assigned by a DNS server to an IP address used in that domain), the operator can use either of the following commands:

Figure 71 Example: of using either a host name or a fully qualified domain name



In the preceding Example:, if the DNS server's IP address is configured on the switch, but a domain suffix is either not configured or is configured for a different domain than the target host, the fully qualified domain name *must* be used.

Note that if the target host is in a domain *other than* the domain configured on the switch:

- The host's domain must be reachable from the switch. This requires that the DNS server for the switch must be able to communicate with the DNS servers in the path to the domain in which the target host operates.
- The fully qualified domain name must be used, and the domain suffix must correspond to the domain in which the target host operates, regardless of the domain suffix configured in the switch.

Example:

Suppose the switch is configured with the domain suffix [mygroup.HP Switch.net](#) and the IP address for an accessible DNS server in this same domain. This time, the operator wants to use the switch to trace the route to a host named "remote-01" in a different domain named [common.group.net](#). Assuming this second domain is accessible to the DNS server already configured on the switch, a `tracert` command using the target's fully qualified DNS name should succeed.

Figure 72 Example: using the fully qualified domain name for an accessible target in another domain

```

HP Switch# traceroute remote-01.common.group.net
[traceroute to 10.22.240.73]
1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.229.3          0 ms          0 ms          0 ms
 2 10.71.217.1         0 ms          0 ms          0 ms
 3 10.0.198.2          1 ms          0 ms          0 ms
[4 10.22.240.73       0 ms          0 ms          0 ms]

```

Fully Qualified Host Name for the Target Host

IP Address for Target Host "remote-01"

Configuring and using DNS resolution with DNS-compatible commands

The DNS-compatible commands include ping and traceroute.)

1. Determine the following:
 - The IP address for a DNS server operating in a domain in your network.
 - The priority (1 to 3) of the selected server, relative to other DNS servers in the domain.
 - The domain name for an accessible domain in which there are hosts you want to reach with a DNS-compatible command. (This is the domain suffix in the fully qualified domain name for a given host operating in the selected domain. See “Basic operation” (page 353).) 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 71 \(page 353\)](#).) 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 73 \(page 355\)](#), 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 73 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 |

| Entity | Identity |
|----------------------------|---------------|
| 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 `docserver` to reach the document server at 10.28.229.219.

Example 191 Configuring switch "A" in Figure 73 (page 355) to support DNS resolution

```
HP Switch(config)# ip dns server-address 10.28.229.10
HP Switch(config)# ip dns domain-name pbs.outdoors.com
```

Example 192 Ping and traceroute execution for the network in Figure 73 (page 355)

```
HP Switch(config)# ping docservr
10.28.229.219 is alive, time = 1 ms
```

```
HP Switch# traceroute docservr
traceroute to 10.28.229.219
      1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.192.2 1 1 ms 0 ms 0 ms
 2 10.28.229.219 2 0 ms 0 ms 0 ms
```

1 First-Hop Router ("B") **2** Traceroute Target

As mentioned under "Basic operation" (page 353), 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 73 (page 355) as a target:

Figure 74 Example: of ping and traceroute execution when only the DNS server IP address is configured

```
HP Switch# ping [docservr.pbs.outdoors.com]
10.28.229.219 is alive, time = 1 ms

HP Switch# traceroute [docservr.pbs.outdoors.com]
traceroute to 10.28.229.219
      1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.192.2 1 1 ms 0 ms 0 ms
 2 10.28.229.219 2 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 75 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
```

DNS Resolver Configuration in the show ip command output

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.

| | |
|----------------|--|
| blink <1-1440> | Blinks the chassis Locate LED for a specified number of minutes (Default: 30 minutes). |
| on <1-1440> | Turns the chassis Locate LED on for a specified number of minutes (Default: 30 minutes). |
| off | Turns the chassis Locate LED off. |

Example 193 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.

12 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 360)).

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.

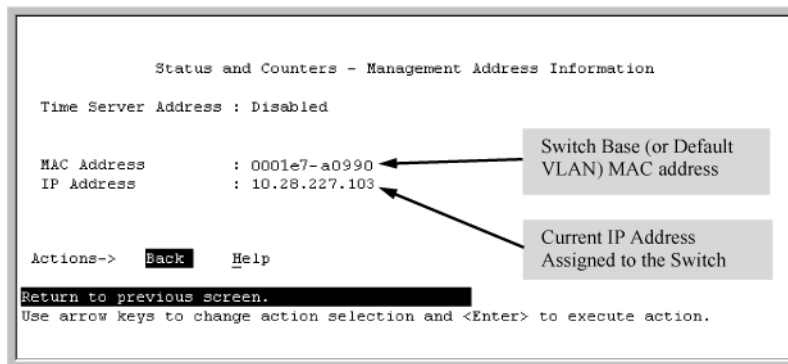
Also, the Base MAC address appears on a label on the back of the switch.

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 76 Example: of the Management Address Information screen



Viewing the port and VLAN MAC addresses

The MAC address assigned to each switch port is used internally by such features as Flow Control and the spanning-tree protocol. Using the `walkmib` command to determine the MAC address assignments for individual ports can sometimes be useful when diagnosing switch operation.

NOTE: This procedure displays the MAC addresses for all ports and existing VLANs in the switch, regardless of which VLAN you select.

- If the switch is at the CLI Operator level, use the `enable` command to enter the Manager level of the CLI.
- 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 77 \(page 361\)](#):

- 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 77 Example: of Port MAC address assignments on a switch

```

HP Switch# walkmib ifphysaddress
ifPhysAddress.1 = 00 12 79 88 b1 ff
ifPhysAddress.2 = 00 12 79 88 b1 fe
ifPhysAddress.3 = 00 12 79 88 b1 fd
ifPhysAddress.4 = 00 12 79 88 b1 fc
ifPhysAddress.49 = 00 12 79 88 b1 cf
ifPhysAddress.50 = 00 12 79 88 b1 ce
ifPhysAddress.51 = 00 12 79 88 b1 cd
ifPhysAddress.52 = 00 12 79 88 b1 cc
ifPhysAddress.53 = 00 12 79 88 b1 cb
ifPhysAddress.54 = 00 12 79 88 b1 ca
ifPhysAddress.55 = 00 12 79 88 b1 c9
ifPhysAddress.56 = 00 12 79 88 b1 c8
ifPhysAddress.57 = 00 12 79 88 b1 c7
ifPhysAddress.58 = 00 12 79 88 b1 c6
ifPhysAddress.59 = 00 12 79 88 b1 c5
ifPhysAddress.60 = 00 12 79 88 b1 c4
ifPhysAddress.61 = 00 12 79 88 b1 c3
ifPhysAddress.62 = 00 12 79 88 b1 c2
ifPhysAddress.63 = 00 12 79 88 b1 c1
ifPhysAddress.64 = 00 12 79 88 b1 c0
ifPhysAddress.65 = 00 12 79 88 b1 bf
ifPhysAddress.66 = 00 12 79 88 b1 be
ifPhysAddress.67 = 00 12 79 88 b1 bd
ifPhysAddress.68 = 00 12 79 88 b1 bc
ifPhysAddress.69 = 00 12 79 88 b1 bb
ifPhysAddress.70 = 00 12 79 88 b1 ba
ifPhysAddress.71 = 00 12 79 88 b1 b9
ifPhysAddress.72 = 00 12 79 88 b1 b8
ifPhysAddress.362 = 00 12 79 88 a1 00
ifPhysAddress.461 = 00 12 79 88 a1 00
ifPhysAddress.488 = 00 12 79 88 a1 00
ifPhysAddress.4456 =
  
```

ifPhysAddress.1 - 4: Ports A1 - A4 in Slot A
(Addresses 5 - 24 in slot A are unused.)

ifPhysAddress.49 - 72: Ports C1 - C24 in Slot C
(In this example, there is no module in slot B.)

ifPhysAddress.362 Base MAC Address (MAC Address for default VLAN; VID = 1)

ifPhysAddress.461 and 488 Physical addresses for non-default VLANs configured on the switch. On the switches covered by this manual, all VLANs use the same MAC address as the Default VLAN. Refer to "Multiple VLAN Considerations" in the "Static LANs (VLANs)" chapter of the *Advanced Traffic Management Guide* for your switch.

Virtual

ifPhysAddress.362 Base MAC Address (MAC Address for default VLAN; VID = 1)

ifPhysAddress.461 and 488 Physical addresses for non-default VLANs configured on the switch. On the switches covered by this manual, all VLANs use the same MAC address as the Default VLAN. Refer to "Multiple VLAN Considerations" in the "Static LANs (VLANs)" chapter of the *Advanced Traffic Management Guide* for your switch.

Virtual

A Network Out-of-Band Management (OOBM)


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 78](#) shows management connections for a typical switch.

Figure 78 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.

[Table 35](#) summarizes the switch management ports.

Table 35 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 |

Table 35 Switch management ports (continued)


| | In band | Out of band | |
|---------------|--|--|---|
| | Networked | Directly connected | Networked |
| Advantages | Allows centralized management | 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) | 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 79, the switches face the hot aisle of the data center, allowing easy connection to the network ports on the backs of the servers.

Figure 79 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 | N/A | yes | N/A |
| SNMP | yes | yes* | yes | yes |

| Application | Inbound OOBM (server) | Outbound OOBM (client) | Inbound data plane (server) | Outbound data plane (client) |
|-------------|-----------------------|------------------------|-----------------------------|------------------------------|
| TFTP | yes | yes | yes | yes |
| HTTP | yes | N/A | yes | N/A |
| SNTP | N/A | yes | N/A | yes |
| TIMEP | N/A | yes | N/A | yes |
| RADIUS | N/A | yes | N/A | yes |
| TACACS | N/A | yes | N/A | yes |
| DNS** | N/A | yes | N/A | yes |
| Syslog | N/A | yes | N/A | yes |
| Ping | yes*** | yes | yes*** | yes |
| Traceroute | yes*** | yes | yes*** | yes |

^{N/A}N/A = not applicable

* *=SNMP client refers to SNMP traps as they originate from the switch.

** ***=DNS has a limit of two servers—primary and secondary. Either can be configured to use the OOBM interface.

*** ***=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.

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 a management VLAN. If you attempt to enable OOBM when a management VLAN 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” \(page 362\)](#)) have added a `listen` keyword with `oobm` | `data` | `both` options to specify which interfaces are active.

Default value is `both` for all servers.

| | | |
|---------|---|---|
| Telnet: | <code>telnet-server [listen <oobm data both>]</code> | <i>Management and Configuration Guide</i> |
| SSH: | <code>ip ssh [listen <oobm data both>]</code> | <i>Access Security Guide</i> |
| SNMP: | <code>snmp-server [listen <oobm data both>]</code> | <i>Management and Configuration Guide</i> |

| | | |
|-------|---|------------------------------------|
| 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>] | Management and Configuration Guide |
| Traceroute: | traceroute [...] [source <ip-address vlan-id oobm>] | Management and Configuration Guide |

Example:

Figure 80 (page 369) 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 80 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
Switch 41(vlan-1)# end
Switch 41# show oobm
```

**Set up IP address on data network.
Exit back to manager context.
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
Switch 41(oobm)# ip default-gateway 10.255.255.1
Switch 41(oobm)# end
Switch 41# ping 10.1.131.44
10.1.131.44 is alive, time = 19 ms
Switch 41# ping 10.1.131.51
```

**Go to OOBM context and
add IP address and
default gateway.
Exit back to manager context.
Ping server in this rack (on data network).
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.
```

13 Support and other resources

Contacting HP

For worldwide technical support information, see the HP Support Center:

<http://www.hp.com/go/hpsc>

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

Receive, by email, support alerts announcing product support communications, driver updates, software releases, firmware updates, and customer-replaceable component information by signing up at <http://www.hp.com/go/myadvisory>.

To change options for support alerts you already receive, click the **Sign in** link on the right.

Typographic conventions


Table 36 Document conventions

| Convention | Element |
|--|---|
| Blue text: Table 36 (page 371) | <ul style="list-style-type: none">• Cross-reference links and e-mail addresses• A cross reference to the glossary definition of the term in blue text |
| Blue, bold, underlined text | email addresses |
| Blue, underlined text: http://www.hp.com | Website addresses |
| Bold text | <ul style="list-style-type: none">• Keys that are pressed• Text typed into a GUI element, such as a box• GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes |
| <i>Italic</i> text | Text emphasis |
| 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 |



WARNING! Indicates that failure to follow directions could result in bodily harm or death.

 **CAUTION:** Indicates that failure to follow directions could result in damage to equipment or data.

 **IMPORTANT:** Provides clarifying information or specific instructions.

NOTE: Provides additional information.

 **TIP:** Provides helpful hints and shortcuts.

14 Documentation feedback

HP is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (docsfeedback@hp.com). Include the document title and part number, version number, or the URL when submitting your feedback.

Index

Symbols

- 802.1X
 - effect, LLDP, 203
 - LLDP blocked, 177
- 802.1X access control
 - authentication failure, SNMP notification, 160
 - SNMP notification of authentication failure, 160
- =prompt, 351

A

- access
 - manager, 150
 - operator, 150
- ACL
 - debug messages, 322
 - dynamic port ACL, 42
 - gateway fails, 281
 - transferring command files, 244
 - troubleshooting, 280
- ACL, IPv4
 - limit, 222
 - RADIUS-assigned, limit, 222
 - scalability, 222
- ACL, IPv6
 - limit, 222
 - RADIUS-assigned, limit, 222
 - scalability, 222
- address
 - network manager, 144
- address table, port, 265
- address, network manager, 144
- advertise location, 189
- ARP
 - maximums, 222
- ARP protection
 - SNMP notification, 153, 160
- authentication
 - notification messages, 153, 160
 - SNTP, 29
 - SNTP client, 26
- authentication trap, 154
- authorized IP managers
 - SNMP, blocking, 143
- auto MDI/MDI-X
 - configuration, display, 59
 - operation, 59
 - port mode, display, 59
- Auto-10, 108, 110, 118
- auto-recovery
 - configuring, 90
 - disabling, 91
 - specified stack member, 91
- auto-TFTP, 228
 - disable, 229, 230
 - disabled, 228

- download to a redundant management system, 228
- downloading software images, 228
- autonegotiate, 189

B

- bandwidth
 - displaying port utilization, 50
 - guaranteed minimum, 131
- blue locator LED, 357
- boot ROM console, 224
- Bootp
 - effect of no reply, 278
- Bootp/DHCP, LLDP, 183
- broadcast
 - limit, 45
- broadcast mode
 - SNTP, 29
- broadcast storm, 107, 286
- broadcast traffic
 - IPX, 45
 - RIP, 45

C

- CDP, 203, 204, 206
- chassislocate
 - LED, 357
- clear
 - statistics
 - global, 264
 - ports, 264
- CLI
 - context level, 52
- command line interface, 52
- Command syntax
 - alias, 349
 - authoritative, 213
 - auto-tftp, 229
 - autorun, 253
 - boot
 - system flash, 227, 236, 238, 243
 - bootfile-name<filename>, 213
- cdp
 - enable, 206
 - mode pre-standard-voice, 207
 - run, 206
- chassislocate, 358
- clear
 - link-keepalive statistics, 68
 - logging, 313
 - statistics, 264
 - trunk-statistics, 264
- copy
 - command-output tftp, 248
 - command-output usb, 248
 - command-output xmodem, 248
 - config xmodem, 242

- crash-data tftp, 249
- crash-data usb, 249
- crash-data xmodem, 249
- event-log tftp, 248
- event-log usb, 248
- event-log xmodem, 248
- fdr-log tftp, 250
- flash tftp, 240
- flash usb, 240
- flash xmodem, 240
- running-config usb, 244
- show-tech, 344
- startup-config usb, 244
- tftp, 241
- tftp command-file, 244
- tftp copy config, 241
- tftp flash, 227, 239
- tftp show-tech, 242
- usb command-file, 247
- usb flash, 237
- usb startup-config, 244
- xmodem, 242
- xmodem command-file, 246
- xmodem flash, 236
- xmodem startup-config, 243
- debug, 322, 327
 - destination, 329
- default-router <IP-ADDR-STR> [IP-ADDR2 IP-ADDR8], 213
- dhcp-server [enable | disable], 212
- dhcp-server pool < pool-name>, 212
- dns-server <IP-ADDR> [IP-ADDR2 IP-ADDR8], 214
- domain-name <name>, 214
- erase
 - startup-configuration, 350
- external-power-supply , 90
 - auto-recovery, 90
 - power-share, 91
 - power-share allow, 93
 - reset, 91
- fault-finder
 - sensitivity action, 292
- fault-finder broadcast-storm, 55
- ignore-untagged-mac, 209
- interface, 52
 - bandwidth-min output, 133, 134
 - flow-control, 53
 - lACP, 118
 - lACP active, 117
 - link-keepalive, 66
 - link-keepalive vlan, 66
 - mdix-mode, 59
 - monitor, 274
 - name, 61
 - poE-allocate-by, 73
 - poE-llDP-detect, 77, 78
 - power-over-ethernet, 72
 - rate-limit all, 127
- interface speed-duplex, 365
- ip
 - dns domain-name, 355
 - dns server-address priority, 354
 - timeP, 35, 36, 37, 38
 - timeP dhcp, 36
 - timeP manual, 36
- ip address, 366
- ip default-gateway, 366
- jumbo
 - ip-mtu, 140
 - max-frame-size, 139
- kill, 349
- lACP, 109
- lease [DD:HH:MM | infinite], 214
- link, 339
- link-keepalive
 - interval, 66
 - retries, 66
- lldp
 - admin-status, 182
 - config basicTlvEnable, 184
 - config dot1TlvEnable port-vlan-id, 185
 - config dot3TlvEnable macphy_config, 185
 - config ipAddrEnable, 183
 - config medPortLocation, 194
 - config medTlvEnable, 192
 - enable-notification, 182
 - fast-start-count, 190
 - holdtime-multiplier, 180
 - refresh-interval, 179
 - run, 179
- log-numbers, 313
- logging, 321, 330, 331, 332, 333
 - facility, 332
 - filter deny, 210
 - filter enable, 210
 - notify, 316
 - priority-descr, 334
 - severity, 334
 - system-module, 335
- logging origin-id, 317
- mac-count-notify
 - traps, 165
- mac-notify traps, 156
 - aged, 156
- mirror-port, 274
- oobm, 364
 - oobm disable, 364
 - oobm enable, 364
 - oobm interface, 365
 - oobm interface speed-duplex, 365
 - oobm ip address, 366
 - oobm ip default-gateway, 366
- page, 350
- ping, 337
- ping6, 337
- power-over-ethernet
 - pre-std-detect, 72
 - redundancy, 75

- threshold, 76
- process-tracking, 258
- reload, 227, 236, 238, 243
- repeat, 350
- setmib
 - lldpnotificationinterval.0 -i, 182
 - lldpReinitDelay.0 -i, 181
 - lldpTxDelay.0 -i, 180
- setup, 350
- sflow
 - destination, 168
 - polling, 168
 - sampling, 168
- sflow destination oobm, 168
- show, 347
 - bandwidth output, 135
 - boot-history, 346
 - cdp, 205
 - cdp neighbors, 205
 - config, 62, 64
 - debug, 325
 - external-power-supply, 98
 - interface, 62, 63
 - interfaces, 47, 263
 - interfaces brief, 59, 199, 261
 - interfaces config, 59
 - interfaces custom, 48
 - interfaces display, 47
 - interfaces transceiver, 297
 - lacp, 115
 - link-keepalive, 68
 - link-keepalive statistics, 68
 - lldp config, 79, 178
 - lldp info local-device, 197
 - lldp info remote-device, 199
 - lldp stats, 200
 - logging, 312
 - mac-address, 265, 359
 - mac-notify traps, 156
 - management, 21, 34, 261
 - monitor, 273
 - name, 62
 - oobm ip, 367
 - power-over-ethernet, 81, 84
 - power-over-ethernet brief, 82
 - rate-limit all, 128
 - resources, 40
 - sflow agent, 169
 - sflow destination, 170
 - sflow sampling-polling, 170
 - snmp-server, 151, 164
 - snmpv3 enable, 145
 - snmpv3 only, 145
 - snmpv3 restricted-access, 145
 - snmpv3 user, 147
 - sntp, 20
 - spanning-tree, 267
 - system, 256
 - tech, 342
 - tech custom, 242
 - timep, 34
 - trunk-statistics, 264
 - trunks, 114
 - vlans, 137, 138
 - vlans ports, 138
- show cpu, 259
- show cpu process, 259
- show debug, 320
- show fault-finder broadcast-storm, 55
- show oobm, 367
- show oobm arp, 367
- show running-config, 320
- snmp-server
 - community, 152
 - enable traps, 160
 - enable traps link-change, 162
 - enable traps mac-count-notify, 165
 - host, 155
 - host inform, 157
 - response-source, 163
 - trap-source, 163
- snmpv3
 - community, 150
 - enable, 145
 - group, 148
 - notify tagvalue, 158
 - only, 145
 - params user, 159
 - restricted-access, 145
 - targetaddress, 158
 - user, 147
- sntp, 22, 25, 26
 - authentication, 29
 - authentication key-id, 27, 28
 - broadcast, 22, 29
 - server, 22, 23, 39
 - server priority, 22, 23, 25, 28
 - unicast, 23, 29
- task-monitor
 - cpu, 257
- test
 - cable-diagnostics, 301
- fttp, 228
- timesync, 26, 35, 37
 - sntp, 21, 22, 23
 - timep, 35, 36
- traceroute, 339
- trunk, 116
- trunk-load-balance, 125
- vlan
 - jumbo, 139
- write
 - terminal, 342
- communities, SNMP
 - viewing and configuring with the menu, 151
- configuration
 - copying, 241
 - impacts of software download on, 224

- port, 44
- port trunk group, 107
- port, duplex, 52
- port, speed, 52
- restoring factory defaults, 350
- SNMP, 144, 149
- SNMP communities, 151
- traffic mirroring, 255
- transferring, 241
- trap receivers, 154
- usb autorun, 250
- configuration file
 - browsing for troubleshooting, 342
- configuration file, multiple
 - copy from a USB device, 244
 - copy to a USB device, 243
 - copy via tftp, 241
 - copy via Xmodem, 242
- console
 - measuring network activity, 278
 - status and counters menu, 255
 - troubleshooting access problems, 277
- contacting HP, 371
- conventions
 - document, 371
 - text symbols, 371
- copy
 - command output, 248
 - config
 - oobm, 241
 - crash data, 249
 - event log output, 248
 - show tech, 344
 - tftp
 - show-tech, 242
 - tftp show-tech, 242
- CPU utilization, 256
- cpu utilization data, 257
- customizing, show command output, 48

D

- date format, events, 304
- debug
 - acl messages, 322
 - compared to event log, 316
 - destination, logging, 322
 - displaying debug configuration, 325
 - forwarding IPv4 messages, 322
 - lldp messages, 323
 - overview, 316
 - packet messages, 322
 - sending event log messages, 317
 - standard event log messages, 322
 - using CLI session, 321
- debug command
 - all, 322, 328
 - cdp, 322, 328
 - configuring debug/Syslog operation, 324
 - destination, 328
 - destinations, 321, 329
 - event log, 335
 - event log as default, 321
 - event types supported, 317
 - ip, 322
 - ip fib, 322
 - ip ospfv3, 322
 - ip pim, 322
 - ipv6 dhcpv6-client, 323
 - ipv6 dhcpv6-relay, 323
 - ipv6 forwarding, 323
 - ipv6 nd, 323
 - lldp, 329
 - operating notes, 335
 - rip, 328
 - security, 323
 - services, 323, 329
 - show debug, 325
 - snmp, 323, 329
 - support for "debug" severity on Syslog servers, 330, 336
- default settings, 17, 52, 53, 59, 66, 82, 185, 228, 229, 253
 - ping, 337, 338
 - security, 251
 - traceroute, 341
- default trunk type, 113
- DHCP
 - address problems, 278
 - effect of no reply, 278
- DHCP server
 - BootP server, 210
 - configuring lease time, 214
 - DHCP request packets
 - ip pools, 211
 - inform packets
 - authoritative, 211
 - authoritative pools, 211
 - dummy pools, 211
 - ip pools
 - authoritative, 211
 - dynamic pool, 210
 - static pool, 210
- DHCP snooping
 - SNMP notification, 153, 160
- DHCP/Bootp, LLDP, 183
- DHCPv4
 - introduction, 210
- DHCPv4 server
 - configuration commands, 212
 - configure authoritative, 213
 - configuring default router, 213
 - configuring DHCP address pool name, 212
 - enable / disable server, 212
 - specify boot file, 213
- DHCPv6
 - client, 323
 - debug messages, 323
 - diagnostics tools, 336

- browsing the configuration file, [342](#)
- displaying switch operation, [342](#), [344](#)
- ping and link tests, [336](#)
- traceroute, [339](#)
- viewing switch operation, [342](#)

DNS

- configuration, [354](#), [356](#)
- configuration, viewing, [356](#)
- configuring domain name, [214](#)
- DNS-compatible commands, [354](#)
- domain name, fully qualified, [353](#), [356](#)
- event log messages, [357](#)
- Example:, [355](#)
- IPv6 DNS resolution, [353](#)
- operating notes, [357](#)
- ping, [354](#)
- resolver, [352](#)
- resolver operation, [353](#)
- secure management VLAN, [357](#)
- server address, DHCP not used, [357](#)
- server IP address, [353](#), [357](#)
- three entries supported, [354](#)
- traceroute, [354](#)
- VLAN, best route selection, [357](#)

DNS ip servers

- configuring, [214](#)

document

- conventions, [371](#)

documentation

- providing feedback on, [373](#)

dot1TlvEnable, [185](#)

download, [238](#)

- software using TFTP, [224](#)
- switch-to-switch, [238](#)
- TFTP, [225](#)
- troubleshooting, [226](#)
- Xmodem, [235](#)

duplex advertisements, [185](#)

duplex information, displaying, [198](#)

duplicate MAC address, [290](#)

Dyn1, [111](#)

E

edge ports, [127](#)

Emergency Location Id Number, [193](#)

event log

- compared to debug/Syslog operation, [316](#)
- debugging by severity level, [321](#), [330](#)
- debugging by system module, [321](#), [330](#)
- generated by system module, [304](#)
- how to read entries, [303](#)
- listing entries, [313](#)
- losing messages, [303](#)
- navigation, [311](#)
- not affected by debug configuration, [335](#)
- security levels, [155](#)
- sending event log messages as traps, [155](#)
- sending messages to Syslog server, [322](#)
- severity level, [304](#), [334](#)

system module, [335](#)

time format, [304](#)

used for debugging, [321](#)

used for troubleshooting, [303](#)

excessive frames, [142](#)

external power

power allocation

- reducing, [93](#)

external power supply, [85](#)

- see *also* redundant power

F

facility

- logging, [321](#)

factory default configuration

- restoring, [350](#)

failover, locator LED, [358](#)

failover, management module, locator LED, [358](#)

failure, switch software download, [226](#)

fan failure, [291](#)

fault finder

- conditions, [296](#)

- enabling, [293](#)

- sensitivities, [296](#)

- thresholds, [293](#)

fault-finder

- transceiver link-flap, [292](#)

- transceiver sensitivities, [292](#)

- warn and disable, [293](#)

- Web interface, [293](#)

fault-tolerance, [108](#)

fiber optics, monitoring links, [65](#)

filter, source-port

- jumbo VLANs, [141](#)

firmware version, [256](#)

flow control

- constraints, [45](#), [52](#)

- effect on rate-limiting, [130](#)

- global, [52](#)

- global requirement, [45](#)

- jumbo frames, [141](#)

- per-port, [45](#), [52](#)

flow sampling, [143](#)

force option

- power allocation, [92](#)

- using, [92](#)

friendly port names, [60](#)

G

gateway

- routing fails, [281](#)

giant frames, [142](#)

guaranteed minimum bandwidth

- apportioning unallocated bandwidth, [132](#)

- configuration, [133](#)

- displaying current configuration, [135](#)

- impacts of QoS queue configuration, [132](#)

- operation, [131](#)

- outbound queue priority, [132](#)

starving queues, 132

H

help

obtaining, 371

HP

Auto-MDIX feature, 59
subscription service, 371
technical support, 371

HP 640 Redundant/External Power Supply Shelf

external power supply, 85
redundant power, 85

I

IDM

resources, 43

IEEE 802.1d, 286

IEEE P802.1AB/D9, 176

IGMP

host not receiving, 282
not working, 282
statistics, 268

Inbound Telnet Enabled parameter, 277

include-credentials, SNMP, 31

informs

sending to trap receiver, 155
SNMP, 156

IP

address maximums, 222
duplicate address, 278
duplicate address, DHCP network, 278
time server address, 21, 35

IP address

for SNMP management, 143

IP routing

debug messages, 322

ip-option with ping, 338

IPv4

static route, maximum, 222

IPv6

debug dhcpv6 messages, 323
static route, maximum, 222

IPX

broadcast traffic, 45
network number, 261

J

jumbo frames

configuration, 137
excessive inbound, 141
flow control, 141
GVRP operation, 137
management VLAN, 141
maximum size, 137, 139
MTU, 137
port adds and moves, 137
port speed, 137
security concerns, 141
through non-jumbo ports, 141

traffic sources, 137

troubleshooting, 142

VLAN tag, 137

voice VLAN, 141

L

LACP

802.1X not allowed, 120

active, 117

blocked ports, 122

default port operation, 120

described, 110, 118

Dyn1, 111

dynamic, 118

full-duplex required, 108, 118

IGMP, 121

no half-duplex, 122

operation not allowed, 283

overview of port mode settings, 108

passive, 117

removing port from active trunk, 117

restrictions, 120

standby link, 119

status, terms, 120

STP, 121

trunk limit, 118

VLANs, 121

with 802.1X, 121

with port security, 121

Layer-3

scalability, 222

link failures

detecting, 65

link speed, port trunk, 108

link test, 336

link-change traps, 153, 162

LLDP

802.1D-compliant switch, 202

802.1X blocking, 177

802.1X effect, 203

advertisement content, 183

advertisement, mandatory data, 183

advertisement, optional data, 184

advertisements, delay interval, 180

CDP neighbor data, 203

chassis ID, 183

chassis type, 183

clear statistics counters, 200

comparison with CDP data fields, 203

configuration options, 174

configuring optional data, 184

data options, 175

data read options, 176

debug messages, 322, 323

default configuration, 174

DHCP/Bootp operation, 177

display neighbor data, 199

enable/disable, global, 179

general operation, 173

- global counters, 201
- holdtime multiplier, 180
- hub, packet-forwarding, 174
- IEEE P802.1AB/D9, 176
- inconsistent value, 181
- information options, 175
- invalid frames, 201
- IP address advertisement, 177, 202
- IP address subelement, 183
- IP address, DHCP/Bootp, 183
- IP address, options, 183
- IP address, version advertised, 183
- mandatory TLVs, 203
- MIB, 174, 176
- neighbor data remaining, 203
- neighbor data, displaying, 199
- neighbor statistics, 201
- neighbor, maximum, 202
- operating rules, 177
- operation, 173
- outbound packet options, 175
- packet boundaries, 174
- packet dropped, 174
- packet time-to-live, 176
- packet-forwarding, 174
- per-port counters, 201
- port description, 184
- port ID, 183
- port speed, 185
- port trunks, 177
- port type, 183
- refresh interval, 179
- reinitialization delay, 181
- remote management address, 176
- remote manager address, 183
- reset counters, 200
- setmib, delay interval, 180
- setmib, reinit delay, 181
- show commands, 177, 178
- show outbound advertisement, 197
- SNMP notification, 175
- SNMP traps, 175
- spanning-tree blocking, 177
- standards compatibility, 176
- statistics, 200
- statistics, displaying, 200
- system capabilities, 184
- system description, 184
- system name, 184
- Time-to-Live, 174
- time-to-live, 174, 180
- transmission frequency, 174
- transmission interval, change, 179
- transmit and receive, 174
- transmit/receive modes, 174
- trap notice interval, 182
- trap notification, 182
- trap receiver, data change notice, 182
- TTL, 174, 176

- VLAN, untagged, 203
- walkmib, 176
- with PoE, 77
- lldp
 - port vlan ID support, 185
- LLDP-MED
 - displaying speed, 198
 - ELIN, 193
 - enable or disable, 174
 - endpoint support, 189
 - fast start control, 190
 - location data, 193, 194
 - medTlvenable, 192
 - Neighbors MIB, 199
 - Voice over IP, 188
- load balancing, 107, 125
- logging
 - facility, 321
 - neighbor-adjacency, 321
 - priority-desc, 321
 - udp, 332
- logging command, 327
- logical port, 112
- loop, network, 107

M

- MAC address, 256
 - displaying detected devices, 359
 - duplicate, 286, 290
 - port, 359
 - VLAN, 360
- Management Information Base, 143
- management module failover, locator LED, 358
- management port, 362
- management VLAN, 143
 - DNS, 357
- manager access, 150
- manager password
 - SNMP notification, 153, 160
- max frame size, jumbo, 139
- maximums, 223
- MDI/MDI-X
 - configuration, display, 59
 - operation, 59
 - port mode, display, 59
- media type, port trunk, 108
- MIB
 - HP proprietary, 143
 - listing, 143
 - standard, 143
- monitoring
 - links between ports, 65
 - locator LED, 357
- multiple forwarding database, 261, 265
- multiple VLAN, 143

N

- N+1
 - redundant power, 89

- navigation, event log, 312
- network management functions, 144, 150
- network manager address, 144
- network slow, 278
- non-PoE
 - configuration
 - examples, 93
 - redundant power, 89
- notifications
 - authentication messages, 153, 160
 - configuring trap receivers, 154
 - link-change traps, 153
 - network security, 160

O

- OOBM
 - sflow, 168
- oobm
 - address config, 366
 - client commands, 368
 - copy config to remote host, 241
 - default gateway config, 366
 - enable/disable, 364, 365
 - server commands, 367
 - show arp, 367
 - show commands, 367
 - show config, 367
- operating system, 224
- operation not allowed, LACP, 283
- operator access, 150
- OS, 224
 - version, 239
- OSPF
 - debug messages, 322

P

- packet
 - debug messages, 322
- password
 - disables usb autorun, 254
 - SNMP notification, 160
 - SNMP notification for invalid login, 153
- pattern matching, show command output, 347
- pbr
 - debug messages, 323
- ping, 337, 352, 354
 - ip-option, 338
- ping test, 336
- PoE
 - advertisements, 193
 - allocate-by, 71
 - allocation
 - controlling, 73
 - benefit of LLDP-MED, 188
 - changing the threshold, 76
 - configuration
 - multiple switches, 96
 - XPS port, 94
 - configuration options, 70
 - configuring, 72
 - priority, 72
 - detection
 - LLDP TLV advertisement, 77
 - detection status, 83
 - DLC, 77
 - enable or disable operation, 72
 - enabling, disabling ports
 - allocating power using LLDP, 77
 - enabling, disabling redundancy, 75
 - EPS, defined, 70
 - Event Log messages, 106
 - fault, 74
 - IEEE 802.3at std, 77
 - LLDP detection, enabling or disabling, 77, 78
 - lldp negotiation, 78
 - manually configuring power levels, 74
 - MPS
 - absent cnt, 85
 - needed power for PoE+, 71
 - other fault, 84
 - over current cnt, 84
 - overview of status, 82
 - PD support, 71
 - poe-lldp-detect command, 77
 - port priority, 71
 - power denied cnt, 84
 - pre-standard devices
 - enabling support, 72
 - prioritization, 71
 - priority, port, 71
 - RPS, defined, 70
 - setting allocation, 74
 - short cnt, 85
 - slot-id-range option, 77
 - status, 191
 - terminology, 70
 - threshold, power, 75
 - TLVs, 78
 - usage, 71
 - using LLDP, 77
 - viewing
 - LLDP port configuration, 79
 - using LLDP information, 79
 - viewing global power status, 81
 - viewing status
 - all ports, 82
 - specific ports, 84
- PoE configuration
 - implementing, 105
- PoE configurations
 - security features
 - applying, 106
- PoE ports
 - VLANs
 - assigning, 106
- PoE power
 - maximum
 - determining, 86

- PoE traffic
 - priority policies
 - assigning, 106
- PoE+
 - LLDP, 77
- policy enforcement engine
 - described, 42
 - displaying resource usage, 42
- poll interval, 18
- port
 - address table, 265
 - blocked by UDLD, 66
 - configuration, 44
 - configuring UDLD, 66
 - context level, 52
 - counters, 262
 - counters, reset, 262
 - enabling UDLD, 67
 - fiber-optic, 45
 - MAC address, 359, 360
 - management, 362
 - menu access, 45
 - traffic patterns, 262
 - transceiver status, 50
 - trunk, 110
 - utilization, 50
 - CLI, 50
- port configuration, 107
- port names, friendly
 - configuring, 61
 - displaying, 62
 - summary, 60
- port security
 - port trunk restriction, 108
 - trunk restriction, 112
- port trunk, 107, 108
 - bandwidth capacity, 107
 - caution, 107, 112, 117
 - default trunk type, 113
 - enabling UDLD, 67
 - IGMP, 112
 - limit, 107
 - limit, combined, 118
 - link requirements, 108
 - logical port, 112
 - media requirements, 110
 - media type, 108
 - monitor port restrictions, 112
 - nonconsecutive ports, 107
 - port security restriction, 112
 - removing port from static trunk, 116
 - requirements, 111
 - spanning tree protocol, 111
 - static trunk, 111
 - static trunk, overview, 108
 - static/dynamic limit, 118
 - STP, 111
 - STP operation, 111
 - traffic distribution, 111
 - Trk1, 111
 - trunk (non-protocol) option, 110
 - trunk option described, 123
 - types, 110
 - UDLD configuration, 66
 - VLAN, 112
 - VLAN operation, 111
- port trunk group
 - interface access, 107
- port-based access control
 - event log, 283
 - LACP not allowed, 120
 - troubleshooting, 283
- port-utilization and status displays, 50
- power
 - information
 - viewing, 98
 - requirements, 106
- power levels, configuring, 74
- power supplies, 85
 - see *also* PSUs
- power supply settings
 - restoring default, 91
- power-over-ethernet, 70
- power-share option, 92
- ProCurve
 - Auto-MDIX feature, 59
 - HP, URL, 143
- ProCurve Manager
 - reading USB autorun files, 252
 - required for USB autorun, 250
 - security concerns when deleting public community, 144
 - SNMP and network management, 143
- prompt, =, 351
- PSU
 - external and internal combined, 88
- PSUs, 85
 - see *also* power supplies
 - supported, 86
- public SNMP community, 144, 150

Q

- QoS, 42
- Quality of Service
 - queue configuration, 132

R

- RADIUS-assigned ACLs
 - resources, 42
- rate display for ports, 50
- rate-limiting
 - caution, 127
 - configuration, 127
 - displaying configuration, 128
 - edge ports, 127
 - effect of flow control, 130
 - effect on port trunks, 129
 - how measured, 130
 - ICMP, 127

- intended use, 127
- note on testing, 131
- operating notes, 129
- optimum packet size, 130
- per-port only, 127
- purpose, 127
- traffic filters, 130
- redundancy, 75
 - locator LED, 358
- redundant power, 85
 - see also external power supply
 - N+1, 89
 - non-PoE, 89
- reset
 - port counters, 262
- resetting the switch
 - factory default reset, 350
- resource monitor
 - event log, 43
- resource usage
 - displaying, 40
 - insufficient resources, 43
- restricted write access, 150
- RFCs, 143
 - RFC 1493, 143
 - RFC 1515, 143
 - RFC 2737, 176, 177
 - RFC 2863, 176, 177
 - RFC 2922, 176
- RIP
 - broadcast traffic, 45
 - debug messages, 323
- RMON, 143
- RMON groups supported, 167
- router
 - maximum routes, 222
 - OSPF area maximum, 222
 - OSPF interface maximum, 222
 - RIP interface maximum, 222
 - supported routes, 222
- routing
 - gateway fails, 281
 - traceroute, 339
- S**
- scalability, 222
- SCP/SFTP
 - enabling, 230
 - session limit, 233, 234
 - troubleshooting, 234
- secure copy, 230
- secure FTP, 230
- secure management VLAN, DNS, 357
- security
 - enabling network security notifications, 160
 - USB autorun, 253
- Self Test LED
 - behavior during factory default reset, 351
- serial number, 256
- setmib
 - delay interval, 180
 - reinit delay, 181
- severity level
 - event log, 304
 - selecting Event Log messages for debugging, 334
- sFlow, 143
 - configuring via the CLI, 168
 - sampling-polling information, 170
 - show commands, 169
- sflow
 - CLI-owned versus SNMP-owned configurations, 168
 - OoBM, 168
- show
 - custom option, 48
 - displaying specific output, 347
 - pattern matching with, 347
- show cpu, 257
- show debug, 325
- show interfaces
 - dynamic display, 47
- show interfaces display, 346
- show management, 35
- show power-over-ethernet command
 - examples, 102
- show running-config command
 - Example:, 104
- show tech, 242, 342
- slow network, 278
- SNMP, 143
 - ARP protection events, 153
 - authentication notification, 153, 160
 - CLI commands, 150
 - communities, 144, 150
 - configuring with the menu, 151
 - mapping, 149
 - configure, 144
 - configuring security groups, 158
 - configuring SNMPv3 notification, 158
 - configuring SNMPv3 users, 158
 - configuring trap receivers, 154
 - configuring trap receivers, 154
 - DHCP snooping events, 153
 - different versions, 153
 - enabling informs, 156
 - enabling SNMPv3, 158
 - fixed traps, 154
 - invalid password in login, 153
 - IP, 143
 - link-change traps, 153, 162
 - manager password change, 153
 - network security notification, 160
 - notification, LLDP
 - SNMP notification, 175
 - public community, 144, 150
 - supported notifications, 153
 - system thresholds, 154
 - traps, 66, 143, 153
 - walkmib, 361

- well-known traps, 154
- SNMP trap, LLDP, 182
- SNMPv3
 - "public" community access caution, 145
 - access, 144
 - assigning users to groups, 146
 - communities, 149
 - enable command, 145
 - enabling, 144
 - group access levels, 148, 149
 - groups, 148
 - network management problems with snmpv3 only, 145
 - restricted-access option, 145
 - set up, 144
 - users, 144
- SNTP
 - authentication command, 29
 - authentication mode, 27
 - broadcast
 - mode, 29
 - broadcast mode, 17, 21
 - broadcast mode, requirement, 17
 - client authentication, 26
 - configuration, 18
 - disabling, 24
 - display config information, 30
 - display statistics, 30
 - event log messages, 40
 - include-credentials, 31
 - key-id, 27, 28
 - key-value, 27
 - menu interface operation, 39
 - operating modes, 17
 - poll interval, 25
 - priority, 25, 29
 - server priority, 25
 - show authentication, 30
 - trusted key, 28
 - unicast mode, 17, 23, 29
 - unicast time polling, 38
 - unicast, replacing servers, 38
 - viewing, 18, 20
- software, 224
- software image, 224
- software version, 256
- source port filters
 - jumbo VLANs, 141
- spanning tree
 - fast-uplink, troubleshooting, 287
 - problems related to, 286
 - show tech, copy output, 343
 - using with port trunking, 111
- SSH
 - file transfer, 229
 - TACACS exclusion, 232
 - troubleshooting, 234, 287
- standard MIB, 143
- static route, maximum, 222
- statistics
 - clearing, 264
 - SNTP, 30
- subscription service, HP, 371
- switch software, 224
 - download using TFTP, 224
 - download, failure indication, 226
 - download, troubleshooting, 226
 - download, using TFTP, 224
 - software image, 224
 - version, 226, 235
- symbols in text, 371
- Syslog, 316
 - "debug" severity level as default, 336
 - adding priority description, 334
 - compared to event log, 316
 - config friendly descriptions, 333
 - configuring for debugging, 324
 - configuring server address, 321
 - configuring server IP address, 327
 - configuring Syslog servers and debug destinations, 321
 - control-desc, 333
 - displaying Syslog configuration, 325
 - logging command, 327, 329
 - operating notes, 335
 - overview, 316
 - priority-descr, 334
 - sending event log messages, 317
 - server configuration, 330
 - severity, "debug", 330
 - specifying severity level events for debugging, 334
 - specifying system module events for debugging, 335
 - user facility as default, 336
 - using event log for debugging, 321, 330
- syslog message sender
 - display identification, 319
- Syslog messages
 - hostname, 317
- system module
 - selecting event log messages for debugging, 335

T

- TACACS
 - SSH exclusion, 232
- task monitor, 257
- taskusage -d, 257
- taskUsageShow, 257
- technical support
 - HP, 371
- Telnet
 - troubleshooting access, 277
- text symbols, 371
- TFTP
 - auto-TFTP, 228
 - auto-TFTP feature, 228
 - auto-TFTP, disable, 229, 230
 - copy command output, 248
 - copy crash data, 249
 - copy event log output, 248
 - copying a configuration file, 241

- copying software image, 240
- disable, 230
- disabled, 228
- download software using CLI, 227
- downloading software using console, 225
- enable client or server, 228
- enabling client functionality, 228
- enabling server functionality, 228
- switch-to-switch transfer, 238
- troubleshooting download failures, 226
- uploading an ACL command file, 244
- using to download switch software, 224
- threshold setting, 144, 150
- thresholds, SNMP, 154
- time format, events, 304
- time protocol
 - selecting, 18
- Time-to-Live
 - LLDP, 174
- TimeP
 - assignment methods, 17
 - disabling, 37
 - poll interval, 37
 - server address listing, 21, 34
 - show management, 35
 - viewing and configuring, menu, 34
 - viewing, CLI, 34
- timesync, disabling, 37
- TLV advertisement, 185
- TLVs, mandatory, 203
- traceroute, 354
 - blocked route, 342
 - fails, 341
- traffic
 - port, 262
- traffic monitoring, 143, 144, 150
- transceiver
 - error messages, 51
 - fault sensitivities, 292
 - fault-finder, 292
 - fiber-optic, 45
 - flapping, 292
 - link-flap, 292
 - view status, 50
- trap
 - CLI access, 154
 - configuring trap receivers, 154
 - security levels, 155
- trap notification, 182
- trap receiver, 144
 - configuring, 154
 - sending event log messages, 155
 - sending SNMPv2 informs, 155
 - SNMP, 154
 - up to ten supported, 154
- traps, 154
 - arp-protect, 161
 - auth-server-fail, 161
 - dhcp-snooping, 161
 - dynamic-ip-lockdown, 161
 - fixed, 154
 - link-change, 161, 162
 - login-failure-mgr, 161
 - password-change-mgr, 161
 - port-security, 161
 - snmp-authentication, 161
 - threshold, 154
- troubleshooting
 - ACL, 280
 - approaches, 276
 - browsing the configuration file, 342
 - configuring debug destinations, 321
 - console access problems, 277
 - diagnosing unusual network activity, 278
 - diagnostics tools, 336
 - displaying switch operation, 342, 344
 - DNS, 352
 - fast-uplink, 286
 - ping and link tests, 336
 - resource usage, 41
 - restoring factory default configuration, 350
 - spanning tree, 286
 - SSH, 287
 - SSH, SFTP, and SCP Operations, 234
 - switch software download, 226
 - switch won't reboot, shows = prompt, 351
 - traceroute, 354
 - unusual network activity, 278
 - using CLI session, 321
 - using debug and Syslog messaging, 316
 - using the event log, 303
 - viewing switch operation, 342
 - web browser access problems, 277
- trunk, 107
 - L4 load balancing, 125
 - load balancing, 125
 - number supported, 108
- TTL
 - LLDP, 174
- typographic conventions, 371

U

- UDLD
 - changing the keepalive interval, 67
 - changing the keepalive retries, 67
 - configuring for tagged ports, 67
 - enabling on a port, 67
 - operation, 66
 - overview, 65
 - viewing configuration, 68
- UDP
 - logging messages, 332
- undersize frames, 142
- Uni-directional Link Detection, 65
- unicast mode
 - SNTP, 29
- unrestricted write access, 150
- unusual network activity, 278

up time, [256](#)

URL

ProCurve, [143](#)

USB, [254](#)

autorun, [250](#)

AutoRun file, [250](#)

command file, [250](#)

configuring passwords, [254](#)

creating a command file, [251](#)

enabling or disabling, [254](#)

LED indications, [252](#)

report outputs, [252](#)

required software versions, [250](#)

secure-mode, [253](#)

security, [253](#)

troubleshooting, [252](#)

auxiliary port, [236](#), [250](#)

auxiliary port LEDs, [252](#)

copy command output, [248](#)

copy configuration file to/from a USB device, [243](#)

copy crash data, [249](#)

copy event log output, [248](#)

copy software image to a USB device, [240](#)

devices with secure partitions not supported, [237](#)

flash drives must be formatted, [237](#)

supported capabilities, [237](#)

uploading an ACL command file, [247](#)

viewing flash drive contents, [237](#)

users, SNMPv3, [146](#)

utilization, port, [50](#)

V

version, OS, [239](#)

version, switch software, [226](#), [235](#)

view

transceiver status, [50](#)

VLAN

address, [143](#)

configuring UDLD for tagged ports, [67](#)

device not seen, [289](#)

event log entries, [304](#)

IP address maximum, [222](#)

jumbo max frame size, [139](#)

link blocked, [286](#)

MAC address, [360](#)

management and jumbo frames, [141](#)

management VLAN, SNMP block, [143](#)

maximums, [222](#)

multiple, [143](#)

port configuration, [289](#)

secure management VLAN, with DNS, [357](#)

switch software download, [224](#)

tagging broadcast, multicast, and unicast traffic, [289](#)

VoIP

LLDP-MED support, [188](#)

W

walkmib, [176](#), [361](#)

web browser interface

troubleshooting access problems, [277](#)

web site, HP, [143](#)

web sites

HP subscription service, [371](#)

write access, [150](#)

X

Xmodem

copy command output, [248](#)

copy crash data, [249](#)

copy event log output, [248](#)

copying a configuration file, [242](#)

copying a software image, [240](#)

uploading an ACL command file, [246](#)

using to download switch software, [235](#)

XPS

additional PoE power, [86](#)

configuring, [89](#)

enabling and disabling

power, [89](#)